

Capacity Planning of DiffServ Networks with Best-Effort and Expedited Forwarding Traffic

Kehang Wu and Douglas S. Reeves *

Departments of Electrical and Computer Engineering and Computer Science
North Carolina State University
kwu@unity.ncsu.edu, reeves@eos.ncsu.edu

Abstract

For networks providing a specific level of service guarantees, capacity planning is an imperative part of network management. Accurate dimensioning is especially important in DiffServ networks, where no per-flow signaling or control exists.

In this paper, we address the problem of capacity planning for DiffServ networks with only Expedited Forwarding (EF) and best effort (BE) traffic classes. The problem is formulated as an optimization problem, where the total link cost is minimized, subject to the performance constraints of both EF and BE classes. The variables to be determined are the routing of EF traffic, and the discrete capacities of the links.

We show that Lagrangean Relaxation and subgradient optimization methods can be used to solve the problem. Computational results show that the solution quality is verifiably good while the running time remains reasonable on practical-sized networks. This represents the first work for capacity planning of multi-class IP networks with convex performance constraints.

1 Introduction

Capacity planning is the process of designing and dimensioning the networks to meet the expected demands. If networks need to stay ahead of the growth of user demand while still be able to provide a satisfactory service, capacity planning is indispensable. Recent years have witnessed the

*This work is supported by DARPA and AFOSR (under grants F30602-99-1-0540 and F49620-99-1-0264).

spectacular growth of the Internet traffic. But the nature of offering only best effort (BE) service makes capacity planning a straightforward matter[15]. Quality of service (QoS) is the ability of a network element to have some level of assurance that its traffic and service requirements can be satisfied. With the popularizing of e-commerce and new value-added services over IP, like Voice over IP, QoS has become a must. Capacity planning will be an imperative part of IP network management to support various qualities of service.

Differentiated services (DiffServ) is regarded as one of the key components for providing QoS in the Internet [6] [24]. The essence of DiffServ is prioritization. The DiffServ Code Point (DSCP) field in the headers of IP packets is marked at the edge of the network. Routers within the core of the network forward packets using different predefined per-hop behaviors (PHBs), according to their DSCP field. To prepare for the deployment of DiffServ, it is necessary to study the capacity planning problem in the context of multiple class-of-service networks. Since there is no signaling or per-flow control, accurate dimensioning of the network is particularly important for achieving performance guarantees.

The IETF DiffServ working group has standardized two PHBs: Expedited Forwarding (EF) and Assured Forwarding (AF). The EF PHB can be used to build a low loss, low latency, low jitter, assured bandwidth, end-to-end service, such as virtual leased line (VLL), through a DiffServ Domain [29]. Because of the great value of this service, EF PHB is very likely to be the first PHB to be put into action.

Operations Research techniques traditionally were applied to a system only after it had already been in place and well modeled. DiffServ is still under active research. New models and revisions of existing standards are actively being proposed. Performance issues related to the traffic aggregation and the interaction between multiple classes are not fully understood. Those factors impose challenges on the formulation of capacity planning problems. Given the urgency for the need of capacity planning in DiffServ network on the one hand, the difficulty and complexity of the problem

on the other hand, we have to resort to simplifications and approximations.

In this paper, we address the problem of capacity planning for DiffServ networks with only Expedited Forwarding (EF) and best effort (BE) traffic classes. We will study other PHBs in the future. The problem is formulated as an optimization problem, where we jointly select the route for each EF user demand pair, and assign a capacity value for each link to minimize the total link cost, subject to the performance constraints of both EF and BE classes. While the performance constraint of EF traffic is only represented by a bandwidth requirement, the performance constraint of BE classes is characterized by the average delay in each link. Queueing is modeled as M/G/1 strict priority queues. Our intention is to not only define the capacity planning problem for the DiffServ networks and disclose feasible solutions, but also provide helpful insights for capacity planning of other QoS architectures.

