Grooming Algorithms

Mohammad Hadi

mohammad.hadi@sharif.edu

@MohammadHadiDastgerdi

Spring 2022

Overview

End-to-End Multiplexing

End-to-End Multiplexing

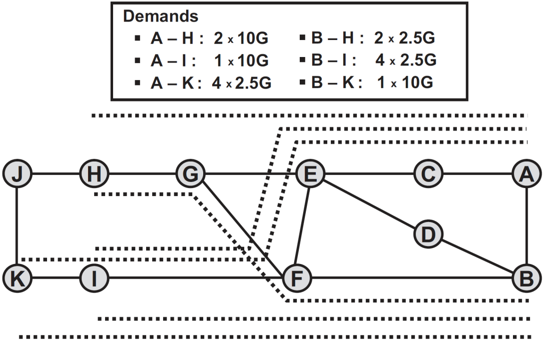

Figure: The line rate is 40 Gb/s and the sub-rate demands are as shown. With end-to-end multiplexing, demands with the same source and destination are bundled together. Six wavelengths with the overall utilization of 27% are required to carry the traffic in this example, as shown by the dotted lines. Without multiplexing, the utilization is 12%. The multiplexing function is most commonly accomplished via a WDM transponder equipped with multiple client-side feeds, known as muxponder. The multiplexed traffic streams should preferably have the same QoS and protection level. The end-to-end multiplexing can be performed using first fit decreasing bin-packing methodology, where the sub-rates are sorted from highest to lowest bit-rate and each sorted sub-rate is multiplexed into the lowest numbered line-rate bin that still has room for it.

Grooming

End-to-End Multiplexing

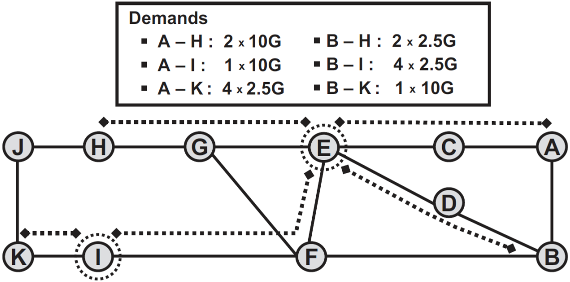

Figure: The line rate is 40 Gb/s and the sub-rate demands are as shown. Grooming is used, with intermediate grooming at Nodes E and I, to pack the wavelengths more efficiently. Five grooming connections are formed as shown by the dotted lines. The grooming design requires five connections in contrast to the six connections required for multiplexing. Second, the groomed wavelengths are on average 75% filled, in comparison to 27% for the multiplexed wavelengths. Finally, in terms of wavelength-link units, grooming requires 9 units whereas multiplexing requires 21 units.

Grooming-Node Architecture

Grooming Switch at the Nodal Core

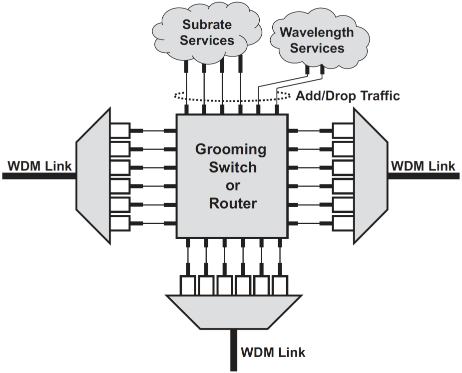

Figure: Grooming switch or router at the nodal core. All traffic entering the node, whether transiting traffic or add/drop traffic, is processed by the grooming switch or router. This architecture is similar to the O-E-O architecture and has all the attendant scalability issues on cost, physical size, power consumption, and heat dissipation. The grooming switch can be a SDH/OTN switch providing circuit switch-based grooming or an IP router providing packet switch-based grooming.

