Packing and Least-Loaded Based Routing in Multi-Rate Loss Networks*

Ibrahim Matta
College of Computer Science
Northeastern University
Boston, MA 02115
matta@ccs.neu.edu

Marwan Krunz
Department of Electrical & Computer Eng.
University of Arizona
Tucson, AZ 85718
krunz@ece.arizona.edu

Abstract

We examine various schemes for dynamically routing virtual circuits (VCs) in a multi-class network. A VC setup request may be rejected by admission control because resources are either unavailable or being reserved for future incoming VCs. We examine least-loaded based schemes, which attempt to balance the load among available routes. In addition, we examine packing based schemes, which attempt to reduce bandwidth fragmentation possibly at the expense of load balancing. Our simulation results show that under skewed workload, our packing based scheme outperforms a traditional least-loaded based scheme in terms of revenue (or equivalently, network utilization). Under uniform workload, both schemes provide similar revenue.

1 Introduction

We consider a network that uses the Virtual Path (VP) concept. This concept is often used to simplify network management and to increase the apparent direct connectedness of the network [1, 4, 9]. Typically, a VP is installed between two nodes (switches) over a sequence of physical links, and bandwidth is allocated to it. Thus a virtual fully-connected network can be overlayed over the physical network, where the VPs constitute the (virtual) links connecting the network nodes. Simple routing schemes that only consider paths with one link (called *direct routes*) and two links (called *alternative routes*) are then used. For a fully-connected network with N nodes, each pair of nodes has one direct route and $N-2$ two-link alternative routes. A number of such routing schemes were designed for telephone networks [3] and recently for ATM networks [10, 4, 6, 5].

*I. Matta was partially supported by NU-RSDF grant 377090.

We consider a network that supports $S \geq 2$ classes of connections (or *virtual circuits*). A virtual circuit (VC) of class s requires the reservation of a certain amount of bandwidth b_s that is enough to ensure a given quality-of-service (QOS). This bandwidth can be thought of either as the peak transmission rate of the VC or its "effective bandwidth" [2] varying between the peak and average transmission rates. Without loss of generality, we assume the bandwidths requested by different classes are distinct and the classes are indexed in increasing order of their requested bandwidths, i.e., $b_1 < b_2 < \cdots < b_S$.

To support a class-s VC, the VC has to be setup on some path from the source to the destination; the QOS demand (b_s) is allocated on each of the links for the lifetime of the VC. Since network bandwidth is limited, some requests for VC setup are denied (blocked) by the admission control algorithm. The objective of the routing algorithm is to choose routes that result in high successful VC setup rate (or equivalently, high carried VC load) while maximizing the utilization of network resources (or equivalently, revenue).

Routing schemes are commonly based on the least-loaded with trunk reservation concept [8]. Here each link has a Trunk Reservation (TR) value associated with it. A two-link alternative route is said to be *TR-permissible* if, for each of its links, there is still a certain amount of idle bandwidth available beyond the corresponding trunk reservation level. For example, consider a link (on an alternative route) that has idle bandwidth of 100 units and TR value of 10 units, then the idle bandwidth considered available is 100 - 10 = 90 units. When a new VC arrives, it is setup on the direct route between the VC's source and destination provided it can support the VC's bandwidth requirement. Otherwise, the VC is setup on the *least-loaded* TR-permissible alternative route if there is at least one that can support the VC. Thus, the scheme attempts to evenly distribute the load among the alternative routes. If the direct route and all the two-link alternative routes are unavailable, the VC is blocked.

Clearly, routing the VC on a two-link alternative route consumes twice as much bandwidth. This is why trunk reservation is used in order to discourage using two-link routes, and thus reserve some amount of bandwidth for future direct VCs.

Recently [4], it has been recognized that in order to maximize the utilization of available resources, a routing policy in a heterogeneous (multi-rate) environment should implement *packing* of narrowband VCs (having relatively small bandwidth requirement) on some VPs in order to leave room on other VPs for wideband VCs. The scheme proposed in [4] attempts to pack class-s VCs by keeping in perspective only the next higher class of VCs. In this paper, we extend the scheme in order to account for *all* higher classes. In [4], the packing based scheme and the traditional least-loaded based scheme were compared under uniform workload on a 6-node network supporting only two classes of VC. In this paper, we consider larger networks supporting more than two classes and subjected to skewed workload.