Although there is no previous work specifically targeting the dimensioning and routing issues of DiffServ networks, there is a lot of work dealing with design and planning issues of communication networks, especially for connection-oriented networks (see [20] and references therein). Traditionally, network design has been focused on optimizing either the network cost or performance by tuning the network topology, link capacity, and routing strategies; see, for example, [11] [9] [12] [25] [21] [19] [14]. More recently, network survivability has attracted attention as well [17] [5] [22].

The literature focusing on the routing problem, where link capacities are given, is abundant [8] [28] [26]. But because the routing and link dimensioning problems are closely related to each other, it is inappropriate to separate them. Papers where the routing and capacity assignment problems are treated simultaneously include [11] [10] [13][23] [19] [20] [2]. Gerla and Kleirock [11] presented heuristic methods based on the flow deviation algorithm [7]. As pointed out by Gavish [10], the weakness of this approach is that there is no means to evaluate the solution quality. Gavish and Neuman [10] formulated the problem as a non-linear integer programming problem, and proposed a Lagrangean relaxation based approach. The networks studied in [11][10] only include

one traffic class, though. Medhi and Tipper [20] proposed four approaches for reconfigurable ATM networks based on the Virtual Path concept. Even though ATM networks includes multiple classes, Medhi presented a model that assumes the deterministic multiplexing of different virtual paths, which results in linear performance constraints. The novel aspect of our DiffServ Network capacity planning problem is the fact that two traffic classes, with independent behaviors and performance requirements, share the same capacity resource, which results in a complex non-linear performance constraint.

The remainder of this paper is organized as follows. In Section 2, notation and detailed assumptions and models are presented. The problem definition is given in Section 3. Section 4 shows a Lagrangean relaxation of the original problem, and describes the subgradient procedure to solve the resultant dual problem. Section 5 presents some numerical results on the use of the method. The paper is concluded in Section 6.

2 Notation and Models

The following notation will be used throughout the paper:

K set of EF traffic demand node pairs

L set of links in the network

J set of possible candidate paths for an EF demand pair

x_{kj} path routing variable; 1 if demand pair $k \in K$ uses path $j \in J$, 0 otherwise.

$$\forall k \in K, \sum_j x_{kj} = 1$$

δ_j^l link-path indicator; 1 if path j uses link $l \in L$, 0 otherwise

α_k EF traffic demand of node pair $k \in K$

β_{ef}^l total EF traffic demand of link $l \in L$

β_{be}^l average BE traffic load on link $l \in L$

- d_l average delay experienced by BE traffic on link $l \in L$
- u_l number of units of capacity needed on link $l \in L$. u_l is integer
- γ size of a unit of capacity
- C_l cost of a unit of capacity on link $l \in L$
- \tilde{y}, \tilde{y}^2 the first and second moment of packet size, (unit: bit & bit²)

In the DiffServ capacity planning problem, we are given a network, which is defined by a set of link L , candidate paths J , and link-path indicators $\{\delta_j^l\}$. We are also given the EF user demand $\{\alpha_k\}$, and the average BE traffic $\{\beta_{be}^l\}$. The projected EF user demands are in the form of origin-destination (O-D) pairs, while BE user demands are defined link by link. The goal of the capacity planning problem is to find the minimum cost network that satisfies the projected user demands of both EF and BE user classes. The variables here are the path routing variable, $\{x_{kj}\}$, which shows whether O-D pair k will use the candidate path j , and the discrete link capacity, $\{u_l\}$. Non-biforked routing model is used for EF class where the traffic from the single EF demand will follow the same path between the origin and the destination. We assume that the cost is only a linear function of link capacity.

Charny reported in [4] that in order to provide the guaranteed delay service using FIFO, the utilization of EF traffic must be limited to a factor smaller than $1/(H-1)$, where H is the number of hops in the longest path of the network. Other implementations of packet scheduling may improve the upper bound on the EF utilization. In this paper, we assume that the projected EF user demand is much smaller than the BE user demand, so there is no concern about this limit on the EF utilization.

