A Hybrid Algorithm for Lightpath Assignment

Neena R. Kaushik, Silvia M. Figueira, and Stephen A. Chiappari
Santa Clara University
Santa Clara, CA 95053, USA
{nrkaushik,sfigueira,schiappari}@scu.edu

Abstract–**This paper focuses on the assignment of multiple-wavelength lightpaths and presents a computationally simple algorithm to increase the number of simultaneous lightpath assignments in optical networks. Several approaches for lightpath assignment have been proposed in the literature, and most of them involve some variation of selecting or rejecting edge-disjoint paths for lightpath allocation, followed by assigning wavelengths in order. We have developed the Hybrid algorithm which either selects or rejects an edge-disjoint path, depending on the length of that path. We show using simulations that the algorithm performs significantly better (up to ~1,300%) than the consistent selection or rejection of edge-disjoint paths, and it performs comparably with (and in some cases even better than) an algorithm that makes lightpath selections dynamically based on the current network state.**

Index Terms–**Lightpath Assignment, Routing and Wavelength Assigning.**

I. Introduction

Optical network reservation, either on-demand or in-advance, requires lightpath assignment. This assignment is done in terms of time-slots/lightpaths, i.e., end-to-end lightpaths are reserved for a specific time-slot. The assignment of lightpaths in optical networks consists basically of routing and wavelength assigning, an NP-complete problem, which has been extensively studied for on-demand reservations. This problem has been studied both with and without wavelength converters, and lower and upper bounds on the blocking probability have been found for the number of wavelengths [13]. In fact, in [14], the lower and upper bounds are defined based on the traffic pattern using integer linear programming and, in [2], algorithms for static traffic are analyzed and lower bounds are obtained.

Lightpath assignment in an optical network consists of two main steps: choosing the route and selecting the wavelengths within the route. Given the two main steps, generally, any assignment approach will either (1) select a route and then the wavelength within the route or (2) select a wavelength and then the route within which that wavelength is available. Different selection processes may be employed in each approach, optimizing specific quantities in each case. When assigning multiple-wavelength lightpaths, one can extend the first approach to (3) select a set of routes and then the wavelengths within the routes, while one can extend the second approach to (4) select a set of wavelengths and then the routes within which those wavelengths are available. In approach number 3, the wavelengths are assigned in a pre-determined order within the set of routes via prioritizing the selection of edge-disjoint paths, whereas, in approach number 4, the wavelengths are assigned in a pre-determined order over the set of routes via prioritizing the rejection of edge-disjoint routes. The assignment of wavelengths in order, so that the least possible number of wavelengths is used, is a form of *packing*, which is in fact a basic technique used for lightpath assignment, on top of which several approaches have been built (see Section II).

We have analyzed and compared the blocking probability obtained when wavelengths are packed and the selection of edge-disjoint routes is prioritized via selection or rejection. The blocking probability was obtained with LRSS, our Lightpath Request Scheduling Simulator. LRSS is a simulator written in C, which takes as input the network topology and a synthetic trace of requests for lightpaths, both in the form of ASCII text files. It accepts any network topology, over which it simulates different wavelength assigning algorithms.

We have used LRSS to compare two systematic assigning strategies used in the Balancing and Concentrating algorithms [3], which follow approaches 3 and 4 (as described above), respectively. Our experiments with LRSS led to the conclusion that each approach behaves better than the other one under special conditions, and that the length of the shortest edge-disjoint path between the source and destination in the requests definitely affects the behavior of the algorithms. Based on this conclusion, we have developed the *packing-based* Hybrid algorithm, which *balances* or *concentrates* each allocation depending on whether the length of each edge-disjoint path between the source and destination is less than or equal to a pre-determined *cut-off* length.

Note that our Hybrid algorithm is based on the systematic and basic Balancing and Concentrating algorithms defined in [3]. Enhancing the Hybrid algorithm to include optimization techniques (such as selecting the least loaded path) in the selection of the route and/or wavelength is possible, but further experiments are required to assess the benefits obtained by each extension.

This paper is organized as follows. Section II describes related work, Section III presents the Balancing and Concentrating algorithms, Section IV introduces the Hybrid algorithm, Section V presents the effects of nodal degree and wavelength converters on the Hybrid algorithm, Section VI compares the Hybrid Algorithm with a previously proposed “dynamic routing” algorithm [4], and Section VII concludes.

II. Related Work

A range of heuristics has been proposed for the RWA (Routing and Wavelength Assignment) problem. In order of performance, the following approaches have been proposed: *basic*, *porder*, *color*, *lpcolor*, *least-loaded*, *aurpack*, *aurexhaustive*. As indicated by the average blocking probability, the spreading of wavelengths generally provides the worst allocations for wavelength assignment. Pack-based strategies perform better, and with *packing*, the performance improves from using wavelengths in order (*porder*), to using the most used wavelength (*pcolor*), to using the routes with the most used wavelengths (*lpcolor*), to choosing the path which is left with the highest capacity after the assignment (*least loaded*), to using *aurpack* and *aurexhaustive* [6].

In [6], Hyytia and Virtamo show the behavior of these algorithms with simulations for a hypothetical network topology connecting various cities in Finland. They show that the heuristic choice influences the blocking probability. The authors also introduce a first-iteration policy to improve the blocking probability of the various algorithms. The iteration policy runs a simulation, then finds the cost and decides whether to accept or reject the connection along that path. It is important to note that the first-iteration policy is prohibitively computationally intensive.

In [10], performance-efficient algorithms specific for a wide range of topologies are presented. In [1], the routing and wavelength assignment problem is solved with greedy algorithms, which are specifically applicable to networks of trees. In [7], the lightpath assignment in an optical network is considered in the route, wavelength, and time domains. The property that all lightpaths cannot be needed at the same time is taken into consideration to assign wavelengths using a branch-and-bound and tabu search method, which achieves better performance than the fixed-alternate routing followed by first-fit wavelength assignment strategy.