We also propose another scheme based on the least-loaded concept. This scheme, however, would sometimes choose a two-link alternative route over the direct route, if the former is significantly attractive. We compare the various schemes by simulation. Our results show that under skewed workload, our proposed schemes, the extended packing based and the modified least-loaded based, give higher revenue at the expense of higher VC blocking probability (or low carried VC load). However, under uniform workload, all schemes provide similar revenue. We also discuss the effect of trunk reservation on fairness among different VC classes.

The paper is organized as follows. Section 2 presents the different routing and admission control algorithms that will be examined. Section 3 compares the various algorithms using discrete-event simulation. Section 4 concludes with future work. Details omitted because of lack of space can be found in [7].

2 Routing and Admission Control Algorithms

In this section, we present the routing and admission control algorithms that will be examined. We first introduce the following definitions.

- A route is said to be *QOS-permissible* if it has sufficient idle capacity to carry the VC.
- In this paper, we use two definitions for the *TR-permissibility* of a two-link alternative route. For simplicity, we will assume that all links have the same TR value.

  Definition 1. An alternative route is said to be TR-permissible if its idle capacity minus the reservation threshold is greater than or equal to the requested bandwidth of the incoming VC [4].

  Note that the idle capacity should then exceed a certain amount of bandwidth that varies depending on the class of the incoming VC. This further discourages higher VC classes (with higher bandwidth requirements) from using alternative routes. We thus refer to this as "class-dependent reservation". Also note that if an alternative route is TR-permissible then it is also QOS-permissible, and hence allowable.

  Definition 2. An alternative route is said to be TR-permissible if only when it carries at least one direct VC on one of its links, the idle capacity must be greater than or equal to a reservation threshold that is independent of the class of the incoming VC.

  This definition of TR-permissibility requires that switches keep track of the number of direct VCs on outgoing links. This avoids unnecessary reservations for direct VCs when not present. Also, since the reservation does not depend on the class, we ensure that all classes are treated fairly concerning the use of alternative routes. We refer to this as "class-independent reservation".

- A two-link alternative route is said to be *allowable* if it is both QOS-permissible and TR-permissible.

2.1 Algorithms with Class-Dependent Reservation

In the following three algorithms, TR-permissibility is defined as given by Definition 1.

2.1.1 Least-Loaded Routing (LLR)

The Least-Loaded Routing (LLR) algorithm is a common algorithm that attempts to evenly distribute the load among the alternative routes. The following steps are executed when a new VC arrives:

1. Set up the VC along the direct route if the direct route is QOS-permissible. Otherwise, go to step 2.
2. If no allowable alternative routes are available, then the VC request is rejected. Otherwise, set up the VC on the allowable alternative route with the largest idle capacity, i.e. the least loaded.

2.1.2 Most-Loaded Routing (MLR)

The Most-Loaded Routing (MLR) algorithm attempts to pack VCs on some alternative routes by favoring the most utilized. Thus, it attempts to leave other alternative routes very lightly loaded, and thus increase the chance of admitting incoming VCs with large bandwidth requirements.

2.1.3 Multi-Rate Least-Loaded Routing with Packing (MLLRP)

The Multi-Rate Least-Loaded Routing with Packing (MLLRP) algorithm proposed by Gupta [4] attempts to pack class-s VCs by considering the next higher class of VCs. The basis of the algorithm is as follows: if a class-s VC cannot be routed on the direct route, then route the VC on the allowable alternative route that has the largest *incremental capacity*. The incremental capacity of an alternative route is defined as the maximum number of class-s VCs the route can support after carrying the maximum possible number of class-$\bar{s}$ VCs, where $\bar{s} = min(s+1, S)$.

2.2 Algorithms with Class-Independent Reservation

In this subsection, we propose two algorithms for which TR-permissibility is defined as given by Definition 2.

