Capacity Allocation in Optical Networks under Dynamic Lightpath Demands

Gangxiang Shen and Rodney S. Tucker
ARC Special Research Centre for Ultra-Broadband Information Networks, Dept. of Electrical and Computer Engineering
University of Melbourne, Melbourne, VIC, 3010 Australia; email: g.shen, r.tucker@ee.unimelb.edu.au

Abstract: We present a new approach for greenfield and incremental optical network capacity allocation under dynamic lightpath demands. The method can improve lightpath blocking performance by one or two orders of magnitude compared to conventional designs.

OCIS codes: (060.4250) Networks; (060.425) Networks, network optimization

1. Introduction

Efficient planning of optical networks requires an allocation of network capacity subject to certain economic or capacity constraints so that a forecasted network (lightpath) demand matrix or other performance objectives can be satisfied. Much research on optical networks has concentrated on static design, which aims to either minimize the maximal number of wavelengths in each fibre as in [1], or to minimize the maximal network congestion level as in [2], etc. However, when serving dynamic lightpath demands, a design simply based on a static forecasted demand matrix cannot perform well [3]. New approaches that take into account the lightpath demand dynamicity and uncertainty need to be developed. So far, though there is much research performed to evaluate lightpath blocking performance under dynamic lightpath demands [1][4], little research has been performed dedicated to planning and allocating network capacity for an optical network under lightpath demand dynamicity and uncertainty.

In this paper, we develop a new recursive approach for optical network planning and capacity allocation under dynamic lightpath demands. The novelty of this approach is that it integrates capacity allocation and performance evaluation in the context of dynamic lightpath demands. Specifically, a recursive process is proposed that uses performance parameters such as link capacity utilization to guide network capacity allocations. We consider two network design scenarios: greenfield and incremental. We carry out simulation studies to show that the proposed approach is effective for planning and capacity allocation of optical networks. The lightpath blocking probability of greenfield networks designed in this way are up to one or two orders of magnitude better than using conventional static designs.

2. Capacity allocations for greenfield and incremental optical networks

The input parameters of a greenfield optical network design include (i) a network topology G, (ii) a forecasted lightpath load matrix (in Erlang) T, (iii) a total network capacity budget C_{total}, and (iv) a lightpath connection routing algorithm R. The objective is to find an optimal allocation for the total capacity budget C_{total} such that the resultant design can achieve the best blocking performance when serving the lightpath load matrix T and full wavelength conversion capabilities are available at each node. The lightpath load (in Erlang) is a statistical mean of the number of lightpath connections requested between each pair of nodes. The total network capacity budget is a constraint on the total number of wavelength units for a network design. Each wavelength can carry one of data rates such as 2.5 Gb/s, 10 Gb/s, or 40 Gb/s.

An intuitive approach to solving the above capacity allocation problem is first to enumerate all possible capacity allocations, then to evaluate blocking performance for each of them, and finally to find the one that achieves the best blocking performance. Nonetheless, there is a huge number of such allocations even for a medium-size network. Thus, an exhaustive search is not scalable. We propose an alternative new recursive approach, which can solve the problem quickly and efficiently by incorporating a load-balancing effort in the design. The flowchart of the recursive approach is shown in Fig. 1.

Step 1 sets up initial parameters for the design. Step 2 makes an initial uniform capacity allocation. Step 3 runs an iteration of lightpath service provisioning to find a blocking performance of the current capacity allocation. Step 4 judges whether the current capacity allocation is good enough to terminate the recursive process. If so, Step 5 terminates the process. Otherwise, Step 6 computes average capacity utilization on each link according to the simulation data in Step 3, where A is the total number of lightpath request events (e.g., 10^5) that are simulated, $u_i^a(n)$ is real-time capacity utilization on link i when the a^{th} arrival event occurs in the n^{th} iteration. The key of the algorithm is the capacity reallocation equation in Step 7, in which the product $U_i(n) \cdot C_i(n)$ measures the average number of utilized wavelengths on link i in the previous simulation iteration, which implies the degree of network capacity desired on the link. If a link is not assigned sufficient capacity, its capacity utilization is high, and thus a larger weight and more capacity are assigned to the link in the capacity reallocation step, and vice versa. We have proved that the algorithm eventually converges.