To simplify our problem, we also assume that there is no multiplexing gain for EF traffic. This allows us to have a linear capacity requirement for the aggregation of EF class traffic.

$$\beta_{ef}^l = \sum_{k \in K} \alpha_k \sum_{j \in J} \delta_j^l x_{kj} \quad (1)$$

According to [29], we know that, from the bandwidth requirement perspective, an EF class will

be given at least the requested bandwidth to ensure its performance guaranteed. The performance constraint of EF traffic can be described by:

$$\beta_{ef}^l \leq u_l \gamma \quad (2)$$

The performance of BE traffic is evaluated on a link-by-link basis. The value $\tilde{y}/u_l \gamma$ stands for the average transmission delay. We use $\frac{\tilde{y}}{u_l \gamma}$ as the basis for the delay bound. Let $d_{lmax} = g_l \frac{\tilde{y}}{u_l \gamma}$, where g_l is a parameter defined by the network designer. The value of g_l of course should be greater than 1. We assume that the performance of BE traffic is satisfactory if:

$$d_l \leq d_{lmax} \quad (3)$$

We further assume strict non-preemptive priority queue is used in the routers to support the EF class. A priority queue is considered as the canonical example of an implementation of EF [29]. All EF packets share a single FIFO queue which has the highest priority. All other packets are sent to the second priority queue. Also assume that every router is modeled as a M/G/1 system with Poisson packet arrivals and an arbitrary packet length distribution. Note that the priority queue gives the most preferable treatment to the EF packets possible among all schedulers, it will also have the greatest impact on the performance of BE traffic. Given the same performance constraint of BE traffic, the choice of priority queue implies a larger link capacity requirement.

From [16] by Kleinrock, we obtain the performance constraint for BE traffic:

$$\frac{\tilde{y}}{u_l \gamma} + \frac{\tilde{y}^2}{2\tilde{y}} \frac{\beta_{ef}^l + \beta_{be}^l}{(u_l \gamma - \beta_{ef}^l)(u_l \gamma - \beta_{ef}^l - \beta_{be}^l)} \leq g_l \frac{\tilde{y}}{u_l \gamma} \quad (4)$$

Note that to have a meaningful solution, $u_l \gamma > \beta_{ef}^l + \beta_{be}^l$, which means that the constraint (2) is redundant. With some simplification, (4) yields:

$$u_l \gamma \geq f(\beta_{ef}^l) \quad (5)$$

$$\text{where } f(\beta_{ef}^l) = \frac{2\beta_{ef}^l + \beta_{be}^l + \frac{\tilde{y}^2(\beta_{ef}^l + \beta_{be}^l)}{2(\tilde{y})^2(g_l - 1)} + \sqrt{(2\beta_{ef}^l + \beta_{be}^l + \frac{\tilde{y}^2(\beta_{ef}^l + \beta_{be}^l)}{2(\tilde{y})^2(g_l - 1)})^2 - 4\beta_{be}(\beta_{be} + \beta_{ef})}}{2}$$

Let $\theta = \frac{\widetilde{y}^2}{2(\widetilde{y})^2(g_l-1)}$. It can be shown that:

$$\frac{df(\beta_{ef})}{d\beta_{ef}} = 1 + \frac{\theta}{2} + \frac{1}{4} \frac{2(\theta^2 + 4\theta)\beta_{ef} + (2\theta^2 + 6\theta)\beta_{be}}{\sqrt{(\theta^2 + 4\theta)\beta_{ef}^2 + (2\theta^2 + 6\theta)\beta_{be}\beta_{ef} + (\theta + 1)^2\beta_{be}^2}} > 0 \quad (6)$$

$$\frac{d^2f(\beta_{ef})}{d\beta_{ef}^2} = \frac{2\theta\beta_{be}^2}{[(\theta^2 + 4\theta)\beta_{ef}^2 + (2\theta^2 + 6\theta)\beta_{be}\beta_{ef} + (\theta + 1)^2\beta_{be}^2]^{\frac{3}{2}}} > 0 \quad (7)$$

which shows that $f(\beta_{ef}^l)$ is a convex function.

