

Differentiated Static Resource Allocation in WDM Networks

James Yiming Zhang¹, Jing Wu², Gregor v. Bochmann¹ and Michel Savoie²

1: University of Ottawa, Ottawa, Ontario, Canada, K1N 6N5; {yizhang, bochmann}@site.uottawa.ca

2: Communications Research Centre Canada, Ottawa, Ontario, Canada, K2H 8S2; {jing.wu, michel.savoie}@crc.gc.ca

Abstract - We present a study on the static resource allocation in lightpath routed WDM networks, where each request is associated with a service grade. The goal is to maintain certain acceptance ratios for the requests of all grades, as well as to minimize the resource consumption. We propose a model of static Grade-of-Service (GoS) differentiation as minimizing the total rejection and cost penalty. Then, we use the Lagrangian relaxation and subgradient methods to solve the problem. The results of using static GoS differentiation are presented.

I. INTRODUCTION

Grade-of-Service (GoS) is important in the design of Wavelength Division Multiplexing (WDM) networks, since optical networks serve an increasing number of services, each having different requirements. In this paper, we focus on point-to-point lightpath routed WDM networks. Existing static Routing and Wavelength Assignment (RWA) algorithms [1] assume all requests have the same priority (grade). There is no differentiated service grade. However, with the network evolution from single to multiple services, it is important to provide a controlled GoS.

The challenge in providing a controlled GoS is that certain service levels should be maintained for requests of every grade. Although a high-grade request should have a better chance to be accepted, by no means a low-grade request can only be accepted after all high-grade requests are accepted. Requests of different grades generally share the same pool of resources. After allocating resources for all high-grade requests, resources may be insufficient to maintain the service level of the low-grade. In contrast, rejecting a small set of high-grade requests could make critical resources available for low-grade requests, thus service levels are maintained for every grade. As such, service differentiation trade-offs must be studied. Specifically, the selection and the resource allocation of the accepted requests, and their impact on the design objectives should be investigated.

Static GoS differentiation uses distinct mechanisms than dynamic GoS differentiation. We study the static GoS differentiation, where all requests must be handled together as a whole. None of the existing dynamic GoS differentiation mechanisms [2-5] can be readily adapted for the static case.

We model the static GoS differentiation using a penalty/price based optimization formulation. Assigning a high rejection penalty to a request makes the request less likely to be rejected than others. By assigning a proper relative rejection penalty for each request, a desired GoS differentiation can be achieved. In our previous work [6], we

optimized static resource allocation for a single-grade of service by maximizing the revenue, which was also modeled as minimizing the rejection penalty. However, we did not discriminate lightpaths using different network resources, and did not study the service differentiation. Here, we propose a new formulation considering both GoS differentiation and resource consumption.

This paper is organized as follows: In Section 2, the network model and assumptions are summarized. In Section 3, we propose a formulation of static GoS differentiation, followed by a solution based on the Lagrangian relaxation and subgradient methods presented in Section 4. In Section 5, we present the results of using static GoS differentiation. Conclusions are given in Section 6.

II. NETWORK MODEL AND ASSUMPTIONS

We consider a WDM mesh network of N nodes interconnected by E fibres. Each fibre has W bi-directional Wavelength Channels (WCs). E is the set of all fibre links in the network. V is the set of all nodes in the network. L is the set of all lightpath requests. More than one lightpath can be set up between a node pair. Using a wavelength converter installed at an intermediate node, two lightpaths of different wavelengths (colors) can be chained together.

We use a model of a wavelength converter with a limited conversion degree. An input lightpath of wavelength c at node i can be converted to any wavelength in the set $I_i(c)$ by using a converter of a certain type. ν is the degree of wavelength conversion, defined as the number of possible output wavelengths of a wavelength converter for a given input wavelength. If a given input lightpath of wavelength c needs conversion, such a conversion is only possible when a converter of index c is available.

