On Concentrating Regenerator Sites in ROADM Networks

Balagangadhar G. Bathula[1], Rakesh K. Sinha[2], Angela L. Chiu[2], Mark D. Feuer[2], Guangzhi Li[2], Sheryl L. Woodward[2], Weiyi Zhang[2], Keren Bergman[1], Inwoong Kim[3], and Paparao Palacharla[3]
[1]*Columbia University, New York,* [2]*AT&T Labs – Research, New Jersey,* [3]*Fujitsu Lab of America, Texas*
Email: bbathula@ee.columbia.edu

Abstract: Concentrating optical regenerators in a subset of nodes provides significant savings. We present a heuristic for choosing this subset. In the studied cases our results have 1 or 2 more sites than the minimum.

OCIS codes: (060.4256) Networks, network optimization, (060.1155) All-optical networks

I. Introduction

The advent of Colorless Non-directional Reconfigurable Optical Add/Drop Multiplexers (CN-ROADMs) has significant implications for network design and operation. In a network with CN-ROADMs, spare regenerators can be pre-deployed within the nodes to enable fast provisioning of services and dynamic reconfigurations. They also enable recovery from network failures without the need of manual intervention. This dramatically reduces mean time to repair and provides operational savings as well. These savings can be enhanced if regenerators are only deployed at a *subset* of the network nodes referred as regenerator sites (RS), as concentrating the pre-deployed regenerators in a subset of nodes can potentially reduce the total number of spare regenerators required due to better sharing of spare regenerators from random demands. It also provides operation efficiencies, as fewer truck rolls should be needed. In this paper, a new routing-constrained regenerator location problem is defined. We show that this problem is NP-hard and present a heuristic solution which can be tuned to accommodate different priorities or costs. For example, regenerator sites can be selected to support minimal latency or minimal cost paths as discussed below. In addition to selecting a set of regenerator sites, the heuristic also constructs a lower bound on the size of the optimal set. Our results with various network topologies, reach distances, and cost metrics show that this heuristic gives near-optimal solutions in most studied cases.

II. Problem Definition

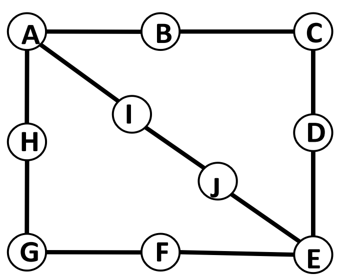

Fig. 1: An Example Network

Regenerator location problems have been studied previously [1,2,3]. The overall goal is to minimize the number of regenerator locations while still being able to route a circuit between any node pair. The problem has been studied in two flavors: (a) the *unconstrained*-routing regenerator location problem (URLP); and (b) the *explicit*-routing regenerator location problem (ERLP). URLP does not limit circuit routing in any way. While the solution may be optimal in number of regenerator locations, individual circuits may incur high cost as a result of using longer routes or using more regenerators. ERLP constrains each route to a specified (typically a min-distance) path, but individual circuits may use more regenerators than necessary. Paper [1] proves that URLP is NP-hard. They also propose and compare three heuristics for URLP. Paper [2] uses a biased random-key genetic algorithm to solve the URLP with improved results. Paper [3] proves hardness of several different variants of the regeneration placement problem and gives approximation algorithms with worst case performance guarantees.

We know that the number of regenerators used by a circuit is a major part of its cost. Neither ERLP nor URLP considered this important metric. Here, we define a new constrained-routing regenerator location problem (CRLP): minimizing the number of regenerator locations with constraints on circuit routing. For example, we can constrain circuits to use only paths with the minimum number of regenerators, or minimum distance paths, or minimum cost paths where the cost is the sum of regenerator and wavelength-km costs. Formally: *we are given a network topology with link distances and a reach distance (maximal optical distance without requiring regeneration). We are also given a set of constraints that restricts the set of paths between any node pair. The goal is to find a minimum set of regenerator sites (RS) such that between each node pair, at least one constrained route is reachable using the regenerators in RS.*

Having a generalized definition of constraints allows us to consider different design priorities and cost metrics. Fig.1 shows a simple network with 10 nodes and 11 links. We assume that the optical reach is 2 hops. For URLP, we can place all regenerators at just three RSs, A, E, I, since a path between any node pair can be constructed using this subset. For ERLP, we specify min-hop paths for each node pair. Because of the strict constraints on routing, we need to place regenerators at five locations: A, C, E, G, I. For CRLP with the constraint of min-regenerator paths, we have a bit more freedom to select routes and we need only four locations: A, J, D, F. This example confirms that URLP, with the most freedom to select routes, attains the fewest RSs, while ERLP, with no freedom, requires more, and the CRLP solution lies in between.

We have proven that CRLP is NP-hard by reduction from a vertex cover problem. Due to space limitation, we omit the hardness proof in this paper. We have also solved this optimally with an integer linear program (ILP) but the ILP ran out of memory on very large networks. In contrast, our heuristic (described below) took under 2 seconds on our largest network and gave near-optimal solutions. We can also adapt the heuristic easily to work with alternative metrics. So while tuning the ILP remains a viable option, we feel that the running time, performance, and the adaptability of our heuristic makes it an attractive approach.

