Trading Network Management Complexity for Blocking Probability when Placing Optical Regenerators

Marcio S. Savasini*, Paolo Monti†, Marco Tacca†, Andrea Fumagalli† and Helio Waldman*
*Optical Networking Lab
Faculdade de Engenharia Eletrica e de Computacao
State University of Campinas, Campinas, SP, Brazil
{savasini, waldman}@decom.fee.unicamp.br
†Open Networking Advanced Research (OpNeAR) Lab
Erik Jonsson School of Engineering and Computer Science
The University of Texas at Dallas, Richardson, TX, USA
{paolo, mtacca, andreaf}@utdallas.edu

***Abstract*—Optical signal regenerators (3R) are required to overcome the adverse effect of fiber and other transmission impairments. 3R units may be placed either at every node (*full placement*) or at some selected nodes (*sparse placement*) of the optical network. It has been argued [1] that while the latter placement strategy may not be optimal in terms of the total number of 3R units required to support a given set of static traffic demands, it offers a number of practical advantages over the former, e.g., a contained complexity of network management in terms of signaling overhead.**

In this paper the full and sparse placement strategies are compared in a dynamic optical network, whereby lightpaths are set up and torn down to best fit the offered changing demands. The study shows that the blocking probability due to the lack of available 3R units achieved by the sparse placement strategy may be comparable to the one achieved by the full placement strategy. Surprisingly, it may even be lower in some cases, thus providing an additional motivation in favor of the sparse placement strategy. The study also shows that the algorithm used to choose the nodes where to place the 3R units must be designed carefully. Two placement algorithms are compared, reporting differences in signaling overhead level as high as 6 times (when achieving a desired level of lightpath connectivity) and differences in blocking probabilities as high as two orders of magnitude (when using the same level of signaling overhead).

***Index Terms*—Blocking Probability, Regenerator Placement, k-connectivity, Network Management, Optical Networks.**

I. Introduction

Optical networks provide high bit-rate transport capabilities to routers and other electronic nodes. Wavelength Division Multiplexing (WDM) technology enables one fiber to carry multiple parallel optical channels [2]. Every time the optical signal is converted to electronics and vice-versa, a transponder is required. Optical circuit (or lightpath [3]) switching has the potential to reduce the necessity for transponders, thus enabling the optical signal to be routed through multiple intermediate nodes without the need of being converted back to electronics. The adverse effect of fiber and other transmission impairments, however, may limit the span of a lightpath (i.e., the number of network elements that can be traversed). As a consequence, a node is able to establish lightpaths to reach only a subset of other nodes. These nodes are said to belong to the node *transparency island* (TI) [4]–[6].

A lightpath not using any regeneration unit can only connect node pairs within the TI boundary. In order for a lightpath to connect two node that are not in each others TI, regeneration units are needed at some intermediate nodes. Optical signal regeneration becomes then a crucial component to circumvent physical impairments in the optical medium. One option for regeneration is reamplification, reshaping and retiming (3R). 3R regeneration may occur in the electronic (OEO conversion) or in the optical domain [7]. Multiple regeneration (or 3R) units may be needed along any given lightpath, one 3R unit every time the quality level does not fulfill the requirements.

The regenerator placement problem has some similarities with the wavelength converter placement problem [8], [9]. Solutions devised to deal with the latter problem may help the former, although they cannot be directly applied in every case. Different strategies for placing the 3R units may be used. The most common strategies deal with minimizing the number of 3R units needed in the network, or with minimizing the number of 3R nodes [10]–[14], i.e., nodes where 3R units are available. A recent work [1] shows that if the network designer places the 3R units on a per-path-basis, for each static lightpath in the network, a smaller number of 3R units is required when allowing any node in the network to carry 3R units, compared to the case where only a subset of nodes is permitted to have 3R units.

However, the same study [1] argues that network management and operation becomes more complex when all nodes are 3R nodes (i.e., *full placement*). For example, limiting the number of 3R nodes (i.e., *sparse placement*) could reduce both the complexity of the routing protocol and the connection

This research was supported in part by NSF Grant No. CNS-043593 and CAPES Process No. BEX4403/05-3.

setup signaling overhead level. In other works [10], [12], [14], a particular emphasis is given on where to place regenerators in order to minimize the number of rejected connections in a dynamic scenario, i.e., lightpaths are setup and torn down to best fit the offered changing demands. These studies show that the algorithm used to choose the position of the 3R nodes may have a significant impact on the value of the network blocking probability.