Our traffic model assumes lightpath requests are given, and are fixed over time. If multiple lightpaths are set up between a node pair, they are not restricted to take the same route. The acceptance and routing of requests between the same node pair but in opposite direction are independent.

III. FORMULATION OF STATIC GRADE-OF-SERVICE DIFFERENTIATION

We formulate the static GoS differentiation as a minimizing-penalty problem, where the rejection of requests, and the use of resources are penalized. When a request is rejected, certain potential revenue is lost. Thus, the rejection penalty is the amount of its potential revenue. When a request is accepted, the resource consumption penalty is the cost of the resources that are used by the lightpath provisioned for the

request. In this way, our objective essentially becomes the profit maximization.

Our objective function is formulated as:

$$\min_{A, \Delta, \Phi} \{J\}, \quad \text{with } J \equiv \sum_{s_{sdn} \in L} [(1 - \alpha_{sdn})P_{sdn} + \alpha_{sdn}C_{sdn}] \quad (1)$$

where s_{sdn} represents a request or a lightpath provisioned for the request. α_{sdn} is a binary integer variable representing the admission status of s_{sdn} . $\alpha_{sdn} = 1$, if s_{sdn} is admitted; otherwise, $\alpha_{sdn} = 0$. P_{sdn} is the penalty for rejecting s_{sdn} . C_{sdn} is the cost of the resources used by s_{sdn} . J is the primal function of the original problem. The design variables are: A , Δ and Φ :

A is the variable set $\{\alpha_{sdn}\}$, representing the admission status of all requests.

Δ is the variable set $\{\Delta_{sdn}\}$, representing the wavelength assignment for all lightpaths. Δ_{sdn} is the variable set $\{\delta_{ij\lambda}^{sdn}\}$ for a given s_{sdn} , which represents the wavelength assignment at all links for s_{sdn} . $\delta_{ij\lambda}^{sdn}$ is a binary integer variable representing the use of $w_{ij\lambda}$ by s_{sdn} . $\delta_{ij\lambda}^{sdn} = 1$, if $w_{ij\lambda}$ is used by s_{sdn} ; otherwise, $\delta_{ij\lambda}^{sdn} = 0$. $w_{ij\lambda}$ is the wavelength channel λ on e_{ij} , $0 < \lambda \leq W$. e_{ij} is the fibre link between node pair (i, j) , $e_{ij} \in E$.

Φ is the variable set $\{\phi_{sdn}\}$, representing the converter assignment for all lightpaths. Φ_{sdn} is the variable set $\{\phi_{i,ab}^{sdn}\}$ for a given s_{sdn} , which represents the converter assignment at all intermediate nodes of s_{sdn} . $\phi_{i,ab}^{sdn}$ is a binary integer variable representing the use of a wavelength converter by s_{sdn} at node i to convert an incoming lightpath of wavelength a to an outgoing lightpath of wavelength b . $\phi_{i,ab}^{sdn} = 1$, if s_{sdn} uses such a converter; otherwise, $\phi_{i,ab}^{sdn} = 0$.

As an example, we define C_{sdn} as the cost of using converters ($c_{i\lambda}$ is the cost of using a wavelength converter of index λ at node i) and WCs ($d_{ij\lambda}$ is the cost of using $w_{ij\lambda}$ on e_{ij}):

$$C_{sdn} = \sum_{e_{ij} \in E, 0 < \lambda \leq W} d_{ij\lambda} \delta_{ij\lambda}^{sdn} + \sum_{i \in V, 0 < \lambda \leq W} \sum_{0 < a \leq W} c_{i\lambda} \phi_{i,\lambda a}^{sdn} \quad \forall s_{sdn} \in L \quad (2)$$

The optimization is subject to the following constraints:

1) Lightpath flow continuity constraint

$$\sum_{j \in V, 0 < c \leq W} \delta_{ijc}^{sdn} - \sum_{j \in V, 0 < c \leq W} \delta_{jic}^{sdn} = \begin{cases} \alpha_{sdn} & \text{if } i = s \\ -\alpha_{sdn} & \text{if } i = d \\ 0 & \text{otherwise} \end{cases} \quad \forall s_{sdn} \in L \quad (3)$$

If a lightpath is admitted, the lightpath must be continuous along its path, and terminates at its two end nodes.