In [5], Ho and Mouftah define a method for path selection, which is based on exchanging information about critical links in the network and avoiding those links during wavelength assignment. They show that this method reduces the blocking probability compared to a fixed-wavelength-assignment scheme. This method has an overhead imposed by the exchange of network link-state information. In [4], the authors show that allocating paths dynamically, based on a function dependent on the number of hops and current state of the wavelengths assigned, lowers the blocking probability. In [8], the wavelength assignment is done by graph-coloring optimizations and these graph coloring strategies achieve a performance which is close to the lower bound for their particular topologies. In [16], a *Relative-Least-Influence* algorithm is presented, which decides the wavelength assignment by taking the entire network state into account. This algorithm achieves lower blocking probability and is fair compared to the *Relative-Capacity-Loss*, *Max-Sum*, and the *Least-Influence* algorithms, which have been previously studied and which make a decision on the wavelength assignment depending on the current network state.

Our work differs from previous works on lightpath assignment because our Hybrid algorithm selects/rejects the alternate edge-disjoint path(s) based on the *cut-off* value (to be explained below) solely, and is therefore computationally simple. Note that we have compared the Hybrid algorithm with the algorithm proposed in [4] in Section VI.

III. Balancing vs. Concentrating Algorithms

In this section, we define the terms associated with lightpath assignment, explain the basic Balancing and Concentrating algorithms for lightpath assignment, and finally highlight specific cases in which each algorithm performs better. Note that the Balancing and Concentrating algorithms pack the wavelengths in a systematic manner, and most heuristic-based lightpath assignment algorithms are based on a variation of the balancing and concentrating strategies.

A. Definitions

1. *Lightpath Request*: Each request is specified by a source-destination pair, a time-slot, and a number of wavelengths. It can be granted or rejected for the specified time-slot depending on the current availability. Note that granting a request means assigning one lightpath from source to destination for each wavelength requested, during the time-slot specified.
2. *End-to-End Path*: A path that starts from the specified source and ends at the specified destination. The minimum number of links in an end-to-end path is one.
3. *Lightpath Assignment*: A request that is granted for an end-to-end path.
4. *Lightpath Rejection:* A lightpath request that is not granted for an end-to-end path.
5. *Edge-Disjoint Path*: Paths are edge-disjoint if they share no edges.

B. The Algorithms

- *Balancing Algorithm:* This algorithm finds all the edge-disjoint paths from the source to the destination. It tries to assign the first wavelength available along the shortest edge-disjoint path, the first wavelength available along the second shortest edge-disjoint path, and so on, along the nth edge-disjoint path. It then tries to assign the second wavelength available in the same order, and so on, up to the mth wavelength, until all the lightpaths requested are allocated.
- *Concentrating Algorithm:* This algorithm finds all the edge-disjoint paths from the source to the destination. It tries to assign all the wavelengths along the shortest

path. Then it tries to assign all wavelengths available along the second shortest edge-disjoint path, and so on, sequentially until all the lightpaths requested are assigned.

TABLE I

SCENARIOS USED IN THE EXPERIMENTS

Number	Experiment	Reservation Inter Arrival Time	Number of wavelengths
1	High traffic with constant wavelength requests	5 minutes	Constant: 1
2	High traffic with heavy-tailed wavelength requests	5 minutes	Zipf's Generalized: exponent = 3 capacity = 4
3	High traffic with uniform wavelength requests	5 minutes	Uniform: [1-4]
4	High traffic with constant wavelength requests	5 minutes	Constant: 4
5	Medium traffic with constant wavelength requests	15 minutes	Constant: 1
6	Medium traffic with heavy-tailed wavelength requests	15 minutes	Zipf's Generalized: exponent = 3 capacity = 4
7	Medium traffic with uniform wavelength requests	15 minutes	Uniform: [1-4]
8	Medium traffic with constant wavelength requests	15 minutes	Constant: 4
9	Low traffic with constant wavelength requests	30 minutes	Constant: 1
10	Low traffic with heavy-tailed wavelength requests	30 minutes	Zipf's Generalized: exponent = 3 capacity = 4
11	Low traffic with uniform wavelength requests	30 minutes	Uniform: [1-4]
12	Low traffic with constant wavelength requests	30 minutes	Constant: 4

We have analyzed the Balancing and Concentrating algorithms extensively for 4-node topologies, as shown in Fig. 1. The 4-node ring is basically a ring, a common configuration used in optical metro networks. The 4-node ring with one chord increases the paths available, increasing the capacity of the 4-node ring. The 3-node-ring-with-a-spike topology introduces a concept of asymmetry in the network, which is useful in comparing the wavelength-balancing and wavelength-concentrating algorithms. These topologies comprehensively cover the 4-node-based vertex graph for a variety of different features and also resemble a practical network, the OMNInet [12], which motivated the use of a four-node graph in our experiments.

For the scenarios given in Table I, the Balancing algorithm performs better, i.e., gives lower blocking probability when the number of two-hop source-destination requests is fixed to twenty-five, fifty, or seventy-five percent compared to the number of one-hop source-destination requests (see Figs. 2-4). The Balancing algorithm also performs better when the diameter link of the OMNInet topology is used for satisfying the lightpath requests, but is not requested as an end-to-end lightpath (see Fig 5). However, the Concentrating algorithm performs better when the requests were mainly for one-hop end-to-end lightpaths. The analysis leads to Observation I, which states an important result.

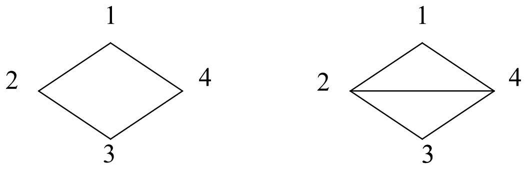

4-node ring

4-node partial mesh

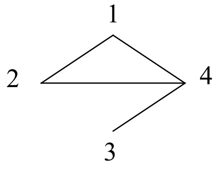

3-node ring with a spike

Figure 1. Topologies Analyzed.

Observation I: *The Balancing and Concentrating algorithms can each result in increased lightpath allocations, depending both on the source-destination pair of the lightpath request and on the topology.*

Explanation: Along a ring, balancing "short requests" and concentrating "long requests" lead to fewer open paths, which leads to higher blocking probability. For example, consider a 32-node ring with a capacity of four wavelengths. Assume one-hop requests between neighbor nodes only and consider three requests for four wavelengths between nodes (1,2), (2,3), (3,4) respectively. In this case, the Concentrating algorithm can satisfy the requests but the Balancing algorithm cannot, because Balancing (1,2) and (2,3) closes the way between 3 and 4. Consider again a 32-node ring with a capacity of four wavelengths. Two requests are now for four wavelengths between nodes (1,8) and (2,10), respectively. In this case, the Balancing algorithm can satisfy the requests, but the Concentrating algorithm cannot, because concentrating between 1 and 8 closes the way between 2 and 10.