The objective of this paper is to analyze the network performance, in terms of blocking probability, when the (sparse) regenerator placement strategies introduced in [15] is used in a dynamic scenario. The proposed strategy (k-CD3S) minimizes the number of 3R nodes used in the networks in order to guarantee a given level of connectivity k, i.e., it guarantees the ability to set up a lightpath between any node pair in the network when up to $(k-1)$ 3R nodes are concurrently malfunctioning. In addition, providing a 3R k-connected network may ease the search for an available path upon the arrival of a demand, i.e., a k-connected network guarantees at least k 3R node-disjoint paths between every source-destination pair.

The value of the blocking probability experienced by incoming demands is compared for different values of the 3R node connectivity, i.e., for different values of k. In general, blocking originates from either lack of 3R units or lack of wavelength resources. In order to assess the impact of 3R node placement strategies, blocking due to lack of wavelength resources is set to negligible values by making the number of wavelengths very large. A first set of results compares the k-CD3S placement strategy with a full placement strategy, which allows all nodes to have a regenerator. A second set compares the k-CD3S placement strategy with the sparse strategy presented in [10]. A fair comparison among different configurations is carried out by setting the same regeneration costs: different configurations maintain the total number of 3R units in the network. The larger the number of 3R nodes, the fewer 3R units per node.

Obtained results show that sparse placement of 3R units may not only yield values of blocking probabilities close to those achieved by the full placement strategy, but, in some instances, sparse placement of 3R units may even yield lower blocking probabilities. The study also shows that the algorithm used to choose the nodes where 3R units are placed needs to be carefully designed. When compared to other sparse placement strategies, k-CD3S shows improvement in signaling overhead level as high as 6 times (when achieving a desired level of guaranteed connectivity) and improvements in blocking probability as high as two orders of magnitude (when using the same level of signaling overhead).

II. Regenerator Placement Strategy

This section describes the network model and the 3R node placement strategy (For a more complete description of the placement strategy, the reader is referred to [15]).

Consider a network topology modeled as a graph $G(\mathcal{N}, \mathcal{A})$, where $\mathcal{N}$ is the set of nodes in the network and $\mathcal{A}$ is the set of directed (fiber) links connecting the nodes. Each node is uniquely identified, e.g., node i is denoted as N_i. The link connecting node N_i to N_v is denoted as $l_{(i,v)}$. At a given bitrate, one may find the boundaries of node i TI as:

$$TI_i^{(r)} = \cup_{j \in \mathcal{N}} \left(C_{(i,j)}^{(r)} \right), \tag{1}$$

where $C_{(i,j)}^{(r)}$ is the set of all physical links and optical nodes which can be used to establish a lightpath, with transmission rate r, between N_i and N_j without requiring 3R. Set $C_{(i,j)}^{(r)}$ defines a subgraph of $G(\mathcal{N}, \mathcal{A})$.

Using each node TI connectivity information, a second graph $G'^{(r)}(\mathcal{N}, \mathcal{A}'^{(r)})$, called the connectivity graph, is created. Set $\mathcal{A}'^{(r)}$ is defined as follows:

$$\mathcal{A}'^{(r)} = \bigcup_{i \in \mathcal{N}} \bar{l}_{(i,j)}^{(r)}, \ \forall j \in TI_i^{(r)}. \tag{2}$$

A link $\bar{l}_{(i,v)}^{(r)} \in \mathcal{A}'^{(r)}$ represents a possible lightpaths, connecting N_i and N_v, that can be established without the need of any 3R units. Notice that $\mathcal{A}'^{(r)}$ is a function of the employed transmission rate r. For sake of simplicity, the index r will be dropped in the remainder of the paper, as it is assumed that all lightpaths have the same rate r.