2) Exclusive WC usage constraint

$$\sum_{s_{sdn} \in L} \delta_{ij\lambda}^{sdn} \leq 1 \quad \forall e_{ij} \in E, 0 < \lambda \leq W \quad (4)$$

A WC may be used by no more than one lightpath.

3) Limited wavelength conversion degree constraint

$$\delta_{ij\lambda}^{sdn} = \sum_{y \in V, w_{jy\lambda} \in I_j(\lambda)} \delta_{jyx}^{sdn}$$

$$\forall s_{sdn} \in L, 0 < \lambda \leq W, \forall e_{ij} \in E, i \neq s, j \neq d \quad (5)$$

A lightpath may only be converted to one of the wavelengths that are allowed by the converters.

4) Converter amount constraint

$$\sum_{s_{sdn} \in L} \sum_{0 < a \leq W} \phi_{i,\lambda a}^{sdn} \leq F_{i\lambda} \quad \forall i \in V, 0 < \lambda \leq W \quad (6)$$

The total number of the incoming lightpaths of wavelength λ at node i that use conversion is limited by the number of installed converters. $F_{i\lambda}$ is the number of wavelength converters of index λ at node i .

IV. A SOLUTION BASED ON THE LAGRANGIAN RELAXATION AND SUBGRADIENT METHODS

The complexity of the formulated problem is very high. For a network with N nodes, E links, W wavelengths and L requests, the problem defined in the previous section contains $|A| + |\Delta| + |\Phi| = L + E \times W \times L + N \times W$ binary integer variables, where $|\bullet|$ denotes the number of elements in the set. For example, a small network with $E=21$, $N=14$, $W=20$ and $L=268$, has 113,108 variables in total. Finding the exact optimum for a problem of this size is hardly possible for the computing facilities available today. Therefore, finding a sub-optimal solution within a reasonable computation time is a practical choice, while knowing the proximity of the sub-optimal solution to the real optimum will be an additional advantage. In this paper, we develop a solution method that applies the Lagrangian Relaxation (LR) and subgradient methods. Our method can find a sub-optimal solution for fairly large networks, and meanwhile the proximity of the sub-optimal solution to the real optimum can be evaluated by a bound. In this paper, we will apply our method to a network with 22 nodes, 35 links, 32 WCs, and 352 requests, which leads to 395,296 design variables in total. The results for 14-node and 28-node networks are not presented, but they demonstrate similar trends.

A. Decomposing the Problem Based on Lagrangian Relaxation

The LR method is used to derive the Dual Problem (DP) of the original problem. The resource constraints are relaxed so that the resource allocation to each lightpath becomes independent. We relax the exclusive WC usage constraint (4) and converter amount constraint (6). Accordingly, additional elements are added into the primary function J by using the corresponding Lagrange multipliers $\xi_{ij\lambda}$ and $\pi_{i\lambda}$. This leads to the following Lagrangian dual problem:

$$\max_{\xi, \pi \geq 0} (q) \leq \min_{A, \Delta, \Phi} \left\{ \sum_{s_{sdn} \in L} [(1 - \alpha_{sdn})P_{sdn} \right.$$

$$\begin{aligned}
& + \alpha_{sdn} \left(\sum_{e_{ij} \in E} \sum_{0 < \lambda \leq W} d_{ij\lambda} \delta_{ij\lambda}^{sdn} + \sum_{i \in V} \sum_{0 < \lambda \leq W} \sum_{0 < a \leq W} c_{i\lambda} \phi_{i,\lambda a}^{sdn} \right) \\
& + \sum_{e_{ij} \in E} \sum_{0 < \lambda \leq W} \xi_{ij\lambda} \left(\sum_{s_{sdn} \in L} \delta_{ij\lambda}^{sdn} - 1 \right) \\
& + \sum_{i \in V} \sum_{0 < \lambda \leq W} \pi_{i\lambda} \left(\sum_{s_{sdn} \in L} \sum_{0 < a \leq W} \phi_{i,\lambda a}^{sdn} - F_{i\lambda} \right) \} \quad (7)
\end{aligned}$$

subject to the constraints (3) and (5). q is the dual function $q(\xi, \pi)$, defined as the infimum of the Lagrangian function.