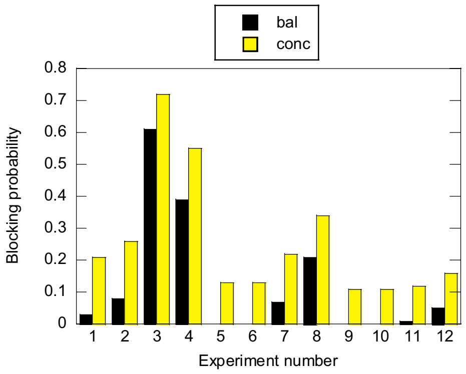

Figure 2. 25% of the requests for the link between nodes 1 and 3.

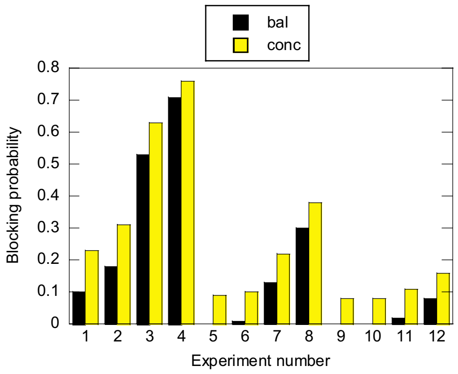

Figure 3. 50% of the requests for the link between nodes 1 and 3.

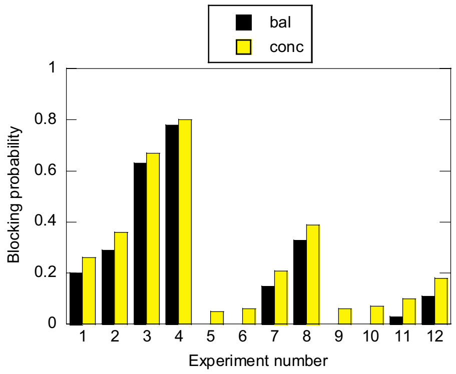

Figure 4. 75% of the requests for the link between nodes 1 and 3.

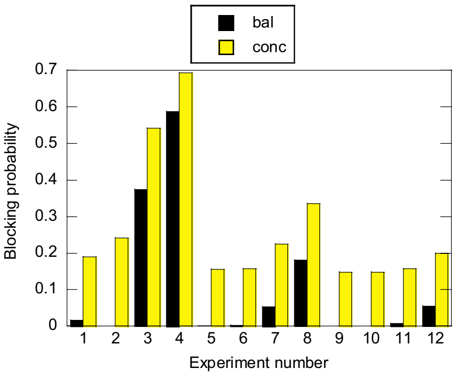

Figure 5. Blocking probability when there are no requests for the 2-hop path.

IV. The Hybrid Algorithm

To take advantage of the result stated in Observation I, we have developed the Hybrid algorithm for lightpath assignment. Our algorithm is a combination of the balancing and concentrating schemes. This algorithm either concentrates the assignments along the shortest paths or balances them along all the edge-disjoint paths, depending on the length of each path being considered.

Hybrid Algorithm

```
begin
  firstpass = 1;
  while (firstpass is not equal to 3)
    for i = 1 to number of wavelengths
      for j = 1 to number of edge-disjoint paths
        if first pass is equal to 1 and edge-disjoint path has more
                than x hops
          continue
        if wavelength[i] is available for all segments in the
                edge-disjoint path[j]
          Allocate wavelength[i] for all segments in
                edge-disjoint path[j]
          Increment number of allocated requests by 1
        if all requests are satisfied
          return the number of requests satisfied
      end (for edge-disjoint paths loop)
    end (for number of wavelengths loop)
    Increment firstpass by 1
  end (while loop)
  return the number of requests satisfied
end (begin)
```

The Hybrid algorithm, shown above, *packs* wavelengths by making allocations in two passes and depends on a non-negative integer parameter x. In the first pass, it avoids those edge-disjoint paths which have more than x hops, i.e., it achieves its concentrating effect in the first pass, by assigning only along those edge-disjoint paths which have fewer than x hops. In the second pass, it assigns the remaining requests along all edge-disjoint paths. Note that x represents the *cut-off* value between balancing and concentrating.

We use the terminology Hybrid and Hybrid-n as follows:

- Hybrid: the algorithm balances along the edge-disjoint paths during the first pass, only if the number of hops in the edge-disjoint paths is less than or equal to $N/2$, where N is the total number of nodes.
- Hybrid-n: the algorithm balances along the edge-disjoint paths during the first pass, only if the number of hops in the edge-disjoint path is less than or equal to n.

In a ring topology, every source-destination pair has two alternate paths, and the minimum length of the second-shortest path, which is edge-disjoint from the shortest path, is $N/2$. In a ring-with-chords topology, the minimum length of the second path, which is edge-disjoint from the shortest path, can be less than $N/2$. However, we have used *N*/2 as the minimum *cut-off,* in our analysis for both ring and ring-with-chords topologies, to study the effects of the Hybrid algorithm on both the ring and ring-with-chords topologies using the same *cut-off* parameter.

A. Gain with the Hybrid Algorithm

In subsection B, we present the percentage improvements obtained with the Hybrid algorithm over the concentrating and balancing approaches for the 32-node ring, 32-node ring with chords (see Fig. 9), and the National Lambda Rail (see Fig. 16) topologies. This improvement is calculated using the formula for percentage gain given below (see Table II for definitions of abbreviations and acronyms):

$$\text{Percentage Gain} = ((BP_x - BP_{hybrid})/BP_{hybrid})*100$$

TABLE II

ABBREVIATIONS AND ACRONYMS

Term	Definition
BP	Blocking Probability
BP_x	Blocking probability obtained with algorithm x where x can be any routing and wavelength assignment algorithm.

Note that in our experiments, for the 32-node ring and 32-node ring with chords, we defined the *cut-off* for the Hybrid algorithm to be half the length of the ring. For the National Lambda Rail (NLR), since it is not a regular ring, we have specified the *cut-off* as absolute values, 12 and 4, for traffic within the entire NLR and within individual rings of the NLR, respectively.