A desired level of network connectivity k can be achieved by forming a k-connected, k-dominating 3R-node set (k-CD3S) of $G'(\mathcal{N}, \mathcal{A}')$. Let $\mathcal{S}_{3R}$ be this set. Then $\mathcal{S}_{3R}$ must meet the following two constraints:

(a) k-dominating constraint: each node that is not in the 3R-node set must be intra-TI connected to at least k 3R nodes, i.e.,

$$\sum_{j \in \mathcal{S}_{3R}} \mathcal{L}_{(i,j)} \geq k, \ \forall i \in (\mathcal{N} \backslash \mathcal{S}_{3R}) \tag{3}$$

where $\mathcal{L}_{(i,j)}$ is a binary variable defined as:

$$\mathcal{L}_{(i,j)} = \begin{cases} 1 & \text{if } \bar{l}_{(i,j)} \in \mathcal{A}' \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

(b) k-node connectivity constraint: subgraph $\tilde{G}(\mathcal{S}_{3R}, \tilde{\mathcal{A}})$, where

$$\tilde{\mathcal{A}} = \{\bar{l}_{i,j} \in \mathcal{A}' \ : \ i, j \in \mathcal{S}_{3R}\} \tag{5}$$

must be k-node connected. A graph is defined to be k-node connected if and only if the removal of any of its $k-1$ nodes does not cause a partition [16].

Set $\mathcal{S}_{3R}$ defines which nodes are 3R nodes. The minimization of the number of those nodes can reduce implementation costs, as well as network management and operation complexity.

A centralized algorithm [15] is used to find a sub-optimal 3R-node placement such that $\mathcal{S}_{3R}$ satisfies both constraints (a) and (b), while minimizing the number of 3R nodes. The algorithm works into two steps. First, an initial solution for the k-CD3S problem is found by selecting a number of nodes that are potential candidates for the 3R-node set. Then a greedy algorithm is applied to prune from the initial set of 3R nodes as many nodes as possible without violating the k-connected, k-dominating constraints on k-CD3S.

III. Routing Algorithm

This section describes the algorithm used to compute a route for all arriving demands. Let $p_{(a,c)}$ be a path from N_a to N_c on graph $G'(\mathcal{N}, \mathcal{A}')$. Let $V_{(a,c)}$ and $L_{(a,c)}$ be the set of nodes and links belonging to $p_{(a,c)}$, respectively. The cost of $p_{(a,c)}$ is defined as follows:

$$Cost_{p_{(a,c)}} = \sum_{\forall d \in V_{(a,c)}} Cost_d + \sum_{\forall \bar{l}_d \in L_{(a,c)}} Cost_{\bar{l}_d}, \quad (6)$$

where

$$Cost_d = \begin{cases} 1 & \text{if } d \in \mathcal{S}_{3R} \wedge f_d > 0 \\ \infty & \text{otherwise,} \end{cases} \quad (7)$$

and

$$Cost_{\bar{l}_d} = \alpha * h_{\bar{l}_d}. \quad (8)$$

The variable f_d represents the amount of available 3R units at d and $h_{\bar{l}_d}$ is equal to the number of physical hops used by the shortest path, computed on graph $G(\mathcal{N}, \mathcal{A})$, between the two edge nodes of link $\bar{l}_d$. The value of α is topology dependent and is computed considering the value of the TI size.

Using (6), each demand is assigned the path with the minimum cost. If two paths use the same amount of 3R nodes, the number of physical links used by each path serves as tie breaker. Once the minimum cost path is found on graph $G'(\mathcal{N}, \mathcal{A}')$, it is easily mapped into a path on graph $G(\mathcal{N}, \mathcal{A})$ based on the assumption of unlimited wavelength resources. If no path is found, the demand is blocked.

IV. Simulation Results

This section presents a collection of blocking probability results that are obtained by means of the k-CD3S placement strategy presented in Section II.

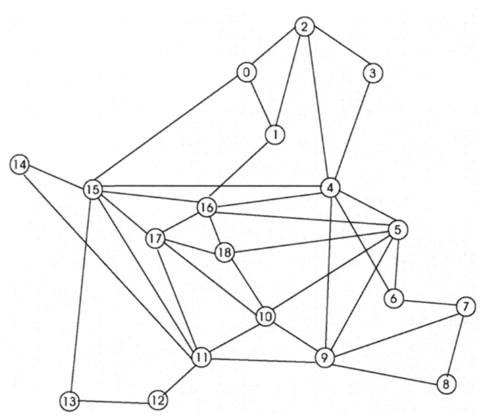

Fig. 1. EON Topology.