Grooming Switch at the Nodal Edge

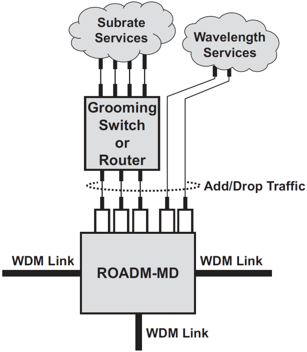

Figure: Grooming switch or router deployed at the nodal edge with a wavelength level switch at the core. The wavelength-level switch can provide optical bypass. Only the subrate services that need to be groomed at the node or that need to be added/dropped at the node are processed by the grooming switch or router, yielding a more scalable architecture. The wavelength services originating at the node feed directly into the core switch. If the ROADM is a rigid non-direction-less, a small wavelength-level edge switch can improve the architectural flexibility.

Grooming Switch at the Nodal Edge

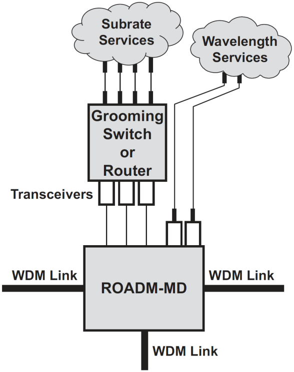

Figure: WDM transceivers can be integrated on the grooming switch or router to improve cost efficiency. This architecture may lead to compatibility issues between the ROADM and grooming switch.

Intermediate Grooming Layer

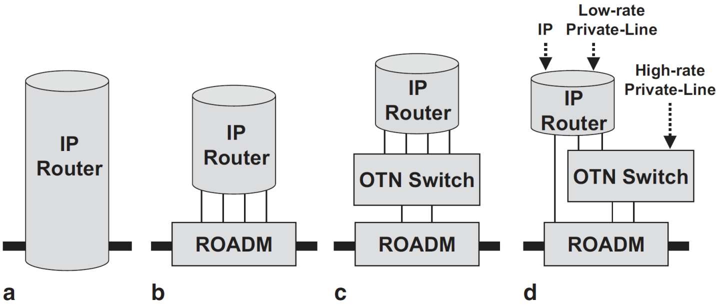

Figure: a All traffic is processed by the IP router. b A ROADM allows optical bypass of the IP router. c Much of the grooming is moved to an intermediate OTN switch. d Improved layering to reduce the amount of traffic passing through both an IP router and an OTN switch. From (a) to (d), the required size of the IP router decreases relative to the amount of network traffic.

Intermediate Grooming Layer

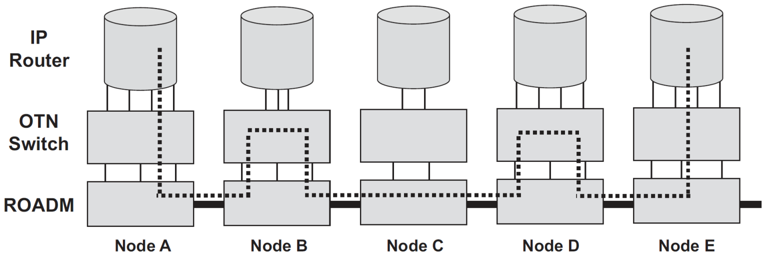

Figure: A connection passes through the IP Router, the OTN switch, and the ROADM at the ingress and egress points. At Nodes B and D, the OTN switch is used to bypass the IP router. At Node C, the ROADM is used to bypass both the OTN switch and the IP router. The main motivations for this architecture are reduced cost and power consumption. For OTN, the switching granularity of the intermediate layer can be as fine as an ODU0 (i.e., 1.25 Gb/s).