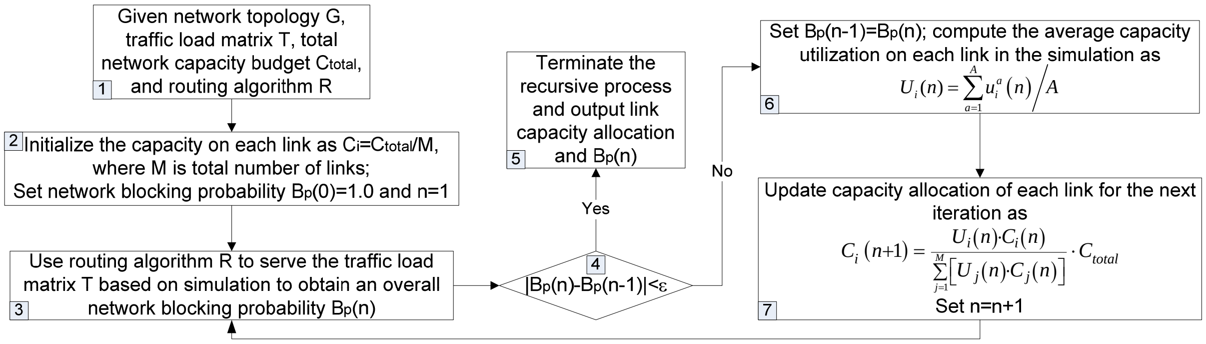

Figure 1: Flowchart for greenfield capacity allocation.

Based on the above greenfield network design, we can also perform an incremental network design. The objective of the incremental design is to efficiently allocate incremental capacity onto an existing network such that the capacity augmentation can bring the largest performance improvement to the network. We have extended the capacity allocation approach for the greenfield networks for this design. The only difference is that here we should efficiently utilize the existing capacity on each link when allocating the incremental capacity onto the network.

3. Simulation studies

To evaluate the performance of the proposed recursive approach, extensive simulations were performed. The following simulation conditions and assumptions were made. The arrival of lightpath requests follows a Poisson distribution and each established lightpath stays in the network according to a negative exponential distribution with a unit mean holding time. The blocking probability was computed based on the simulation of a total of 10^5 lightpath arrival events. For each arrival event, if the lightpath request was blocked, it was dropped. Although the proposed recursive approach can incorporate any routing algorithm, two typical routing algorithms were considered in the studies, including (i) a fixed shortest-path routing [1], and (ii) a *utilization-based* least-load routing algorithm [5], in which real-time link capacity utilization is measured as a cost to search the least congested route for each new service connection. The simulations were performed based on a test network 14-node 21-link NSFNET (as shown in Fig. 2). Finally, full wavelength conversion capability was assumed at each optical switch node in all the studies.

Greenfield design: We first consider the simulations of the greenfield network designs. For this design case, a uniform lightpath load matrix with 1.0 Erlang load per node pair was assumed though simulations were also performed for random lightpath load matrices (the results are similar). We set a total network capacity budget C_{total} for the NSFNET network to be 336 units (21 links × 16 wavelengths per link), which ensures on average each network link has a total of 16 wavelengths. Three capacity allocation cases were studied: (i) uniform capacity allocation, (ii) capacity allocation weighted by offered load accumulated on each link, and (iii) the new recursive approach described in Section 2 above. For the second case, we proportionally allocated more wavelength capacity to the links traversed by more forecasted lightpath load given that the shortest path routing algorithm was applied, but the allocation was performed only once for each service provisioning simulation.

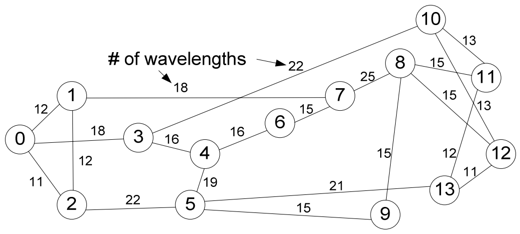

Figure 2: Test network: 14-node 21-link NSFNET.

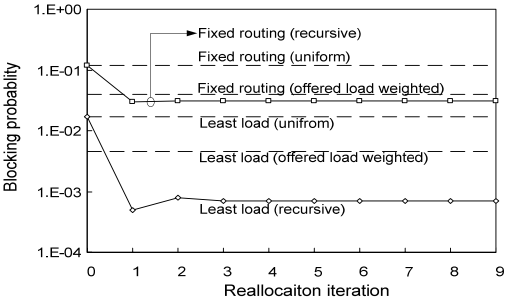

Figure 3: Blocking performance of greenfield network design.

Fig. 2 shows a greenfield network design based on the recursive approach and under the least-loading routing algorithm, in which a value shown by each link is the number of wavelengths allocated to the link. Note that we also found the link capacity allocations for the other design cases, but do not show them here. Fig. 3 compares the blocking performance for the above three capacity allocation approaches. All the curves with legend "recursive" correspond to the new recursive approach, the curves with legend "uniform" correspond to the uniform capacity allocation, i.e., case (i), and the curves with “offered load weighted” correspond to the allocation case (ii) above. In addition, all the curves with legend “fixed routing” are the results based on the fixed shortest-path routing algorithm, while for the curves with legend “least load,” the least-load routing algorithm was used. Based on the curves, we can see that the recursive capacity allocation process converges (with an increasing “reallocation iteration” index on the x-axis). After one iteration, the blocking probability (as shown in Fig. 3) reaches a minimum. At convergence, we can also see that the recursive approach always achieves a better lightpath blocking performance than the other two allocation cases. Under the least-load routing, the blocking performance improvement by the recursive approach is so significant that it outperforms the offered load weighted allocation seven times, and the uniform allocation even more than 200 times. A similar convergence can also be found in Fig. 4, which shows capacity utilization distribution on different links (see “link index” on the x-axis) under the recursive design approach for the least-load routing algorithm. With the recursive capacity allocation, the capacity utilizations converge to an almost identical value on every links, which implies a load balancing in the network.