Results are found for (i) one regular topology, i.e., a 64-node mesh-torus network, (ii) one irregular topology, i.e., the 19-node European Optical Network (EON) in Fig. 1, and (iii) a set of randomly generated topologies, using the Doar and Leslie's formula [17]. The minimum and average nodal degrees for the random generated topologies are 2 and 3, respectively.

The demand arrivals form a Poisson process with rate λ. Source and destination nodes of each demand are randomly chosen using a uniform distribution over all possible node pairs. Once established, a demand remains in the system for a time that is exponentially distributed with parameter $\frac{1}{\mu} = 1$. It is assumed that the signaling latency in the network is negligible, and the correct network status information is available at all nodes.

To provide results that are not dependent upon any specific call admission control, all arriving demands are first stored in a virtual centralized buffer, as shown in Fig. 2. The buffer can store only one demand at a time. If a demand can be established, it is routed in the network and the buffer slot becomes free. However, if not enough resources are available in the network the demand remains in the buffer. While the buffer is storing a particular demand, all arriving demands are blocked an dropped, until enough resources become available in the network to establish the buffered demand.

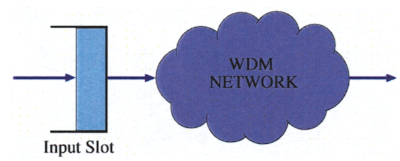

Fig. 2. Input slot buffer.

To provide a comparison benchmark for the k-CD3S placement strategy, results obtained using two additional strategies are also shown. One is a full placement strategy, where each node in the network is a 3R node. The other is a sparse placement strategy called central node first (CNF) [10]. With CNF, each node in the network is ranked based on the frequency it is used by the shortest path between every node pair in the network. The more a node is used, the more central it becomes for the network. 3R nodes are assigned according to the above ranking.

For each experiment, the total amount of 3R units (RU) used in the network is fixed. RU is set to $RU = 2 \times |\mathcal{N}|$. As a result, with the full placement strategy, two 3R units are available at each node. A number of 3R units equal to $\lfloor \frac{RU}{|S_{3R}|} \rfloor$ is assigned to each 3R node. The remaining m 3R units, if any, are allocated according to the following: 3R nodes are ordered according to decreasing nodal degree (calculated on graph $G'(\mathcal{N}, \mathcal{A}')$) and the first m 3R nodes in the list are assigned one additional 3R unit each. In order to give priority to paths using less 3R nodes, the value of α in (8) is set to 0.1.

Unless otherwise specified, blocking probabilities are averaged over 10 different experiments. For each experiment the confidence interval is 10% or better at 95% confidence level.

Fig. 3 shows blocking probability results as a function of the network traffic intensity. Results are obtained using the regular topology. Two placement strategies are compared: k-CD3S and full placement. For k-CD3S different values of the guaranteed level of connectivity k are shown. $k = 12$ corresponds to the case when all nodes are 3R nodes (full placement). In order to have a readable figure, some values of k, i.e., 3, 6, 7, 8, and 10, are not shown.

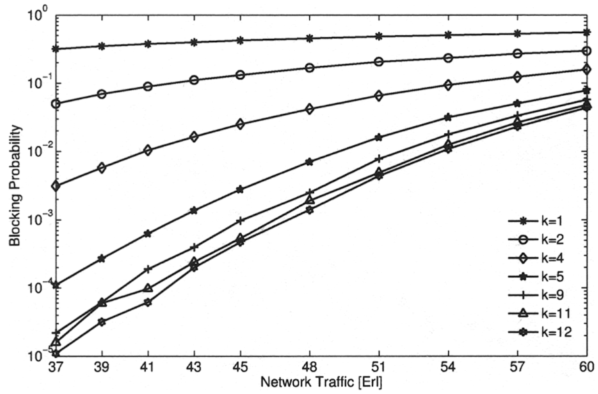

Fig. 3. Torus ($|\mathcal{N}|$ = 64, RU = 128, TI = 2): blocking probability as a function of the network traffic intensity. Two strategies: k-CD3S and full placement (k = 12).

Small values of k lead to high blocking probability levels. This is the case when few 3R nodes have to regenerate (when needed) a large number of demands. As the guaranteed level of connectivity increases, more nodes are chosen to become regeneration points and better performance in terms of blocking probability can be achieved. This is because an increased number of 3R nodes means more path options for each demand.