Intermediate Grooming Layer

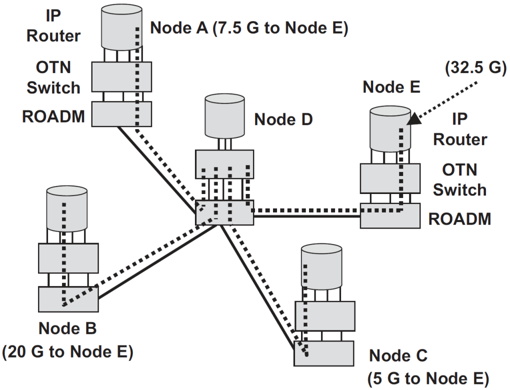

Figure: Assume that the line rate is 40 Gb/s. Nodes A, B, and C exchange less than a wavelength's worth of traffic with Node E. However, by using the OTN switch at Node D to groom this traffic together onto a single wavelength, Nodes A, B, and C can establish efficient IP adjacencies with Node E, which in turn improves jitter, delay, and processing constraints. One disadvantage of the intermediate sub-wavelength architecture is the added complexity and cost of having to manage another layer.

Selection of Grooming Sites

Selection of Grooming Sites

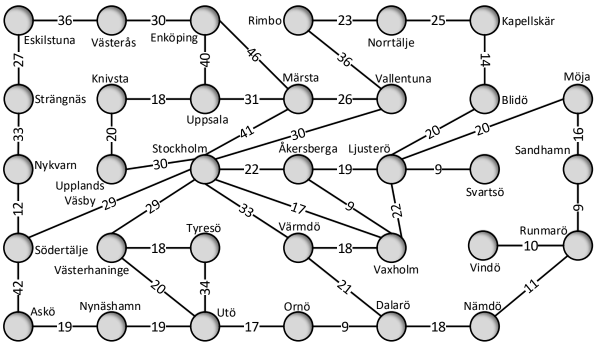

Figure: For economic reasons, every network node is not equipped with grooming switch. The nodes without a grooming switch typically must backhaul their subrate traffic to nearby grooming nodes. Number of grooming sites creates a tradeoff between routing complexity and grooming switch utilization. Subrate traffic source, heavily trafficked nodes, nodal degree, regeneration sites, proximity of the non-grooming nodes affect the selection of grooming site.

Hierarchical Grooming

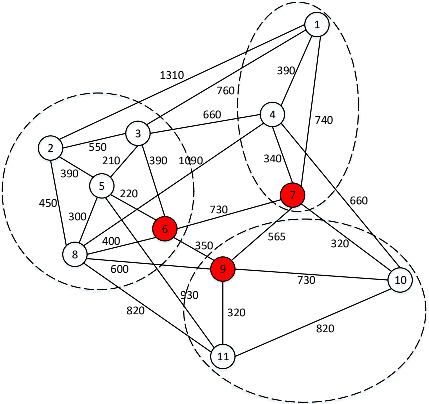

Figure: Selection of grooming sites implicitly establishes a grooming hierarchy in the network, where the network nodes are partitioned into clusters, with one node in each cluster selected as the hub. The bulk of the inter-cluster subrate traffic is first directed to the hub corresponding to the source node's cluster; this traffic is then routed to the hub corresponding to the destination node's cluster and from there to the ultimate destination. Most of the intra-cluster subrate traffic is groomed in the hub as well. The hub-selection scheme can be based on K-Center approach or link capacity approach. tThe objective in the K-center problem is to find a set of K center nodes in a graph that minimizes the maximum distance between any non-center site and the nearest center site.

Backhaul Strategies

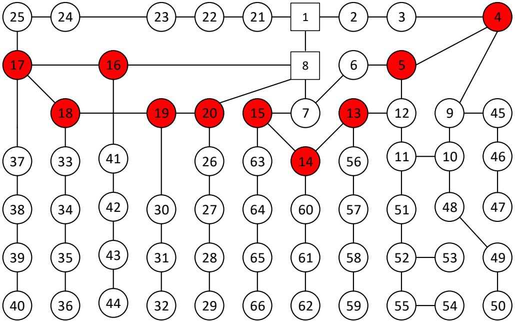

Figure: If only a subset of the network nodes are equipped with a grooming switch, then the remaining nodes with subrate traffic either use end-to-end multiplexing to carry their subrate traffic or they backhaul their subrate traffic to a grooming node. If the backhauling is used, the non-grooming node is said to "home" on a grooming node; the grooming node is referred to here as the "parent" node. There are several criteria that may be used to determine on which node a non-grooming node should home such as backhaul distance, grooming switch size, and reliability.