Note that $\delta_{ij\lambda}^{sdn} = 0$, if $\alpha_{sdn} = 0$. Similarly, $\phi_{i,\lambda a}^{sdn} = 0$, if $\alpha_{sdn} = 0$. Therefore, we have the following two important relations, which lead to the decomposition of DP.

$$\delta_{ij\lambda}^{sdn} = \alpha_{sdn} \delta_{ij\lambda}^{sdn} \quad \forall e_{ij} \in E, 0 < \lambda \leq W \quad (8)$$

$$\phi_{i,\lambda a}^{sdn} = \alpha_{sdn} \phi_{i,\lambda a}^{sdn} \quad \forall i \in V, 0 < \lambda \leq W \quad (9)$$

By using (8) and (9), and removing the terms that are independent of the decision variables, the problem (7) becomes:

$$\min_{A, \Delta, \Phi} \left\{ \sum_{s_{sdn} \in L} [(1 - \alpha_{sdn}) P_{sdn} + \alpha_{sdn} D_{sdn}] \right\} \quad (10)$$

where D_{sdn} is:

$$D_{sdn} = \sum_{e_{ij} \in E} \sum_{0 < \lambda \leq W} (d_{ij\lambda} + \xi_{ij\lambda}) \delta_{ij\lambda}^{sdn} + \sum_{i \in V} \sum_{0 < \lambda \leq W} \sum_{0 < a \leq W} (c_{i\lambda} + \pi_{i\lambda}) \phi_{i,\lambda a}^{sdn} \quad \forall s_{sdn} \in L \quad (11)$$

The relaxed problem (10) can be further decomposed into subproblems, each of which corresponds to one request. Corresponding to s_{sdn} , the request-level subproblem is defined as follows:

$$\min_{\alpha_{sdn}} \left[(1 - \alpha_{sdn}) P_{sdn} + \alpha_{sdn} \min_{\Delta_{sdn}, \Phi_{sdn}} (D_{sdn}) \right] \quad (12)$$

subject to the constraints (3) and (5).

B. Solving the Derived Subproblem and Constructing a Feasible Solution

The formulation of the subproblem (12) is similar to the problem solved by the Modified Minimum Cost Lightpath (MMCSLP) algorithm [6]. The complexity of solving MMCSLP is $O((N+W)NW)$. The subgradient method [7] is used to solve DP.

We use a heuristic algorithm to construct a feasible routing scheme. Generally, a solution to DP might be an infeasible resource allocation, because some constraints are relaxed when creating DP. We use the heuristic algorithm proposed in [6] to obtain a feasible resource allocation. The solution to DP is adjusted by rejecting some accepted requests, and modifying resource allocations for the lightpaths. In the worst case, the computational complexity of the heuristic algorithm is $O(L(NW)^2)$.

C. Evaluating the Constructed Resource Allocation

We use the *duality gap* to evaluate the resource allocation that is constructed by the heuristic algorithm. The upper and

lower bounds of the optimal value J^* of the primary function can be estimated: its upper bound is the value of J corresponding to a feasible resource allocation; and its lower bound is the optimal value q^* of the dual function. The difference $(J^* - q^*)$ is known as the *duality gap*. Moreover, an upper bound of the duality gap is $(J - q^*)$. Using the upper bound of the duality gap as its estimate, a relative duality gap, defined as $(J - q^*)/q^*$, is used as a measure of the sub-optimality of a feasible resource allocation. Thus even without obtaining the exact optimum, we know the distance of a sub-optimal solution from the optimum is within a certain range.