Although a network with more 3R nodes may show lower blocking probability, it may lead to higher signaling and control messages in the control plane. More places where regeneration can occur means that more messages should be exchanged to advertise regeneration resource information.

TABLE I

TORUS ($|\mathcal{N}| = 64$; $TI = 2$): NUMBER OF 3R NODES AS A FUNCTION OF THE CONNECTIVITY LEVEL (k). TWO STRATEGIES: K-CD3S AND CNF

	# 3R Nodes			# 3R Nodes	
k	k-CD3S	CNF	k	k-CD3S	CNF
1	9	23	7	39	48
2	14	26	8	44	53
3	20	29	9	50	55
4	26	34	10	55	59
5	28	36	11	60	63
6	32	46	12	All Nodes	

Table I shows the number of 3R nodes necessary to achieve a required level of connectivity (k). Results are obtained using the regular topology. Two strategies, k-CD3S and CNF, are considered. The table shows that with the k-CD3S placement strategy only 9 3R nodes are required to guarantee a connected network. The CNF placement strategy requires 23. To guarantee a 2-connected network, 14 and 26 3R nodes are sufficient using the k-CD3S and the CNF strategies, respectively.

Considering that the number of messages in the control plane may grow with $|S_{3R}|^2$, the k-CD3S placement strategy reports differences in signaling overhead level as high as 6 times (when achieving the same level of guarantee connectivity) when compared to the CNF placement strategy. Results in Table I also confirm the earlier claim that sparse 3R placement strategies (e.g., k-CD3S with $k = 9$ and $k = 10$) can achieve blocking probability results comparable with the full placement strategy, but with a reduction in signaling overhead (i.e., the k-CD3S strategy requires up to 20% less 3R nodes).

Fig. 4 shows blocking probability results as a function of the network traffic intensity. Results are obtained using the EON topology. Two placement strategies are compared: k-CD3S and full placement. Contrary to the result presented in the previous plot, the figure shows that when all nodes are regeneration points, the worse performance in terms of blocking probability is achieved.

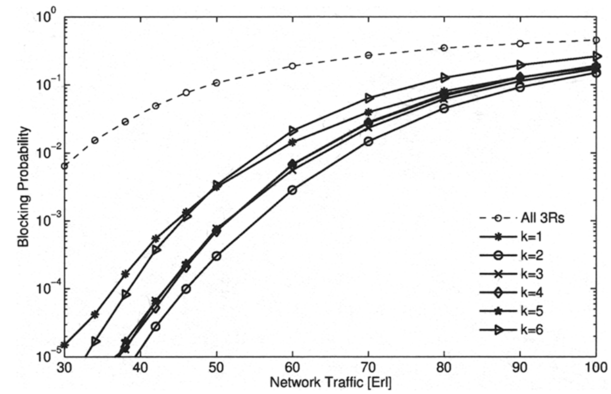

Fig. 4. EON ($|\mathcal{N}|$ = 19, RU = 38, TI = 2): blocking probability as a function of the network traffic intensity. Two strategies: k-CD3S and full placement.

For the EON topology, the best performance in terms of blocking probability are achieved when $k = 2$. This correspond the the case when 3R units are available only in three nodes, i.e., node 4, 9, and 15. With higher connectivity levels, e.g., full placement, regeneration resources are assigned to a larger number of nodes, thus reducing the amount of 3R units per node. Some 3R nodes seem to be less important in the regeneration process, as they are located at the network edges. This result confirms that resources could be better distributed.

TABLE II

EON ($|\mathcal{N}| = 19$; $TI = 2$): NUMBER OF 3R NODES AS A FUNCTION OF THE CONNECTIVITY LEVEL (k). TWO STRATEGIES: K-CD3S AND CNF

	# 3R Nodes - EON	
k	k-CD3S	CNF
1	2 (4 15)	2 (4 15)
2	3 (4 9 15)	3 (4 9 15)
3	5 (4 5 9 11 15)	5 (4 5 9 11 15)
4	6 (4 5 9 11 15 17)	10 (0 2 4 5 6 9 10 11 15 16)
5	8 (4 5 9 10 11 15 16 17)	14 (0 1 2 4 5 6 7 9 10 11 12 13 15 16)
6	11 (0 2 4 5 6 9 10 11 14 15 17)	17 (0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16)

Table II shows the number of 3R nodes necessary to achieve a required level of connectivity (k) when considering the EON topology. Two strategies, k-CD3S and CNF, are considered.