B. Simulation Results

The results presented in this section show the blocking probability gains obtained from simulations using the LRSS with a 32-node ring and ring-with-chords topologies, as well as the topology used in the National Lambda Rail.

TABLE III

REPRESENTATIVE REQUESTS

Variable	Value
Advance reservation request arrival for lightpath assignment	Exponential distribution
Average number of lightpath-request arrivals in a time-slot	10, 12, 15, 20, 30 and 60
Number of lightpath requests for each arrival	Constant = 1, Zipf's distribution Uniform Distribution [1 - 4]
Time-slot length	1 hour

The traffic characteristics are shown in Table III. The traces were generated using FONTS [9]. We consider ring topologies in our analysis because they are the basic building-blocks used to form more sophisticated topologies used in networks, such as the OMNInet [12], SURFnet [15], and National Lambda Rail [11]. The subsections below present a representative set of the experiments performed. We have experimented with a uniform source-destination pair distribution. The goal with these experiments is to show the behavior of the Hybrid algorithm and its gain over the basic and systematic, balancing and concentrating strategies. Note that all pack-based lightpath assignment heuristics are some variation of the balancing and concentrating strategies.

1. 32-Node Ring

For the 32-node ring, in general the Hybrid algorithm obtains positive gains (see Figs. 6-8) over the balancing and concentrating strategies. In some of the cases, negative gains are obtained, but they are significantly smaller than the positive gains. For very low traffic in the constant and Zipf's distribution cases, a significant negative gain is observed over Balancing, i.e., Balancing obtains lower blocking probability compared to the Hybrid algorithm, because, selecting the longer edge-disjoint path with balancing leaves more paths open along other wavelengths, thereby increasing the total number of assignments.

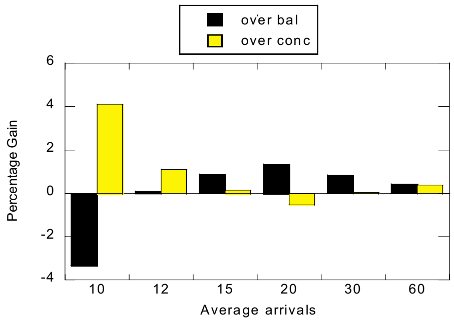

Figure 6. 32-Node Ring: Percentage Gain with Hybrid - Constant.

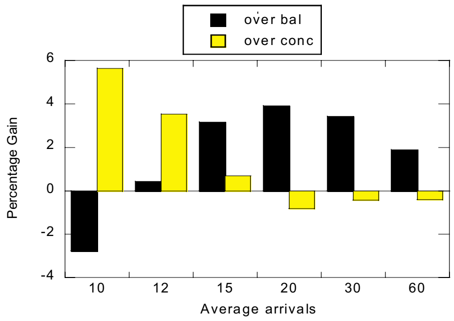

Figure 7. 32-Node Ring: Percentage Gain with Hybrid - Zipf.

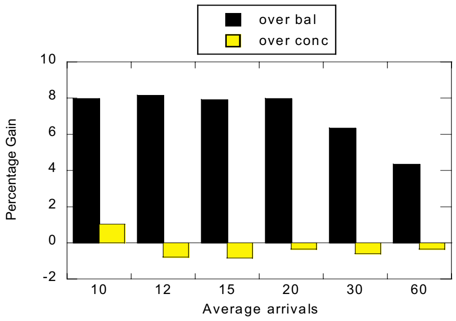

Figure 8. 32-Node Ring: Percentage Gain with Hybrid - Uniform.

2. 32-Node Ring with Chords

Adding chords to a network topology helps to increase the capacity of the network by providing extra paths between pairs of nodes. However, the impact caused by adding the chords depends on the ability of the assignment algorithm to take advantage of them. Fig. 9 shows a set of ring-with-chords topologies.

We have simulated the Balancing, Concentrating, and the Hybrid algorithms for 0, 1, 2, 4, and 6 chords, in a ring-with-chords topology, with an increasing number of wavelengths. The number of wavelengths is increased from 1 to 64, and the traffic used in all the cases is kept the same, which is 60 requests per time-slot. Note that the number of wavelengths is increased up to 64 only, because the blocking probability drops almost to 0 at that level.

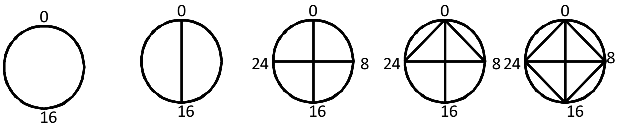

Figure 9. 32-Node Ring with Chords.

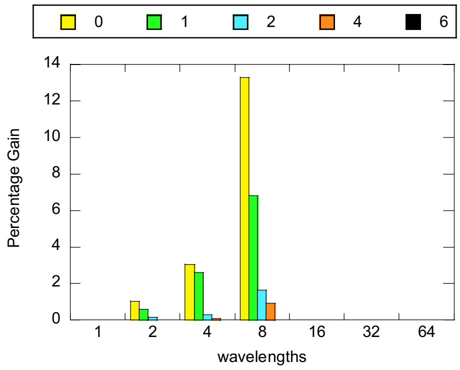

Figure 10. Percentage Gain - Hybrid over Balancing - Constant.

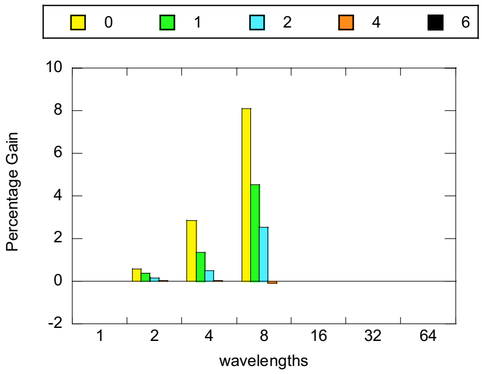

Figure 11. Percentage Gain - Hybrid over Balancing - Zipf.