2.2.1 Multi-Rate Least-Loaded Routing (MLLR)

The Multi-Rate Least-Loaded Routing (MLLR) algorithm forces classes of VC with lower bandwidth requirements on alternative routes when the load of the given direct route reaches a certain threshold. So even if the direct route can support the VC, the MLLR algorithm may route the VC on the least utilized allowable alternative route provided that it is sufficiently "attractive". The alternative route is said to be attractive if its idle capacity exceeds x times the idle capacity of the direct route. The value of x is chosen to be higher for classes with higher bandwidth requirements. In particular, it is chosen to be in proportion to the load/demand imposed by each class and to account for the fact that using an alternative route, instead of the direct route, consumes twice as much network bandwidth. We define the value of x as follows:

$$x = 2^{(10\ \rho_s\ b_s)/(\sum_{k=1}^{S} \rho_k\ b_k)},$$

where ρ_s is the traffic intensity of class-s VCs. We assume here that the ρ_s are known a priori (based on traffic forecasts) or dynamically estimated. Thus, classes of VC with lower bandwidth requirements have higher chances to get routed on alternative routes. This will reserve bandwidth on direct routes for classes of VC with higher bandwidth requirements while causing minimal loss of network resources by routing only classes with lower bandwidth requirements on alternative routes. The following steps are executed when a new class-s VC arrives:

1. If the direct route is not QOS-permissible and there are no allowable alternative routes, then reject the VC. Otherwise, go to step 2.

2. If the direct route is QOS-permissible and the least-utilized alternative route is either not allowable or not attractive, then setup the VC on the direct route. Otherwise, the VC is routed on the least-utilized alternative route.

2.2.2 Generalized Multi-Rate Least-Loaded Routing with Packing (GMLLRP)

As a generalization to MLLRP, we propose the GMLLRP algorithm, which is similar to MLLRP except for the definitions of TR-permissibility and incremental capacity. TR-permissibility is given by Definition 2. Regarding incremental capacity, we extend its definition to not only account for the next higher class of VCs when routing an incoming class-s VC, but *all* higher classes. Thus GMLLRP attempts to pack class-s VCs in order to reduce blocking for all higher classes.

We define the *incremental capacity* of the alternative route as the maximum number of class-s VCs the route can support assuming the maximum possible number of class-$\bar{s}$ VCs are carried, where $\bar{s}$ $(s+1 \leq \bar{s} \leq S)$ is the higher class that minimizes incremental capacity. Note that for a class-s VC request, we only consider the incremental capacity due to classes $s+1, \cdots, S$ since the incremental capacity due to classes less than s is always zero. The following steps are executed when a new class-s VC arrives:

1. Set up the VC along the direct route if the direct route is QOS-permissible. Otherwise, go to step 2.

2. If no QOS-permissible alternative routes are available, then the VC request is rejected. Otherwise, route the VC on the route that has the largest incremental capacity if it is also TR-permissible; in case of ties, the VC is routed on the route with largest idle capacity. If no such route exists, then the VC request is rejected.

3 Performance Evaluation

3.1 Network Model

We consider a fully-connected logical VP network, which could be carved out over an arbitrary underlying physical topology. We assume all VP links have the same total bandwidth and the same reservation threshold. The network is used by a number of VC classes. A class-s VC requires the reservation of b_s units of bandwidth. Class-s VC setup requests arrive to the network according to a Poisson process of rate λ_s. Each class-s VC, once it is successfully setup, has a lifetime of exponential duration with mean $1/\mu_s$.

To model a *skewed* workload, we assume each VC class has different arrival rate and average lifetime. Furthermore, the network is partitioned into two equal groups, each containing half of the total number of nodes N. The source and destination nodes of a VC are chosen randomly from the same group. The group is chosen with some specified probability, p_{skew}. A node in another group may be chosen by the routing algorithm to act as the intermediate node in a two-link path.

3.2 Performance Measures

To evaluate the performance of the algorithms, our main measure is *revenue*, which is defined as

$$revenue = \sum_{k=1}^{S} \rho_k \, (1 - B_k) \, b_k$$

where $\rho_k = \frac{\lambda_k}{\mu_k}$, and B_k is the blocking probability of class k. The reservation threshold parameter is assumed to be set *a priori* such that revenue is maximized. We also define the *carried load* as

$$carried\ load = \sum_{k=1}^{S} \rho_k \, (1 - B_k)$$

The carried load measure gives the average number of VCs carried by the network.

3.3 Numerical Results

We present results for a 20-node network, i.e., $N = 20$. Each VP link has a total of C units of bandwidth. Here we take $C = 96$. We have four classes of VC with $b_1 = 1.3$, $b_2 = 4.1$, $b_3 = 6.7$ and $b_4 = 9.9$. The arrival rates are $\lambda_1 = 0.4\lambda$, $\lambda_2 = 0.3\lambda$, $\lambda_3 = 0.2\lambda$ and $\lambda_4 = 0.1\lambda$, where λ is the total VC arrival rate. The departure rates are $\mu_1 = 0.004$, $\mu_2 = 0.003$, $\mu_3 = 0.002$ and $\mu_4 = 0.001$. Notice that we have chosen the parameters such that the highest class of VC, which might represent large video conferences requiring the largest amount of bandwidth, arrives less often and holds on longer. We take $p_{skew} = 0.8$.