The table shows that, for small values of k, both regenerator placement strategies require the same number of 3R nodes. When the value of k increases, k-CD3S outperforms CNF.

Fig. 5 and Fig. 6 show the performance, in terms of blocking probability, of the two sparse placement strategies (k-CD3S and CNF) when using the same level of signaling overhead, i.e., the same number of 3R nodes.

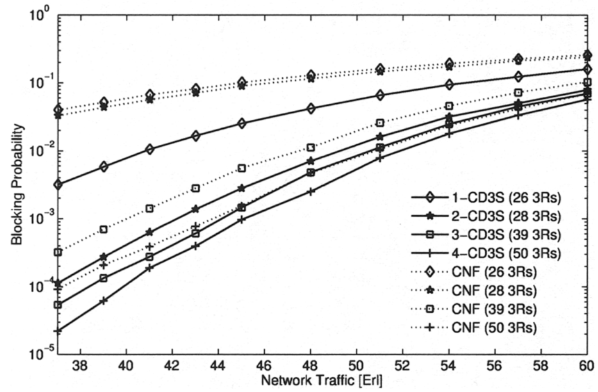

Fig. 5. Torus ($|\mathcal{N}| = 64$, RU = 128, $TI = 2$): blocking probability as a function of the network traffic intensity. Two strategies: k-CD3S and CNF.

Fig. 5 shows the blocking probability results for the 64-node mesh-torus. With less than 23 3R nodes CNF is not able to guarantee connectivity for all pairs. With 26 3R nodes, k-CD3S outperforms more than ten times CNF, while it guarantees $k = 4$. CNF is able to guarantee only $k = 2$. With 28 3R nodes, the difference in performance reach two orders of magnitude, with k-CD3S guaranteeing $k = 5$ and CNF $k = 2$. The performance difference between the two strategies decreases when the number of 3R nodes in the network increases.

For EON topology (Fig. 6), both algorithms perform exactly the same with 2, 3 and 5 3R nodes. For 6, 8 and 11 3R nodes k-CD3S performs ten to hundred times better than CNF.

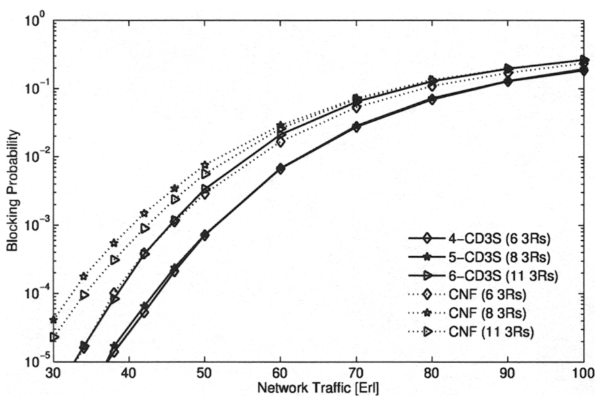

Fig. 6. EON ($|\mathcal{N}| = 19$, RU = 38, $TI = 2$): blocking probability as a function of the network traffic intensity. Two strategies: k-CD3S and CNF.

Table III shows the number of 3R nodes necessary to achieve a required level of connectivity (k) when considering a set of randomly generated topologies with $|\mathcal{N}| = 50$ and $|\mathcal{N}| = 100$. The results are averaged over 20 different topologies for each network size.

For both the 50-node and the 100-node network, k-CD3S requires a smaller set of 3R nodes. On average the difference goes from 50%, for the 50-node topologies, to 80% for the 100-node topologies.

TABLE III

RANDOM TOPOLOGIES ($|\mathcal{N}| = 50, 100$; $TI = 2$): AVERAGE NUMBER OF 3R NODES AS A FUNCTION OF THE CONNECTIVITY LEVEL (k). TWO STRATEGIES: K-CD3S AND CNF

	# 3R Nodes, $\lvert\mathcal{N}\rvert = 50$		# 3R Nodes, $\lvert\mathcal{N}\rvert = 100$	
k	k-CD3S	CNF	k-CD3S	CNF
1	7.4	15.2	15	39.7
2	13.5	21.7	25.1	54.3
3	19.4	32.4	36.7	67.8
4	26.3	40.2	48.5	87.2

Fig. 7 and Fig. 8 show the performance, in terms of blocking probability, of the two sparse placement strategies (k-CD3S and CNF) when using the same level of signaling overhead, i.e., the same number of 3R nodes. Results for the full placement strategy are also reported. The blocking probability values in the figures are obtained using two random topologies, one with 50 nodes and the other with 100 nodes. The two topology are just two examples of the twenty topologies used to generate the numbers in Table III.