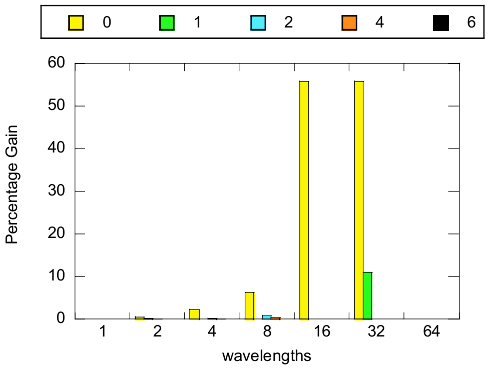

Figure 12. Percentage Gain - Hybrid over Balancing - Uniform.

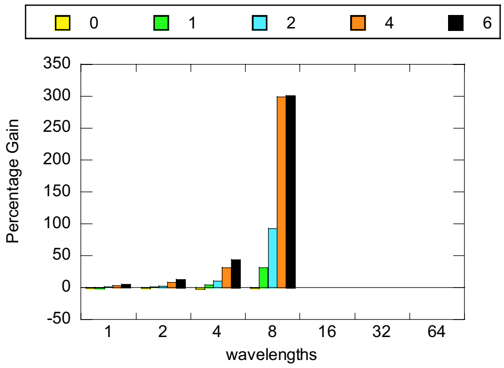

Figure 13. Percentage Gain - Hybrid over Concentrating - Constant.

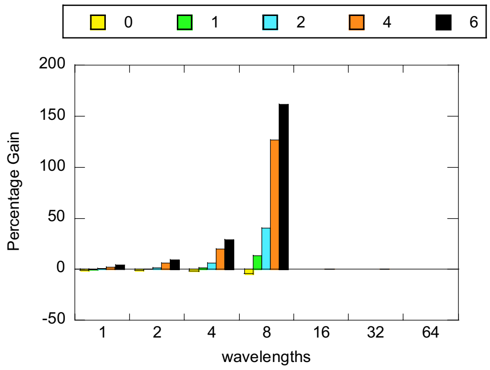

Figure 14. Percentage Gain - Hybrid over Concentrating - Zipf.

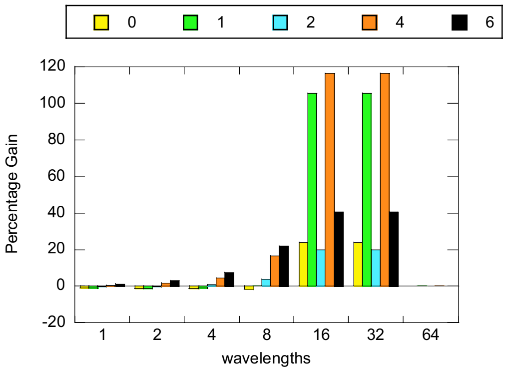

Figure 15. Percentage Gain - Hybrid over Concentrating - Uniform.

In general, the Hybrid algorithm provides positive gains over the balancing (see Figs. 10-12) and the concentrating (see Figs. 13-15) strategies. In some of the cases, we observe negative gains which are significantly smaller than the positive gains. Observe that percentage gains as high as ~300% are obtained over Concentrating (see Fig. 13) indicating that, for a ring with chords, the Hybrid algorithm is a better option. The percentage gain over Concentrating increases as the number of chords increases from 0 to 6. The percentage gains obtained over Balancing (see Figs. 10-12) in most cases are also significant, indicating that, for a ring with chords, it is better to balance along the edge-disjoint path, which has a length less than or equal to half the length of the ring, rather than balance along the edge-disjoint path which can have a maximum length greater than half the length of the ring.

Observe also that the gain increases up to a certain number of wavelengths, and then drops to zero. This can be explained as follows. At a low number of wavelengths, the traffic is high, and therefore all algorithms will have high blocking probability, which results in smaller improvements obtained with Hybrid, and hence low gain. At a high number of wavelengths, since we are using the same traffic, the traffic is low compared to the number of wavelengths, and therefore all algorithms result in minimal blocking probability, leading to no gain at all.

3. National Lambda Rail

The National Lambda Rail (as shown in Fig. 16, which was obtained at [11]) is a research optical-network in the United States. Since the NLR is not a regular ring, we have experimented with an absolute *cut-off* value of 12, which is approximately half the length of the ring.

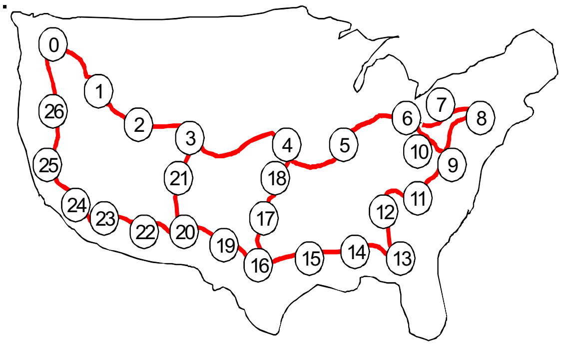

Figure 16. National Lambda Rail.

Figs. 17-19 shows the gain obtained with the Hybrid algorithm when requests are for all source-destination pairs. In general, Hybrid-12 gives positive gains over both the balancing and the concentrating strategies. In the constant case, a negative gain is obtained over balancing, because using the balancing strategy leaves more paths open. Hybrid-12 assigns, in the first pass, requests of length less than and equal to 12, blocking certain paths along all wavelengths. Therefore, with constant number of arrivals, with average requests 10, 12, and 15, negative gains as high as ~25% are observed over the Balancing algorithm.

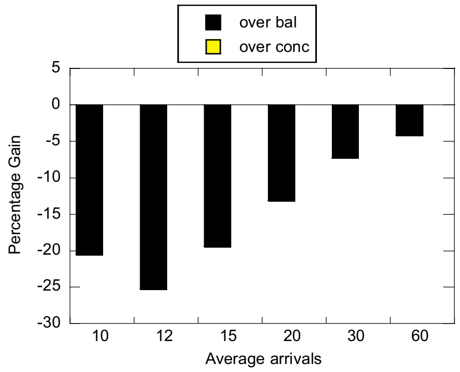

Figure 17. NLR: All possible source-destination pairs - Constant.