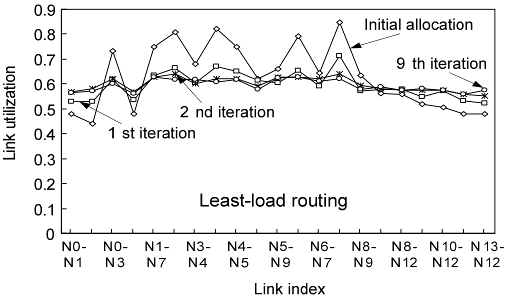

Figure 4: Link utilization distribution of greenfield network design.

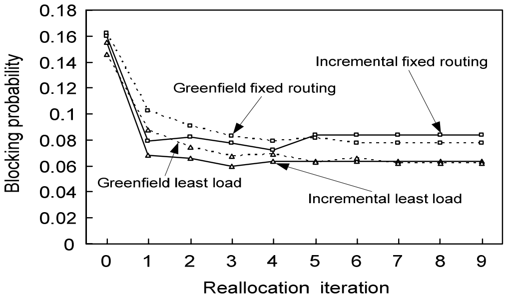

Figure 5: Blocking performance of incremental network design.

Incremental design: We also evaluated blocking performance for the incremental network designs. For this, we employed the greenfield network designs above as initial networks. To devise an incremental traffic environment and lightpath demand uncertainty, the lightpath load between each node pair was randomly changed within a range of [0, 2.8] Erlang, but the total lightpath load of the network was increased to be 127.4 Erlang (NSFNET 91 node pairs × 1.4 Erlang per pair), which is equivalent to on average 1.4 Erlang lightpath load per node pair, a 40% increase of lightpath load per node pair. With the increase of lightpath load, the total capacity of the network was also increased by 10%, which equals 34 units of extra capacity. Also, for comparison, a benchmark blocking performance was evaluated based on a greenfield network design, which was designed subject to the same total capacity as that of the incremental design, i.e., a total of 370 units of capacity, a sum of total existing capacity (336 units) plus extra capacity (34 units). This benchmark greenfield design can function as a standard to measure the effectiveness of the incremental design.

Fig. 5 shows the blocking performance of the incremental design. The point at the zeroth iteration in each curve corresponds to a blocking probability when an initial greenfield network (before the capacity increase) serves the incremental lightpath load matrix. The subsequent points are the results of the recursive designs after more capacity is added to the network. The curves show that the proposed recursive approach is also effective for the incremental network design, which always brings the blocking performance to minimum within the first few iterations. Moreover, comparing the blocking performance of the incremental design and the benchmark greenfield design (with 370 units of capacity), we see that their overall blocking performances are close, which verifies the effectiveness of the recursive approach.

4. Conclusion

We proposed a new recursive approach for planning and capacity allocation of greenfield and incremental optical networks under dynamic lightpath demands. The approach integrates capacity allocation and performance evaluation through a load-balancing effort and is generic enough to efficiently allocate capacity for optical networks under any routing algorithm. Simulation results indicate that the proposed approach can efficiently allocate capacity for greenfield and incremental optical networks, and the outcomes can perform one or two orders of magnitudes better than conventional greenfield designs.

References:

[1] H. Zang, J. P. Jue and B. Mukherjee, “A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks,” *Optical Networks Magazine*, January 2000, pp. 47-60.

[2] B. Mukherjee, D. Banerjee, S. Ramamurthy and A. Mukherjee, “Some principles for designing a wide-area WDM optical network,” *IEEE Transaction on Networking*, vol. 4, no. 5, October 1996, pp. 684-696.

[3] D. Leung and W.D. Grover, “Capacity planning of survivable mesh-based transport networks under demand uncertainty,” *Photonic Network Communications*, vol. 10, no. 2, September 2005, pp. 123-140.

[4] I. Chlamtac, A. Ganz and G. Karmi, “Lightpath communications: an approach to high bandwidth optical WAN’s,” *IEEE Transaction on Communications*, vol. 40, no. 7, July 1992, pp. 1171-1182.

[5] G. Shen, S. K. Bose, T. H. Cheng, *et al.*, “Heuristic algorithm for efficient lightpath routing and wavelength assignment in WDM networks under dynamically varying loads,” *Computer Communications*, vol. 24, no. 3-4, February 2001, pp. 364-373.