III. Heuristic Algorithm

We give a brief summary of the heuristic for the min-regenerator path constraint. We start by augmenting the network graph by adding edges (i, j), whenever nodes i and j are within reach distance. On this augmented graph G, we can route within any pair of adjacent nodes without requiring regeneration. We can prove a 1:1 correspondence between min-regenerator paths in the network and min-hop paths in G, which transforms the CRLP problem into a graph problem of finding the minimum set of nodes covering min-hop paths between all node-pairs. In each step, the heuristic picks the next best RS according to a rank function on its set of candidate nodes, updates its data structures, and repeats these steps until all source-destination pairs have valid min-hop paths using selected RSs. Let P_1 be the set of node pairs without valid paths that contain v in one of their constrained paths and P_2 be the subset of node pairs in P_1 that acquire a valid path as a result of placing an RS on v. Then we define $rank(v)$ as a weighted combination of $|P_1|$ and $|P_2|$.

We now describe two optimization tricks built into the above greedy heuristic description. (1) For any node pair, (a, z), we have to place an RS on all intermediate nodes in one min-hop (a,z) path. We need to decide which min-hop path will be used. Obviously, if node v is in all min-hop paths between a and z then an RS must be placed on node v. Let us denote this set of nodes as R^+. We seed the algorithm by placing an RS on all nodes in R^+. Moreover R^+ also provides a lower bound, as all its members must be included in any solution. In fact, we can improve this lower bound slightly. Once we have placed all the RSs in R^+, we check if all node-pairs have a valid path using the chosen locations. If not, we know that any valid solution must use at least one more location in addition to R^+. (2) It is possible that nodes selected later may cause some previously selected nodes unnecessary. So we do post-processing to check if we can delete any nodes from the output without affecting the quality of paths. Detailed results of the impact of using R^+, different ranking functions, and post-processing will appear in an extended paper.

The running time without post-processing is $O(n^3 \times (\textit{number of RSs in solution} + \textit{avg degree of } G))$ where n is the number of nodes. The space complexity is $O(n^2)$. The asymptotic running time can be improved by use of more sophisticated dynamic graph algorithms but we chose simple algorithms for ease of implementation. The running times (less than 2 seconds, on a generic PC with 2.3 GHz CPU, for our 75 node topology) suffice for our purpose.

IV. Heuristic results

We evaluated our heuristic on several topologies and we report the results for CONUS topology [4], depicted in Fig. 2. CONUS comprises 75 nodes and 99 links, for a total of 2775 node pairs. Results on other topologies have similar observations. We parameterize the cost of a circuit as $c_r \times \textit{number of regenerations in path} + c_m \times \textit{length of the path}$, where c_r (c_m) are unit regenerator (wavelength-km) costs. Three different sets of parameters were considered: minimum-regenerator route ($c_r = 1$; $c_m = 0$), minimum-distance route ($c_r = 0$; $c_m = 1$), and minimum-cost route (c_r =1000; $c_m = 1$). The last one is the most representative of the network cost. The goal is to minimize the number of RSs such that each node pair is able to pick a path according to the selected constraints.

Fig. 3 compares the number of RSs obtained by our heuristic to the lower bound discussed above. We observe that the heuristic solution is close to the lower bound in *min-regeneration* and *min-cost* cases at all optical reach distances. For the min-distance case, we observed a large gap between our solution and our lower bound. As stated previously, our ILP formulation ran out of memory on large input but we did manage to run it for one reach distance, 2000 km, and found that the heuristic solution was within one RS of the optimal ILP solution.

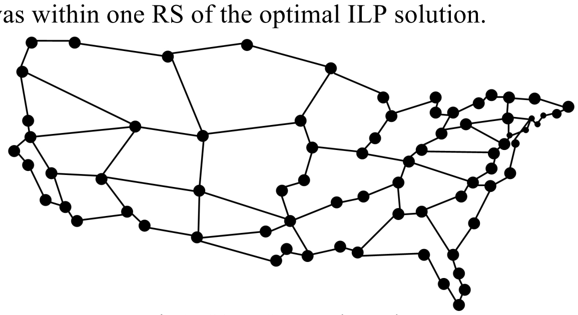

Fig.2: CONUS network topology

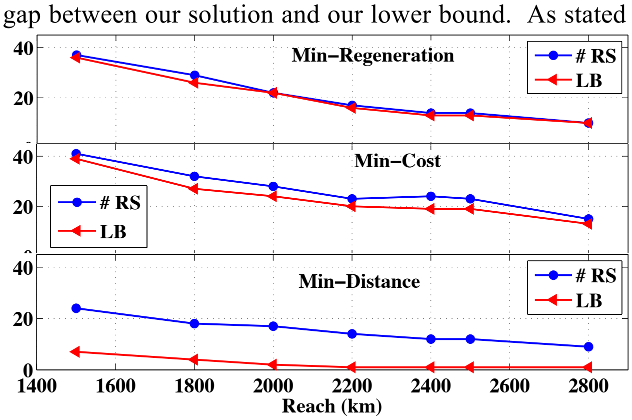

Fig.3: Number of regeneration sites (RS) for different CRLP and their comparison with lower bound (LB).