V. RESULTS OF USING STATIC GoS DIFFERENTIATION

The challenge of solving the static GoS differentiation problem is that the required resources in providing service to a request highly depend on the overall resources usage and availability, which in turn depend on the acceptance of other requests. Thus the cost of a lightpath between a given node pair varies depending on the routing of the lightpath. Therefore, the decision of accepting or rejecting requests and associated resource allocations must be considered together, not separately. This is the reason why the existing dynamic GoS differentiation mechanisms cannot be readily adapted for the static case. From the algorithm design point of view, the searching space of the global optimal solution of a static resource allocation is several orders greater than that of the dynamic case.

The network topology of a sample 22-node network is shown in Figure 1. The lightpath requests are represented by a matrix, where the horizontal/vertical axis represents the source/destination node, and a number in the matrix represents the number of the requested lightpaths between a given node pair. The traffic matrix that is used in the simulation is shown in Table 1.

In the first simulation, we demonstrate the impact of $d_{ij\lambda}$ on the hop-count of lightpaths. In this simulation, we do not differentiate requests and fix the price of lightpaths by assigning all P_{sdn} 's to the same value. As $d_{ij\lambda}$ decreases (shown in Figure 2), lightpaths are able to take more hops (a hop is defined as one fibre link), and thus with the same amount of resources, more lightpaths can be accommodated. Since the revenue of a lightpath (P_{sdn}) is set to 1000, when $d_{ij\lambda} = 510$, no lightpath can be over two hops, because the revenue cannot cover the cost of two hops. The number of one or two-hop lightpaths hardly changes as $d_{ij\lambda}$ changes, because rerouting options are very limited for these lightpaths. However, for a lightpath with three or more hops, there is a variation as $d_{ij\lambda}$ changes, indicating trade-offs exist. For example, when $d_{ij\lambda}$ varies from 250 to 10, with little sacrifice of less lightpaths of two or three hops, our algorithm is able to accommodate many more lightpaths of a larger hop-count. Through simulations of other networks (results are not shown), we observe that the distribution of

lightpath hop-counts strongly depends on the nodal degree and the traffic load of the network. For a highly connected network (i.e., high nodal degree) under a heavy traffic load, when $d_{ij\lambda}$ is small to medium, more requests are accepted by using lightpaths with large (more than four) hops. The reason is twofold: on the one hand a highly connected network provides more options in routing a lightpath, and on the other hand, a heavy traffic load demands more resources, creating severe resource competition. However, regardless of network topology (we simulated 14, 22 and 28 node mesh networks) and traffic model, most of the lightpaths are routed within 4-6 hops. In Figure 3, the trade-offs between the average hop-count and the number of rejected requests are presented. As $d_{ij\lambda}$ increases, the average hop-count drops, while the number of rejected requests increases. There is a change when $d_{ij\lambda}$ reaches 250, because all the lightpaths having more than three hops cannot afford the resource cost any more. There is a sharp change when $d_{ij\lambda}$ reaches 500, after which no lightpath can afford more than a single hop.