3 Problem Formulation

The formal problem definition is presented below. Our goal is to minimize the total costs.

$$\min(\sum_{l \in L} c_l u_l \gamma) \quad (8)$$

Subject to the constraint:

$$u_l \gamma \geq f(\beta_{ef}^l), \quad \forall l \in L \quad (9)$$

$$\text{where } \beta_{ef}^l = \sum_{k \in K} \alpha_k \sum_{j \in J} \delta_j^l x_{kj} \quad (10)$$

$$u_l \geq 0 \text{ and integer, } \forall l \in L \quad (11)$$

$$\sum_{j \in J} x_{kj} = 1, \quad \forall k \in K \quad (12)$$

$$x_{kj} = 0/1, \quad \forall j \in J, k \in K \quad (13)$$

Constraint (9) ensures the performance of EF and BE traffic. (11) imposes a discrete constraint on the link capacities. (12) and (13) ensure that all traffic from one EF O-D pair will follow one single path.

The fact that the function $f(\beta_{ef}^l)$ is increasing in β_{ef}^l enables us to rewrite the constraint (10).

The resulting problem is:

$$\min(\sum_{l \in L} c_l u_l \gamma) \quad (14)$$

$$\beta_{ef}^l \geq \sum_{k \in K} \alpha_k \sum_{j \in J} \delta_j^l x_{kj} \quad (15)$$

and (9)(11)(12)(13).

We refer to the problem defined by (14,9,15,11,12,13) as problem (P) in the rest of this paper.

As can be seen from the above problem formulation, problem (P) is a non-linear integer programming problem, which is difficult in general. Other than the non-linear constraint (9), problem (P) is identical to the multiple choice multiconstrained knapsack problem [27], which is known to be NP-complete. The non-linear constraint (9) leads to a more difficult problem.

4 Lagrangean Relaxation

Lagrangean Relaxation is a common technique for multicommodity flow problems [1]. It has been successfully applied to the capacity planning and routing problems [9] [10] [2] [19] [20]. We describe its use for our problem in this section.

Let $\lambda = (\lambda_l)$ be the dual multiplier associated with the constraint (15). Then the Lagrangean can be expressed as

$$\begin{aligned} L(x, u, \lambda) &= \sum_{l \in L} c_l u_l \gamma + \sum_{l \in L} \lambda_l (\beta_{ef}^l - \sum_{k \in K} \alpha_k \sum_{j \in J} \delta_j^l x_{kj}) \\ &= \sum_{l \in L} (c_l u_l \gamma + \lambda_l \beta_{ef}^l) + \sum_{k \in K} \sum_{l \in L} -\alpha_k \lambda_l \sum_{j \in J} \delta_j^l x_{kj} \end{aligned} \quad (16)$$

The Lagrangean dual problem (D) is then:

$$\max_{\lambda \leq 0} h(\lambda) \quad (17)$$

$$\text{where } h(\lambda) = \min_{x, u} L(x, u, \lambda) \quad (18)$$

Subject to the constraint:

$$u_l \gamma \geq f(\beta_{ef}^l), \quad \forall l \in L \quad (19)$$

$$u_l \geq 0 \text{ and integer}, \quad \forall l \in L \quad (20)$$

$$\sum_{j \in J} x_{kj} = 1, \quad \forall k \in K \quad (21)$$

$$x_{kj} = 0/1, \quad \forall j \in J, k \in K \quad (22)$$

4.1 Solving the Lagrangean Dual Problem (D)

For a given λ , the Lagrangean is separable in x and u . (18) is reduced to solving two independent subproblems,

$$\min_{x,u} L(x, u, \lambda) = \min_u L_1(u, \lambda) + \min_x L_2(x, \lambda) \quad (23)$$

Subproblem (D1):

$$\min_u L_1(u, \lambda) = \min_u \left\{ \sum_{l \in L} (c_l u_l \gamma + \lambda_l \beta_{ef}^l) \right\} \quad (24)$$