Figure 1 shows revenue and carried load for LLR, MLR and MLLRP. LLR performs better or as well as the other two. Figure 2 compares LLR with our proposed schemes, MLLR and GMLLRP. Both outperform LLR in terms of revenue at the expense of carried load as they tend to accept fewer low-bandwidth VCs and more bandwidth-intensive VCs. Thus there is a clear tradeoff between maximizing network utilization and accepting as many VCs as possible. Figure 3 demonstrates this by showing the class blocking probabilities for LLR and MLLR. It also illustrates the fairness of class-independent reservation as it brings the blocking probability of different classes within a smaller range.

4 Conclusions and Future Work

We proposed packing based and least-loaded based VC control algorithms for multi-class VP-based networks. Under skewed workload, it was observed that packing results in increased revenue (or equivalently, network utilization). We showed that a least-loaded based scheme that sometimes selects long alternative routes over short direct routes, especially for low-bandwidth VCs, also results in increased revenue. This is true as long as the longer route is significantly attractive (lightly loaded) that it offsets the loss in network bandwidth. Future work remains to develop dynamic reservation thresholds that are adjusted based on traffic estimates. This is of practical interest when the input traffic is unknown or time-varying.

References

[1] I. Chlamtac, A. Faragó, and T. Zhang. Optimizing the system of virtual paths. *IEEE/ACM Transactions on Networking*, 2(6):581–587, December 1994.

[2] A. Elwalid and D. Mitra. Effective bandwidth of general markovian traffic sources and admission control of high-speed networks. *IEEE/ACM Transactions on Networking*, 1(3):329–343, June 1993.

[3] A. Girard. *Routing and Dimensioning in Circuit-Switched Networks*. Addison-Wesley Publishing Company, 1990.

[4] S. Gupta. *Performance Modeling and Management of High-Speed Networks*. PhD thesis, University of Pennsylvania, Department of Systems, 1993.

[5] S. Gupta, K. Ross, and M. ElZarki. Routing in virtual path based ATM networks. In *GLOBECOM '92*, pages 571–575, 1992.

[6] R.-H. Hwang, J. Kurose, and D. Towsley. MDP routing in ATM networks using virtual path concept. In *IEEE INFOCOM*, pages 1509–1517, Toronto, Ontario, Canada, June 1994.

[7] I. Matta and M. Krunz. Packing and least-loaded based routing in multi-rate loss networks. Technical Report NU-CCS-96-12, College of Computer Science, Northeastern University, Boston, MA 02115, September 1996. Available from http://www.ccs.neu.edu/home/matta.

[8] D. Mitra, R. Gibbens, and B. Huang. Analysis and optimal design of aggregated-least-busy-alternative routing on symmetric loss networks with trunk reservation. In *13th ITC*, Copenhagen, Denmark, 1991.

[9] Y. Sato and K. Sato. Virtual path and link capacity design for ATM networks. *IEEE J. Select. Areas Commun.*, SAC-9:104–111, January 1991.

[10] E. Wong, A. Chan, S. Chan, and K. Ko. Bandwidth allocation and routing in virtual path based ATM networks. In *IEEE ICC*, pages 647–652, 1996.