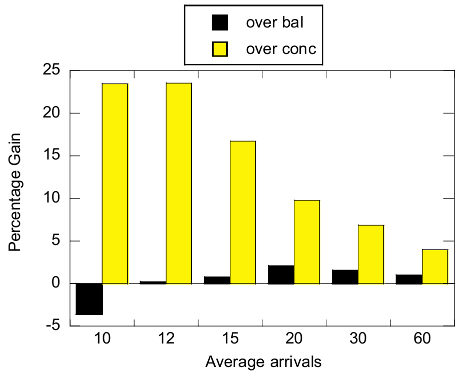

Figure 18. NLR: All possible source-destination pairs - Zipf.

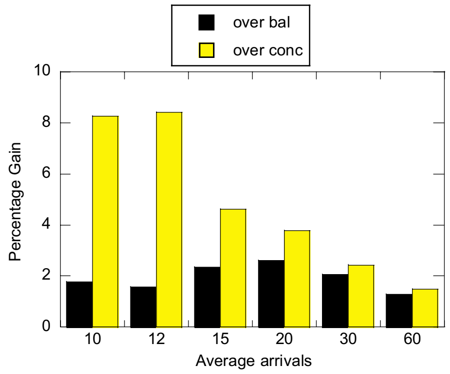

Figure 19. NLR: All possible source-destination pairs - Uniform.

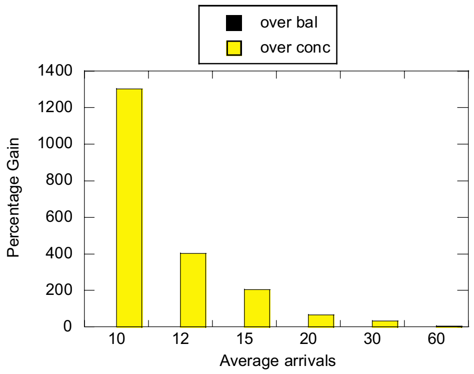

Figure 20. NLR: Traffic within all rings - Constant.

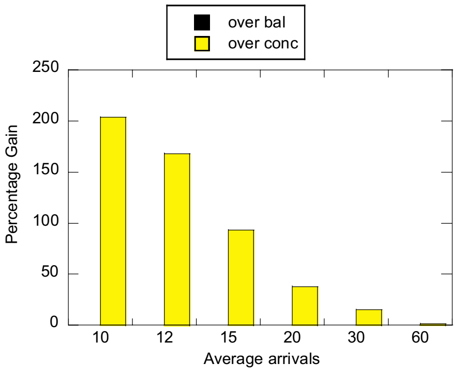

Figure 21. NLR: Traffic within all rings - Zipf.

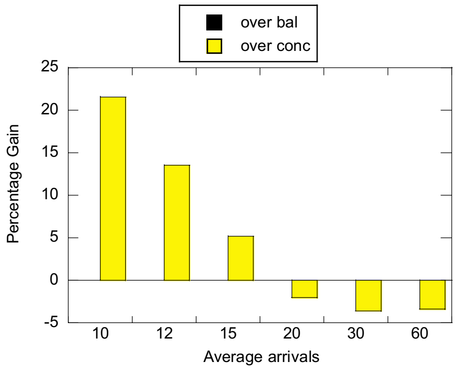

Figure 22. NLR: Traffic within all rings- Uniform.

In Figs. 20-22, we show the effects of assigning equal traffic within individual rings using a *cut-off* value of 4. Note that Hybrid-4 gives a significant positive gain (~1,300%) with constant requests (see Fig. 20) and insignificant negative gain with uniform requests (see Fig. 22) over the concentrating strategy. However, there is no gain over the balancing strategy. We chose the *cut-off* value in this case to be 4, because 4 is close to half the length of the four sub-rings formed by the chords in the NLR topology.

C. Discussion

Our percentage gain calculations show that for the 32-node ring, 32-node ring with chords, and the National Lambda Rail, a significant positive gain (up to ~1,300%) was obtained with the Hybrid algorithm. Note that some negative gains were obtained, in certain cases. However, they are insignificant compared to the positive gain obtained. The behavior of the concentrating, hybrid, and balancing strategies depends on the topology and the specific source-destination pairs requested, and rejecting/selecting edge-disjoint paths accordingly leads to better results. In general, a significant positive gain and an insignificant negative gain over the basic balancing and concentrating strategies imply that the Hybrid algorithm is the right choice for assigning lightpaths in optical networks.

V. Effects of Mean Nodal Degree

We have done extensive simulations with an 8-node ring to study the effect of the nodal degree on the behavior of the Hybrid algorithm. We started with an 8-node ring, and then varied the connectivity of each node from two to seven, by adding chords to the topology. We present the results below, first without wavelength converters, and then with wavelength converters. In both cases, we performed simulations with the following number of wavelengths: 1, 2, 4, 8, 16, and 32. We show the results only with eight wavelengths (using uniform requests with sixty arrivals in a time-slot) due to lack of space (see Figs. 23-25). Observe from Figs. 23-25 that at higher nodal degrees, Hybrid gives positive gains over the Concentrating algorithm both with and without wavelength converters.

A. Results without Wavelength Converters

In Fig. 23, we observe that, with no converters, as the nodal degree increases, the Hybrid algorithm becomes the best option in most cases, except at degree two, where the Concentrating algorithm does better because, in this case, making allocations along shorter edge-disjoint paths leaves more paths open for other allocations

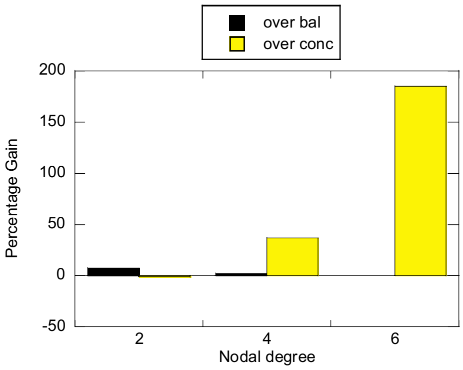

Figure 23. Increasing Nodal Degree: Percentage Gain with Hybrid - no wavelength converters.

The Hybrid and the Balancing algorithms achieve better performance than the Concentrating algorithm as the nodal degree increases, because using the edge-disjoint paths, up to a length of four, leaves more paths open leading to lower blocking probability.

B. Results with Wavelength Converters

In this section, we present results for an 8-node ring with wavelength converters, first by placing wavelength converters at alternate nodes and then by placing wavelength converters at each node.