Subject to the constraint:

$$u_l \gamma \geq f(\beta_{ef}^l), \quad \forall l \in L \quad (25)$$

$$u_l \geq 0 \text{ and integer}, \quad \forall l \in L \quad (26)$$

Subproblem (D2):

$$\min_x L_2(x, \lambda) = \min_x \left(\sum_{k \in K} \sum_{l \in L} -\alpha_k \lambda_l \sum_{j \in J} \delta_j^l x_{kj} \right) \quad (27)$$

Subject to the constraint:

$$\sum_{j \in J} x_{kj} = 1, \quad \forall k \in K \quad (28)$$

$$x_{kj} = 0/1, \quad \forall j \in J, k \in K \quad (29)$$

Subproblem (D1) can be separated into L problems, one for each link. The problem for each link l can be expressed as:

$$\min_u (c_l u_l \gamma + \lambda_l \beta_{ef}^l) \quad (30)$$

Subject to the constraint:

$$u_l \gamma \geq f(\beta_{ef}^l) \quad (31)$$

$$u_l \geq 0 \text{ and integer,} \quad (32)$$

Because $f(\beta_{ef}^l)$ is increasing in β_{ef}^l and $\lambda_l \leq 0$, the problem (30) can be simply rewritten as:

$$\min_u (c_l u_l \gamma + \lambda_l f^{-1}(c_l u_l)), \quad u_l \geq 0 \text{ and integer} \quad (33)$$

The value of u_l^* , which satisfies the equation (34) below, can be obtained numerically through any one-dimensional optimization methods, such as Newton's method [3].

$$\frac{dz(u_l)}{du_l} = 0, \quad u_l \geq 0 \quad (34)$$

$$\text{where } z(u_l) = c_l u_l \gamma + \lambda_l f^{-1}(c_l u_l)$$

Let u_{l1}^* and u_{l2}^* be the two non-negative integers closest to u_l^* . We know that the solution of u_l to the problem (33) is either u_{l1}^* or u_{l2}^* . Among them, the one that minimizes (33) is picked as the final solution.

Subproblem (D2) can be separated into K subproblems, one for each O-D pair. The problem for each O-D k is:

$$\min_x \left(\sum_{l \in L} -\alpha_k \lambda_l \sum_{j \in J} \delta_j^l x_{kj} \right) \quad (35)$$

Subject to the constraint:

$$\sum_{j \in J} x_{kj} = 1 \quad (36)$$

$$x_{kj} = 0/1, \quad \forall j \in J \quad (37)$$

The solution of subproblem (D2) is then easily obtained by setting $x_{kj^*} = 1$ for j^* satisfying:

$$P(j^*) = \min_j (P(j)) \quad (38)$$

where $P(j) = \sum_{l \in L} -\alpha_k \lambda_l \delta_j^l x_{kj}$

4.2 Subgradient Optimization Procedures

We use the subgradient method to search for optimal multipliers λ . (For the detailed description of the subgradient method and choice of parameters, please refer to the book [1] by Ahuja, Magnanti, and Orlin.)

For a given initial λ , once we solve the problem (D), a dual subgradient is computed as follows:

$$\omega_l = \beta_{ef}^l - \sum_{k \in K} \alpha_k \sum_{j \in J} \delta_j^l x_{kj}, \forall l \in L \quad (39)$$

The subsequent values of the Lagrangean multipliers is updated:

$$\lambda_l \leftarrow \min(0, \lambda_l + t\omega_l), \forall l \in L \quad (40)$$

where the step size, θ , is defined by:

$$t = \phi \frac{h^*(u) - h(u)}{\|\omega\|^2} \quad (41)$$

where $h^*(u)$ is the value of the best feasible solution found so far, and ϕ is a scalar between 0 and 2. ϕ is set to 2 initially in our study and is halved if the solution does not improve in 10 iterations.