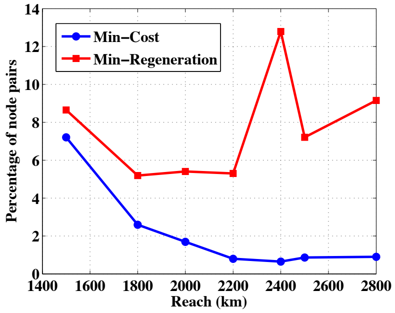

Fig.4: Percentage of node-pair paths with excess length, compared to the shortest-distance path.

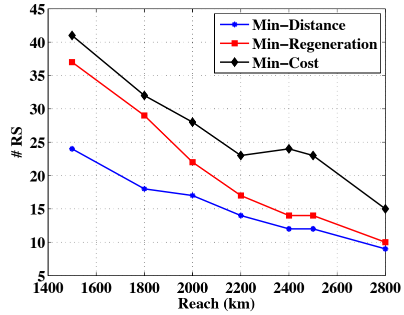

Fig.5: Comparison of number of regeneration sites in the network

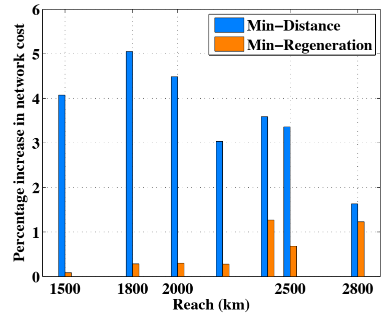

Fig.6: Percentage increase in total cost with one circuit on each route

A min-regeneration path is not necessarily a shortest-path and vice-versa. Thus, routing for min-regeneration results in some paths with longer length, compared to the shortest path for a given node-pair. In Fig.4, we plot the percentage of node-pairs that incur excess length for min-regeneration and min-cost constraints. Min-regeneration case shows consistently higher excess length than min-cost case, especially for a reach of 2400 km or greater. We also generated the distribution of excess length values for the min-regenerator and min-cost cases at reach=2000km. The distance penalty for most of the node pairs was observed to be within 400 km.

In Fig.5, we compare the number of RS obtained for three scenarios discussed above. We observe that the number of RS is the lowest for min-distance case and the highest for min-cost case. However, we have also computed the total network cost (assuming one circuit on each route) for each case, and found that the min-cost case has the lowest network cost, as intended. In Fig.6 we plot the percentage increase in the total network cost (compared to the min-cost case) for the min-distance and min-regeneration constraints. The plotted cost does not include the cost of spare regenerators, which will depend on the individual network operator's practices. Overall, although we see the min-distance case achieves the smallest number of regenerator sites, its total network cost (neglecting spares) is higher (about 5% for reach=1800 km). For the min-regeneration case, the RS set expands, but the network cost is only moderately higher than the min-cost case. For a more complete solution to a real network, one will have to specify the traffic matrix, the spare equipment policy, and other operating practices.

Finding the optical solution using ILP for the CONUS topology is computationally challenging, but our heuristic gives a nearly optimal solution. For all the evaluated topologies, the heuristic is found to give a solution in less than 2 seconds on a 2.3 GHz, Intel Core i5 processor, with 4 GB RAM. We have implemented the heuristic in MATLAB and compared the results with ILP, for both small and large-scale network topologies such as NSFnet, 24-node US network. Our heuristic results match ILP results very well. For most of the compared network topologies, we have exactly the same results.

V. CONCLUSION

To plan pre-deployment of regenerators effectively in practical networks, we have introduced a new constrained-routing regenerator location problem and presented a novel heuristic approach to address it. The heuristic achieves its main objective: to obtain near-optimal regenerator placement solutions with minimal computation time. Results indicate that the heuristic performs well for all three constraints: minimum-distance paths, minimum-regenerator paths and minimum-cost paths.

ACKNOWLEDGEMENT

We gratefully acknowledge many useful suggestions and comments from Robert Doverspike, Pete Magill and Dahai Xu and support from the CIAN NSF ERC (subaward Y503160).

REFERENCES

[1] S. Chen and S. Raghavan, "The regenerator location problem," proceedings of the International Network Optimization Conference, INOC 2007.
[2] Abraham Duarte et al., "Randomized heuristics for the regenerator location problem," Optimization Online, http://www.optimization-online.org/DB_HTML/2010/08/2706.html
[3] Michele Flammini et al., "On the complexity of the regenerator placement problem in optical networks," Proceedings of the twenty-first annual symposium on parallelism in algorithms and architectures, SPAA 2009.
[4] A. L. Chiu et al., "Network Design and Architectures for Highly Dynamic Next-Generation IP-Over-Optical Long Distance Networks," J. Lightwave Technol., 27(12), 1878-1890 (2009).