Converter at Alternate Nodes

In Fig. 24, we observe that the Hybrid algorithm achieves the best performance, and the Concentrating algorithm achieves the worst performance in most cases, except at degree two where it provides the lowest blocking probability. At degree two, it is better to assign the requests along the shortest edge-disjoint paths, rather than making allocations along the longer paths, because the longer path assignment prevents certain source-destination pairs to be allocated. Therefore, the Concentrating algorithm achieves the best performance at degree two in most cases, because this algorithm assigns the requests along the shortest edge-disjoint path first, along all wavelengths.

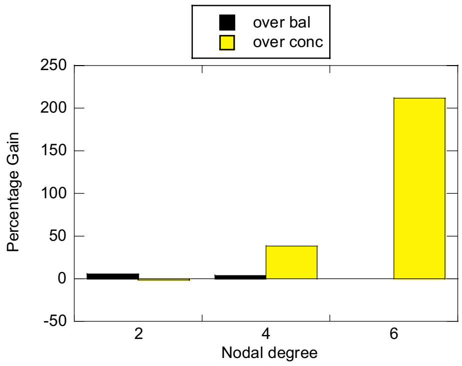

Figure 24. Increasing Nodal Degree: Percentage Gain with Hybrid - wavelength converter at alternate nodes.

Converter at Every Node

In Fig. 25, we observe that as the nodal degree is increases, the Balancing algorithm achieves the best performance in most cases, and the performance of the Hybrid algorithm falls between the performance of the Concentrating and of the Balancing algorithms. The Balancing algorithm achieves the best performance in the presence of wavelength converters because, by balancing along the edge-disjoint paths, it allows all possible combinations of source-destination pairs to be assigned along other wavelengths. The Concentrating and the Hybrid algorithms, by making assignments along shorter paths, along all wavelengths, may prevent certain source-destination pairs to be assigned, thereby leading to higher blocking probability.

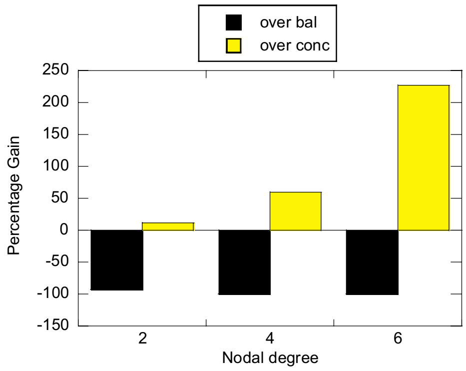

Figure 25. Increasing Nodal Degree: Percentage Gain with Hybrid - wavelength converter at every node.

From the above results, we conclude that the Hybrid algorithm achieves a good performance when no wavelength converters are present, as well as when wavelength converters are present at every other node. Observe that this behavior is depicted in the form of mostly positive gains over both the balancing and the concentrating strategies (see Figs. 23 and 24). When wavelength converters are present at all nodes, the performance of the Hybrid algorithm falls between those of the Concentrating and of the Balancing algorithms. Observe that this behavior is depicted in the form of positive gains over the concentrating and negative gains over the balancing strategies (see Fig. 25).

VI. COMPARISON WITH PREVIOUS WORK

For completeness, we have compared the Hybrid (*cut-off* as half the number of nodes in the topology) algorithm with the “dynamic routing” algorithm proposed in [4], for the on-demand case. In this section, we refer to the “dynamic routing” algorithm described in [4] as h-97. We compared the Hybrid algorithm with the h-97 algorithm, because, in our opinion, the h-97 is currently one of the best heuristic algorithm for routing and wavelength assignment, in spite of its computational complexity, as it selects the route dynamically based on an equation which takes into account the current state of the network:

$$\delta(i,j) = \alpha\beta(i) + (1-\alpha)\Upsilon(j),$$

where $\delta(i,j)$ denotes the value for choosing the i^{th} wavelength on j^{th} route, $\beta(i)$ denotes the number of free links in the entire network corresponding to wavelength i, $\Upsilon(j)$ denotes the number of free wavelengths on route j, and α can take any value between 0 and 1. Note that the route j and wavelength i which has lowest $\delta(i,j)$, is selected first.

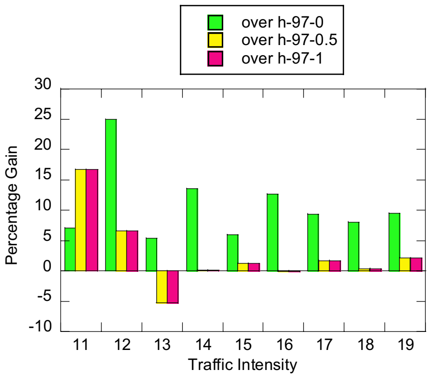

Figure 26. Hypercube - Low Traffic.

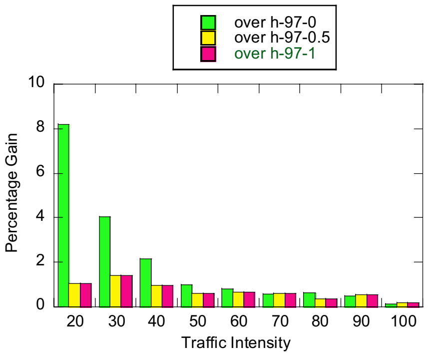

Figure 27. Hypercube - High Traffic.

Our simulation results are presented in this section. Both the inter-arrival and service times follow an exponential distribution. In our simulations, the source-destination requests are uniformly distributed, and the total number of requests in a single trace is 1,000,000. Each request is for a single lightpath. The traffic intensity in our graphs is greater than one, because the source-destination requests are for the entire network, which has four wavelengths in each link. We first compare the algorithms using the hypercube topology, since the hypercube was studied in [4]. For the Hypercube topology, the Hybrid algorithm gives positive gain for most traffic intensities (see Figs. 26 and 27).

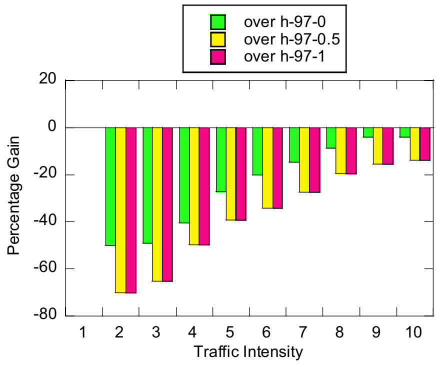

Figure 28. National Lambda Rail - Low Traffic.