At each iteration, the solution of $\{x_l\}$ for the primal problem (P) can be generated from the solution of subproblem (D2). The value of $\{u_l\}$ can be computed according to (9). Consequently, the primary objective function can be obtained. As the iteration proceeds, we store the best solution found so far for the primal problem (P). In this way, we are always able to obtain a feasible solution.

The benefit of Lagrangean optimization procedures is that the solution of dual problem provides the lower bound to the primal problem. Therefore, the solution quality can be assessed by the

duality gap, which is the difference between the solutions of problem (P) and problem (D). Note that because the duality gap is always no smaller than the actual difference between the obtained feasible solution and the optimal solution, it is a more conservative estimation of the solution quality.

5 Computational Results

In this section, we present numerical results based on experimentation. We wished to examine the performance of the algorithm presented for a wide selection of realistic-sized problems. The program is implemented in C and the computational work is performed on a Pentium III 1.13GHz PC with 256M memory, running the Windows 2000 operating system. The implementation is tested for problems up to 150 nodes, 400 links, and 4000 EF O-D pairs.

The network topologies used in the experiment are generated using Georgia Tech Internetwork Topology Models(GT-IMT) [31]. The location of originations and destinations are randomly generated. The candidate paths are calculated using Yen's K-shortest path algorithm [30]. For each O-D pair, 10 candidate paths are generated. Link cost is set to be proportional to its length. If not specified, EF demands are randomly generated with a uniform distribution from 0 Mbps to 10 Mbps, while the average BE traffic load of each link is also uniformly distributed from 30 Mbps to 100 Mbps. The link unit γ is set to be 45 Mbps. We use $\tilde{y} = 4396$ (bits) and $\tilde{y}^2 = 22790170$ (bits²) for all the test cases. They are calculated based on a traffic trace from NLANR Passive Measurement and Analysis project [18].

The objective of our experiment is to evaluate the solution quality and running time of the algorithm. We test the algorithm on 8 different sizes of networks. For each network size, we generate 30 different topologies with the same number of nodes, links, and O-D pairs. Some details of the network topologies are listed in Table 1.

The solution quality is calculated as the percentage difference between the solution of the primal

Table 1: Network topology information

<i>node number</i>	5	10	20	30
<i>link number</i>	14	30	62	86
<i>O-D number</i>	15	50	200	500
<i>Running time (sec)</i>	(0.35, 0.51)	(2.48, 3.69)	(10.29, 12.01)	(25.82, 28.53)
<i>Solution Quality (%)</i>	(0.2,3.7)	(0,5.4)	(0,4.9)	(0.3,4.0)

<i>node number</i>	50	70	100	150
<i>link number</i>	134	188	276	400
<i>O-D number</i>	1000	1500	2000	4000
<i>Running time (sec)</i>	(53.6, 59.0)	(73.8, 81.7)	(103.2, 117.5)	(165.4, 193.3)
<i>Solution Quality (%)</i>	(0.2, 5.8)	(0, 4.8)	(0, 3.9)	(0, 3.1)

problem and the dual problem.

$$\text{Solution Quality} = \frac{s_p - s_d}{s_d} \tag{42}$$

where s_p and s_d are the solutions of primal problem and dual problem respectively.

Table 1 shows the running time and solution quality with 90% confidence interval.

In all 240 test cases, the algorithm converges without difficulty. It is easy to see from the table that the Lagrangean Relaxation together with subgradient method produce reasonable result as the duality gap is bounded by no more than 6%. Given the large number of networks being tested, we believe that the solution should remain good quality for other size of networks.

Because capacity planning is usually performed in the time scale of weeks to months, the running time of the algorithm is not a critical factor. But it is still desirable to know how the running time scales up with respect to the network size. As can be seen from Figure 1, the running time goes up linearly with respect to the network size in our test cases. It is fair to predict that the running time of the algorithm will stay reasonable for larger size of networks.