For both topologies, the higher the number of 3R nodes, the better the blocking performance. In the 50-node network, when using 7 and 14 3R nodes, CNF is not able to guarantee connectivity to all pairs. The performance of the k-CD3S strategy is comparable with the performance of the full placement strategy, when $k = 2$ or $k = 3$.

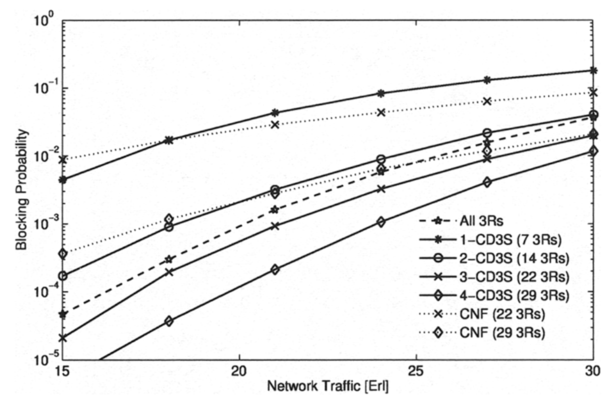

Fig. 7. Random topology ($|\mathcal{N}| = 50$, RU = 100, $TI = 2$): blocking probability as a function of the network traffic intensity. Three strategies: k-CD3S, CNF, and full placement.

In the 100-node network, when $k = 4$ the performance of the k-CD3S strategy is comparable with the performance of the full placement strategy. The CNF placement strategy with

15, 26, and 37 3R nodes is not able to guarantee connectivity to all pairs.

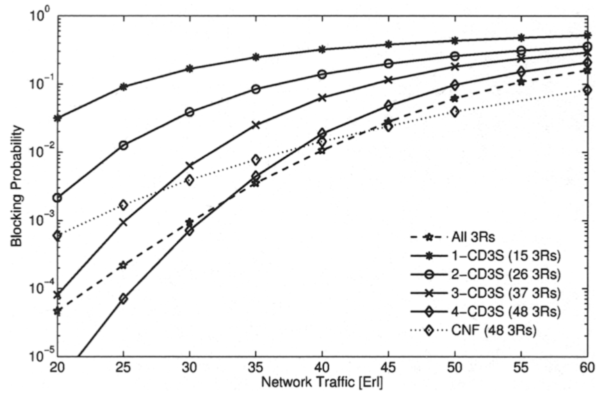

Fig. 8. Random topology ($|\mathcal{N}| = 100$, $RU = 200$, $TI = 2$): blocking probability as a function of the network traffic intensity. Three strategies: k-CD3S, CNF, and full placement.

Note that, in the last two figures, the curves of the CNF placement strategy present a different slope when compared to the ones of the k-CD3S placement strategy. k-CD3S is able to provide more path options for each demand. As a consequence k-CD3S tends to route incoming demands on paths that on average are longer when compared to the path provided by the CNF placement strategy.

V. Conclusion

Restricting the placement of 3R units to a subset of network nodes (sparse placement strategy) has several straightforward advantages, e.g., reduced complexity of the routing protocol and reduced connection setup signaling overhead level. Prior to this study, it was believed that restricting the placement of 3R units to a subset of the network nodes had however a significant drawback, i.e., to support the same amount of static traffic demands more 3R units are required, when compared to the unrestricted placement design (full placement strategy).

Based on an extensive computer simulation campaign, the results presented in this paper seem to indicate that sparse placement of 3R units, to support dynamic traffic demands, may not only yield blocking probabilities values that are close to those achieved by the full placement strategy, but, in some instances, it may also yield lower blocking probabilities. This conclusion seems to add one more advantage to using sparse over full placement strategies.