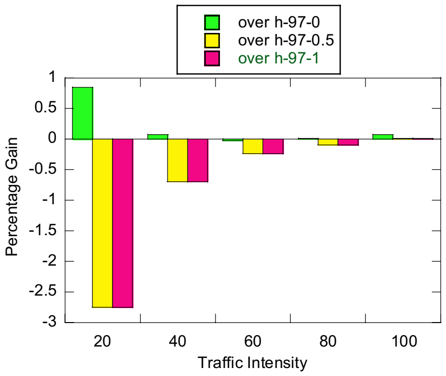

Figure 29. National Lambda Rail - High Traffic.

Observe from Figs. 28 and 29 that, for the National Lambda Rail topology, the Hybrid algorithm gives significant negative gain at low traffic intensities, and negligible negative gain at high traffic intensities.The h-97 algorithm prioritizes the selection or rejection of the edge-disjoint path along a particular wavelength based on both the number of idle wavelengths in the network and the number of free wavelengths on the routes available. Note that it *packs* the wavelengths based on their usage and is a computationally intensive method which requires the update of these two parameters at the end of every assignment.

VII. Conclusion

The Hybrid algorithm *packs* the assignment of wavelengths. It combines the concentrating and the balancing of wavelengths and, in general, gives a significant positive gain over the basic concentrating and balancing approaches. The gain of the Hybrid algorithm over the balancing and concentrating strategies is not a monotonically increasing function, and is low both at low and at high traffic loads. The Hybrid algorithm achieves good performance in the presence of wavelength converters and is comparable to a computationally-intensive algorithm, which makes lightpath assignments dynamically based on the current network state and, therefore, should be used for lightpath assignment in optical networks.

Acknowledgement

We thank Dr. Nirdosh Bhatnagar, from Santa Clara University, for his support. We would also like to thank Dr. Sumit Naiksatam, now at Cisco Systems, for providing the FONTS traces with advance-reservation requests used as input for LRSS. Dr. Naiksatam developed FONTS during his doctoral studies at Santa Clara University.

References

[1] Z. Bian, “Wavelength Assignment Algorithms for WDM Optical Networks”, in *Masters Thesis*, Simon Fraser University, Aug. 2003.
[2] L. Chen and E. Modiano, “Efficient Routing and Wavelength Assignment for Reconfigurable WDM Networks with Wavelength Converters”, in *Proceedings of the IEEE INFOCOM,* pp. 173-186*,* Apr. 2003.
[3] S. Figueira, N. Kaushik, S. Naiksatam, S. Chiappari, and N. Bhatnagar, “Advance Reservation of Lightpaths in Optical Network Based Grids”, in *ICST/IEEE GridNets,* Oct. 2004.
[4] H. Harai, M. Murata, and H. Miyahara, “Performance of Alternate Routing Methods in All-Optical Switching Networks”, in *Proceedings of the IEEE INFOCOM*, pp. 516-524, Apr. 1997.
[5] P. Ho and H.T. Mouftah, “A Novel Distributed Protocol for Path Selection in Dynamic Wavelength-Routed WDM Networks”, in *Kluwer Photonic Network Communications*, Vol. 5, No. 1, pp. 23-32, Jan. 2003.
[6] E. Hyytia and J.Virtamo, “Dynamic Routing and Wavelength Assignment Using First Policy Iteration”, in *Proceedings of the Fifth IEEE Symposium on Computers and Communications*, pp. 146-151, July 2000.
[7] J. Kuri, N. Puech, M. Gagnaire and E. Dotaro, “Routing and Wavelength Assignment of Scheduled Lightpath demands in a WDM Optical Transport Network”, in *IEEE Journal on Selected Areas in Communication,* pp.1231-1240, Vol. 21, No. 8, Oct. 2003.
[8] K. Li, “Performance Evaluation of Heuristic Algorithms for Wavelength Assignment in WDM Optical Networks”, in *Proceedings of the International Conference on Parallel Processing Workshops,* pp. 399-406, Aug. 2004.
[9] S. Naiksatam and S. Figueira, “Elastic Reservations for Efficient Bandwidth Utilization in LambdaGrids”, in *Elsevier’s FGCS - The International Journal of Grid Computing: Theory, Methods and Applications,* Vol. 23, No. 1, pp. 1-22, Jan. 2007.
[10] A. Narula-Tam, P. J. Lin and E. Modiano, “Efficient Routing and Wavelength Assignment for Reconfigurable WDM Networks”, in *IEEE Journal on Selected Areas in Communication,* Vol. 20, No. 1, pp. 75-88, Jan. 2002.
[11] National Lambda Rail: http://www.nlr.net/.
[12] OMNInet: http://www.icair.org/omninet.
[13] R. Ramaswami and K.N. Sivarajan, “Routing and Wave-

length Assignment in All-Optical Networks", in *IEEE/ACM Transactions On Networking*, Vol. 3, No. 5, pp. 489-500, Oct. 1995.

[14] R. Ramaswami and K.N. Sivarajan, "Optical Networks - A Practical Perspective", in *Chapter 8, Morgan Kaufmann*, 1998.

[15] SURFnet: http://www.surfnet.nl/info/en/home.jsp.

[16] S. Xu, L. Li, S. Wang, and C. Chen, "Wavelength Assignment for Dynamic Traffic in WDM Networks", in *IEEE International Conference on Networks,* pp. 375*,* 2000.

Neena R. Kaushik is a Ph.D. candidate at Santa Clara University, California. Her research interests are in queuing theory and scheduling.

Silvia M. Figueira received the B.S. and M.S. degrees in Computer Science from the Federal University of Rio de Janeiro (UFRJ), Brazil, and the Ph.D. degree in Computer Science from the University of California, San Diego. Currently, she is an Associate Professor of Computer Engineering at Santa Clara University. Her research interests are in the areas of parallel and grid computing, particularly in performance modeling, performance prediction, and resource scheduling.

Stephen A. Chiappari chairs the Department of Applied Mathematics in the School of Engineering at Santa Clara University. He received his Ph.D. in mathematics from the University of Illinois in 1990.