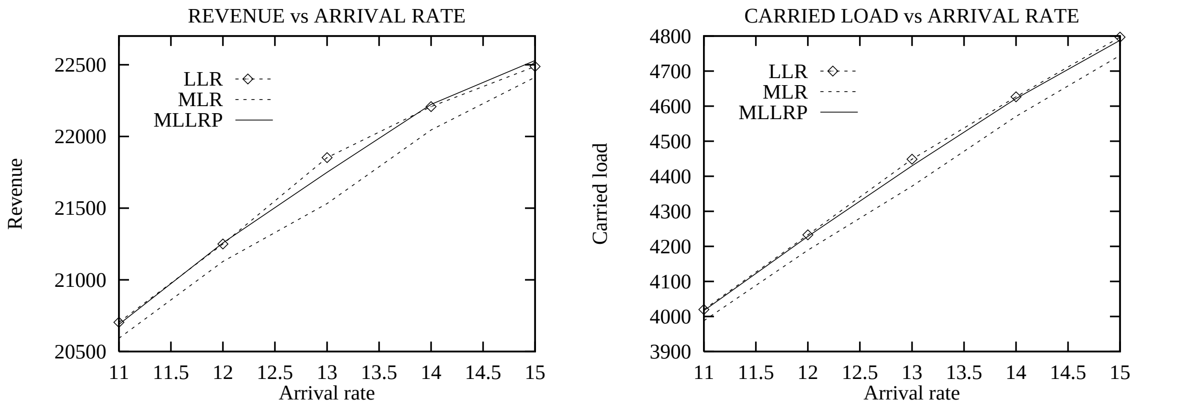

Figure 1. Revenue and carried load versus total VC arrival rate. Skewed workload.

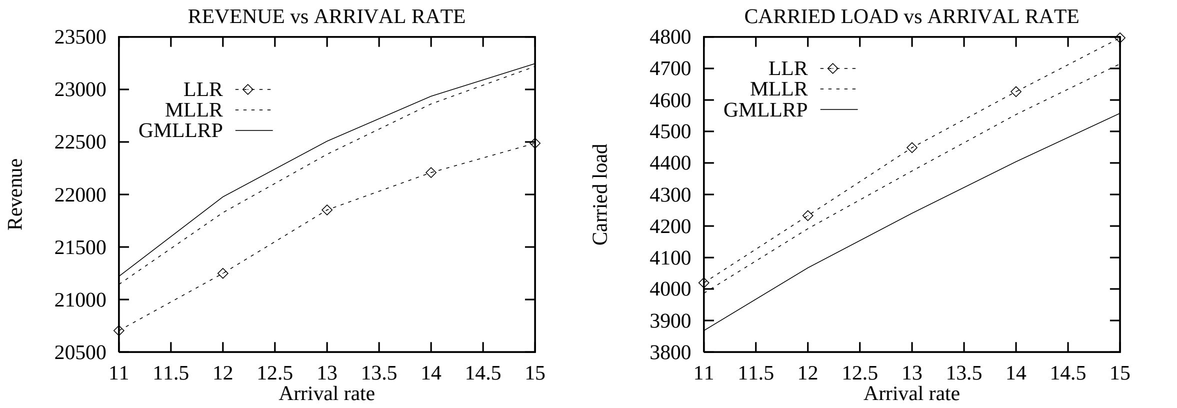

Figure 2. Revenue and carried load versus total VC arrival rate. Skewed workload.

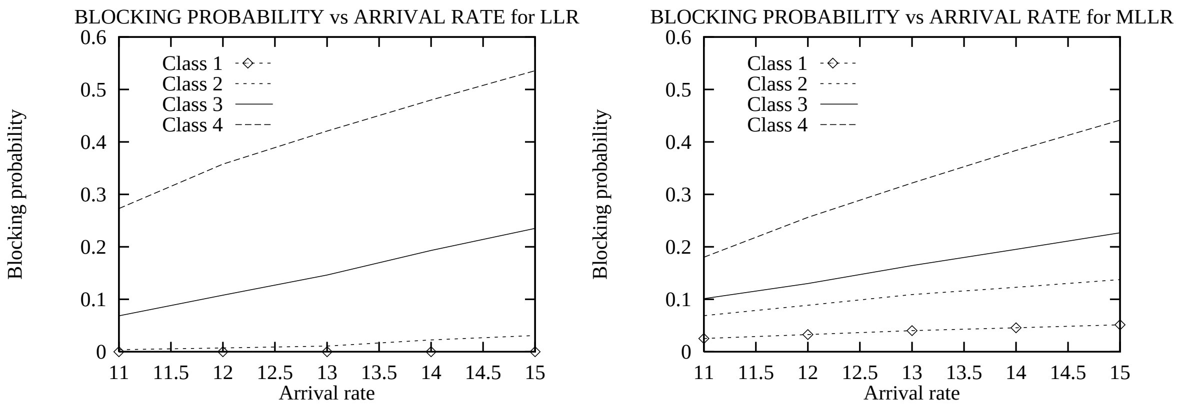

Figure 3. Class blocking probabilities versus total VC arrival rate for LLR and MLLR. Skewed workload.