The second part of the study focused on sparse placement strategies, showing that the placement algorithm must be carefully designed to avoid unnecessary blocking penalties. Two placement algorithms, CNF [10] and k-CD3S [15], were compared. In same cases, k-CD3S showed improvement in signaling overhead level as high as 6 times (when achieving a desired level of guaranteed connectivity) and improvements in blocking probability as high as two orders of magnitude (when using the same level of signaling overhead).

Based on the early encouraging results, further study needs to be carried out. For instance, it will be interesting to investigate a flexible strategy able to determine, for each scenario under consideration, the optimum value of the level of connectivity (k) when the k-CD3S placement strategy is used to support dynamic traffic.

Acknowledgments

The authors would like to thank Dominic Schupke for his valuable input.

References

[1] M. Scheffel, "Regenerator allocation strategies for optical transparency domains considering transmission limitations," in *Proceedings of the IEEE International Conference on Communications (ICC) 2005*, May 2005.

[2] R. Ramaswami and K. N. Sivarajan, *Optical Networks: a Practical Perspective*. Morgan Kaufmann Publishers, 2002.

[3] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high-bandwidth optical WAN's," *IEEE Trans. Commun.*, vol. 40, no. 7, pp. 1171–1182, July 1992.

[4] A. A. M. Saleh, "Islands of transparency - an emerging reality in multiwave optical networking," in *11th Annual Meeting IEEE Lasers and Electro-Optics Society*, 1998.

[5] J. Strand and A. Chiu, "Impairments and other constraints on optical layer routing," Internet Engineering Task Force, RFC 4054, 2005.

[6] S. Das, R. R. Tabrizi, P. Monti, M. Tacca, and A. Fumagalli, "A Link State Advertisement Protocol for Optical Transparency Islands," in *Proc. IEEE Workshop on High Performance Switching and Routing (HPSR)*, May-Jun 2007.

[7] H.-P. Nolting, "All-optical 3R-regeneration for photonic networks," in *ONDM 2003: Proceedings of the 7th IFIP Working Conference on Optical Network Design & Modelling*, 2003.

[8] S. Subramaniam, M. Azizoglu, and A. Somani, "On optimal converter placement in wavelength-routed networks," *INFOCOM*, vol. 1, pp. 500–507, Apr. 1997.

[9] X. Chu, B. Li, and I. Chlamtac, "Wavelength converter placement under different RWA algorithms in wavelength-routed all-optical networks," *IEEE Trans. Commun.*, vol. 51, no. 4, pp. 607–617, Apr. 2003.

[10] X. Yang and B. Ramamurthy, "Dynamic routing in translucent WDM optical networks: the intradomain case," *J. Lightw. Technol.*, vol. 23, no. 3, pp. 955–971, Mar. 2005.

[11] D. Schupke, C. Eder, T. Blankenhorn, and S. Duhovnikov, "A regenerator pool dimensioning method to avoid early blocking," in *IEEE Workshop on High Performance Switching and Routing (HPSR)*, May 2005.

[12] G. Shen, W. D. Grover, T. Hiang Cheng, and S. K. Bose, "Sparse placement of electronic switching nodes for low blocking in translucent optical networks," *Journal of Optical Networking, vol. 1, Issue 12, p.424*, vol. 1, Dec. 2002.

[13] E. Yetginer and E. Karasan, "Regenerator placement and traffic engineering with restoration in GMPLS networks," *Photonic Network Communications*, vol. 2, no. 6, pp. 139–149, September 2003.

[14] S.-W. Kim, S.-W. Seo, and S. C. Kim, "Regenerator placement algorithms for connection establishment in all-optical networks," in *IEEE Global Telecommunications Conference*, vol. 2, Nov. 2000, pp. 1205–1209.

[15] M. S. Savasini, P. Monti, M. Tacca, A. Fumagalli, and H. Waldman, "Regenerator placement with guaranteed connectivity in optical networks," in *ONDM 2007: 11th International Conference on Optical Networking Design and Modeling*, May 2007, pp. 438–447.

[16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms (Second Edition)*. Cambridge, MA: MIT Press, 2001.

[17] M. Doar and I. M. Leslie, "How bad is naive multicast routing?" in *INFOCOM*, vol. 1, March/April 1993, pp. 82–89. [Online]. Available: citeseer.ist.psu.edu/doar93how.html