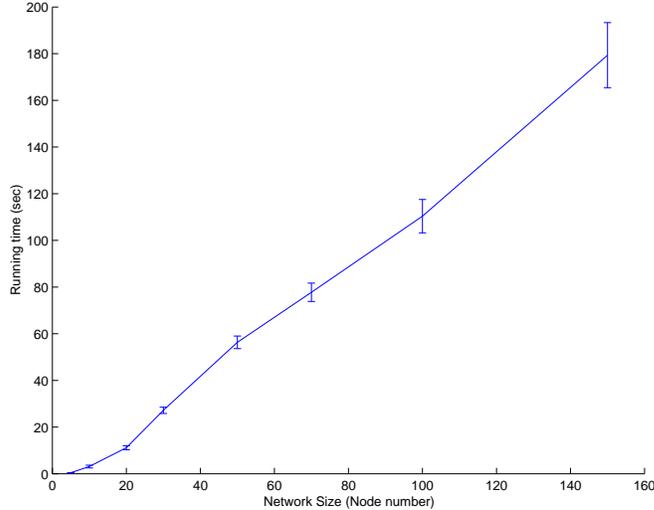


Figure 1: Running time vs. Network Size

6 Conclusions and Future Directions

Services based on the EF PHB are likely to be deployed in the near future. In this paper, we addressed the problem of link dimensioning and routing for DiffServ networks with only EF and BE traffic. We formulate the problem as an optimization problem, where the total link cost is minimized, subject to the performance constraints of both EF and BE classes. The requirement of the EF traffic class only incurs linear link capacity constraints, while the performance guarantee of BE traffic results in nonlinear constraints.

We presented a Lagrangean Relaxation-based method to decompose the original problem. A subgradient method is used to find the optimal Lagrangean multiplier. We investigated experimentally the solution quality and running time of this approach. The results from our experiments indicate that our method is capable of reaching within 6% of the optimum and the running time goes up linearly with the network size.

This paper presents a preliminary investigation of the capacity planning issue for DiffServ networks. The novelty of the problem presented in this paper is that it involves two traffic classes, EF

and BE, which have totally different forms of performance requirements. The problem formulation and solution approaches may be extended to other network architectures and traffic classes.

There is opportunity to extend this work in several directions. We will study types of traffic in addition to EF and BE, such as the Assured Forwarding (AF) class, and the interactions between these classes. We are also investigating how to maintain optimality with minimum change to network design and path selection when traffic loads change incrementally, or failures occur.

Acknowledgement

The authors wish to thank Dr. Cole Smith for his helpful comments and suggestions.

References

- [1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows, Theory, Algorithms and Applications*. Prentice Hall, 1993.
- [2] A. Amiri. A System for the Design of Packet-Switched Communication Networks with Economic Tradeoffs. *Computer Communications*, 21(18):1670–1680, Dec. 1998.
- [3] M. Avriel. *Nonlinear Programming*. Prentice Hall, 1976.
- [4] A. Charny and J.-Y. Le Boudec. Delay Bounds in Network with Aggregate Scheduling. In *First International Workshop on Quality of Future Internet Service*, Berlin, Germany, Sep. 2000.
- [5] O. Crochat and J.Y. Le Boudec. Design Protection for WDM Optical Networks. *IEEE Journal on Selected Areas in Communications*, 16(7):1158–1165, Sep. 1998.
- [6] S. Blake et al. An Architecture for Differentiated Service, IETF RFC 2474. Dec. 1998.