Figure 1. Network topology of a 22-node network

TABLE 1. TRAFFIC MATRIX

0	0	2	0	0	0	2	0	2	0	2	0	2	0	2	2	0	2	0	0	0	2	0
0	0	2	2	0	0	2	0	2	0	0	0	2	0	2	0	2	0	0	2	2	0	2
2	2	0	2	0	0	2	2	0	0	2	2	0	0	2	2	2	0	2	2	0	2	2
2	0	0	2	0	0	2	2	0	2	0	2	0	2	2	0	2	0	2	2	0	2	0
0	2	0	0	0	2	0	2	0	0	2	2	0	0	2	0	0	2	0	2	2	0	2
0	2	0	2	0	0	2	0	0	0	2	2	2	0	0	2	0	0	2	2	0	2	0
1	1	2	0	2	2	0	2	0	1	0	2	2	2	0	0	1	2	0	0	0	2	0
2	0	1	2	0	0	0	0	2	1	0	2	0	0	2	0	2	0	2	0	2	0	2
2	0	0	2	2	0	0	0	1	0	0	0	0	2	0	2	2	2	1	1	0	2	0
0	0	0	1	0	1	0	0	0	0	0	2	0	0	1	0	1	0	0	0	0	0	0
0	0	0	1	0	2	0	0	0	2	0	2	0	0	1	0	0	0	2	2	1	0	0
1	2	0	0	2	1	2	1	1	1	0	0	2	1	0	2	1	0	2	1	2	1	0
1	0	0	1	0	0	1	0	2	0	1	0	0	2	2	1	1	1	0	0	0	0	1
0	0	0	1	0	0	2	0	0	1	0	2	0	0	1	0	0	0	0	0	0	0	1
2	0	0	0	0	1	2	1	1	0	1	0	2	0	0	1	2	0	0	1	1	0	0
0	1	0	2	1	0	2	2	0	0	0	0	1	2	0	0	0	1	0	0	2	0	2
0	2	0	1	0	0	1	0	2	0	0	0	2	0	1	0	0	0	2	0	0	1	0
1	1	2	0	2	2	0	2	0	1	0	2	2	2	0	0	0	2	0	0	0	0	0
2	0	1	2	0	0	0	0	2	1	0	2	0	0	2	0	2	0	2	0	2	0	2
0	0	0	1	0	2	0	0	0	2	0	2	0	0	1	0	0	0	0	2	1	0	0
0	0	0	0	1	0	0	0	0	0	2	0	0	0	1	0	1	0	0	0	0	0	0
1	0	0	1	0	0	1	0	2	0	1	0	2	2	1	1	0	0	0	0	0	0	0

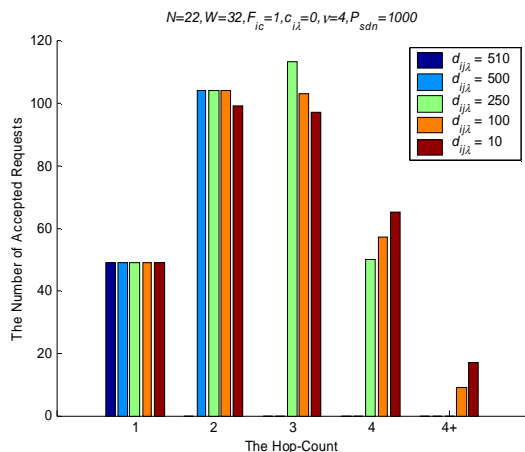


Figure 2. Distribution of lightpath hop-count




Figure 3. Trade-off between average hop count and the number of rejected requests

We investigate the impact of the cost of converters on their usage. As their cost increases, their usage decreases. We quantitatively demonstrate such a trend (shown in Figure 4). When $c_{i\lambda}$ increases from 0 to 25, the use of converters drops most drastically, and after that, it decreases moderately with $c_{i\lambda}$. When it reaches 700, no converter is used. The reason is that when a lightpath uses a converter, the lightpath must have at least two hops. Thus the cost of the lightpath is at least $c_{i\lambda} + 2d_{ij\lambda}$. In the example, after $c_{i\lambda}$ reaches 700, the cost of the lightpath is at least 900, thus no profit can be generated from providing a lightpath using a converter for a request.




Figure 4. The use of converter with respect to its cost

In the next example, we evaluate the impact of resource cost on the acceptance of requests in the presence of GoS differentiation. We differentiate the requests between different node pairs as two grades: Distinct Grade (DG) and Regular Grade (RG). The mask of the DG is shown in Table 2, where 1 represents a DG request, while 0 represents an RG request. As $d_{ij\lambda}$ increases, the rejection of DG and RG requests exhibits different behaviors (Figure 5). Since a DG