- [7] L. Fratta, M. Gerla, and L. Kleinrock. The Flow Deviation Method: An Approach to Store-and-Forward Communication Network Design. *Networks*, 3(2):97–133, Mar. 1973.
- [8] B. Gavish and S. Hantler. An Algorithm for Optimal Route Selection in SNA Networks. *IEEE Transactions on Communications*, COM-31(10):1154–1161, Oct. 1983.
- [9] B. Gavish and I. Neuman. Capacity and Flow Assignment in Large Computer Networks. In *IEEE Infocom '86*, pages 275–284. IEEE, Apr. 1986.
- [10] B. Gavish and I. Neuman. A System for Routing and Capacity Assignment in Computer Networks. *IEEE Transaction on Communications*, 37(4):360–366, Apr. 1989.
- [11] M. Gerla and L. Kleinrock. On the topological design of distributed computer networks. *IEEE Transactions on Communications*, COM-25(1):48–60, Jan. 1977.
- [12] M. Gerla, J. Monteiro, and R. Pazos. Topology Design and Bandwidth Allocation in ATM Nets. *IEEE Journal on Selected Areas in Communications*, 7(8):1253–1262, Oct. 1989.
- [13] H. Gersht and R. Weihmayer. Joint Optimization of Data Network Design and Facility Location. *IEEE Journal on Selected Areas in Communications*, 8(9):149–152, Dec. 1990.
- [14] M. Girish, B. Zhou, and J. Hu. Formulation of the Traffic Engineering Problems in MPLS Based IP Networks. In *Fifth IEEE Symposium on Computers and Communications*, pages 214–219, Jul. 2000.
- [15] S. Keshav. *An Engineering Approach to Computer Networking*. Addison-Wesley, 1997.
- [16] L. Kleinrock. *Queueing Systems, Volume II: Computer Application*. Wiley Interscience, 1976.
- [17] C.L. Monma M. Grottschel and M. Stoer. *Network Models, Handbooks in Operations Research and Management Science*, volume 7, chapter 10, pages 617–672. North-Holland, Amsterdam, 1995.

- [18] NLANR Passive Measurement and Analysis project. <http://pma.nlanr.net/pma/>.
- [19] D. Medhi. Multi-Hour, Multi-Traffic Class Network Design for Virtual Path-based Dynamically Reconfigurable Wide-Area ATM Networks. *IEEE/ACM Transactions on Networking*, 3(6):809–818, Dec. 1995.
- [20] D. Medhi and D. Tipper. Some Approaches to Solving a Multihour Broadband Network Capacity Design Problem with Single-path Routing. *Telecommunication Systems*, 13(2-4):269–291, 2000.
- [21] D. Mitra, J. Morrison, and K. Ramakrishnan. ATM Network Design and Optimization: A Multi-rate Loss Network Framework. *IEEE/ACM Transactions on Networking*, 4(4):531–543, Aug. 1996.
- [22] E. Modiano and A. Narula-Tam. Survivable Routing of Logical Topologies in WDM Networks. In *IEEE Infocom '01*, volume 1, pages 348–357, Apr. 2001.
- [23] T. Ng and D. Hoang. Joint Optimization of Capacity and Flow Assignment in a Packet-switched Communication Network. *IEEE Transaction on Communications*, COM-35(2):202–209, Feb. 1987.
- [24] K. Nichols, V. Jacobson, and L. Zhang. A Two-bit Differentiated Services Architecture for the Internet, IETF RFC 2638. Jul. 1999.
- [25] I.Cidon O.Gerstel and S.Zaks. The Layout of Virtual Paths in ATM Networks. *IEEE/ACM Transactions on Networking*, 4(6):873–884, Dec. 1996.
- [26] H. Pirkul and A. Amiri. Routing in Packet-switched Communication Networks. *Computer Communications*, 17(5):307–316, May 1994.
- [27] P. Sinha and A.A. Zoltners. The Multiple Choice Knapsack Problem. *Operation Research*, 27(3):503–515, May-Jun. 1979.

- [28] D. Tcha and K. Maruyama. On the Selection of Primary Paths for a Communication Network. *Computer Networks and ISDN Systems*, 9(4):257–265, Arp. 1985.
- [29] V. Jacobson, K. Nichols, and K. Poduri. An Expedited Forwarding PHB, IETF RFC 2598. Jun. 1999.
- [30] J.Y. Yen. Finding the K Shortest Loopless Paths in a Network. *Management Science*, 17(11):712–716, Jul. 1971.
- [31] E.W. Zegura, K. Calvert, and M. Jeff Donahoo. A Quantitative Comparison of Graph-based Models for Internet Topology. *IEEE/ACM Transactions on Networks*, 5(6):770–783, Dec. 1997.