Backhaul Strategies

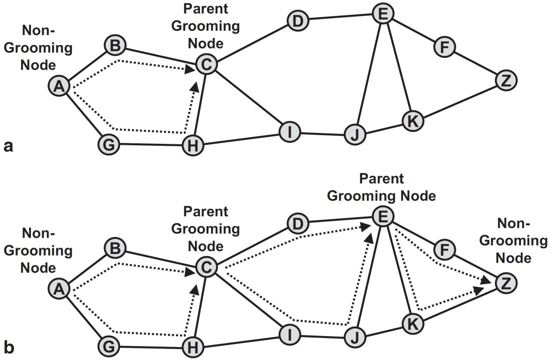

Figure: Node A, which is not equipped with grooming equipment, homes on just a single grooming node, Node C. b An end-to-end path from Node A to Node Z, both of which home on just a single parent grooming node, is protected against failures except at Nodes C and E.

Backhaul Strategies

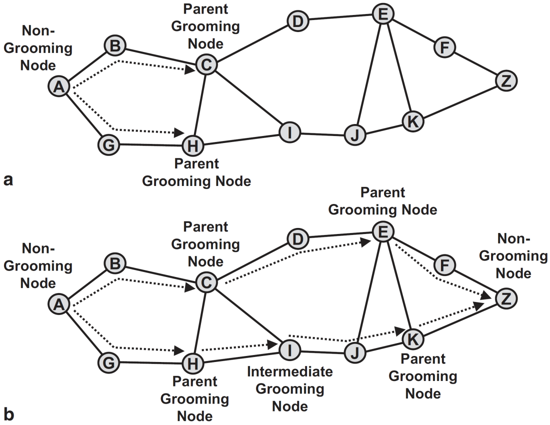

Figure: a Node A homes on both Nodes C and H and delivers its subrate traffic to both parents over diverse paths. b An end-to-end protected path from Node A to Node Z. There are no single points of failure along the path. It is assumed that Node I is used for intermediate grooming.

Grooming Trade-offs

Cost Versus Path Distances

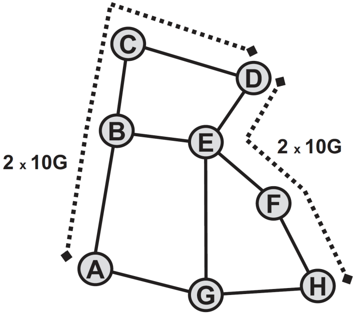

Figure: The line rate is 40 Gb/s. One wavelength between Nodes A and D carries two 10 Gs. One wavelength between Nodes D and H also carries two 10 Gs. If a new 10 G demand is added between Nodes A and H, it can potentially be carried in the existing wavelengths, rather than establishing a new connection directly along A–G–H. This is the lower cost option (at least in the short term) but utilizes a longer path that burns future capacity. Further, this option is somewhat more vulnerable to failure with respect to the new demand and will result in greater latency and jitter.

Cost Versus Capacity

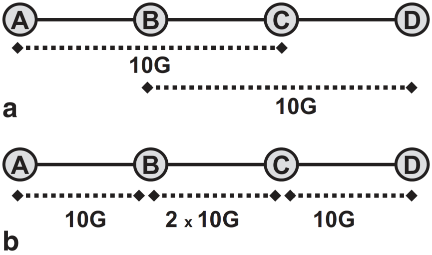

Figure: The line rate is 40 Gb/s. Two 10 G demands are added, one between A and C, and one between B and D. a The two 10 Gs are carried in separate connections between their respective endpoints. This option requires four grooming ports and utilizes two wavelengths on Link BC. b Nodes B and C are used to groom the traffic such that there is a single connection between B and C carrying two 10 Gs. This option occupies just one wavelength on Link BC, but it requires a total of six grooming ports. Here, the extra grooming potentially reduces the reliability of the circuits. Additionally, if the intermediate grooming occurs in an IP router, then there may be additional latency or jitter.

Cost Versus Protection Capacity

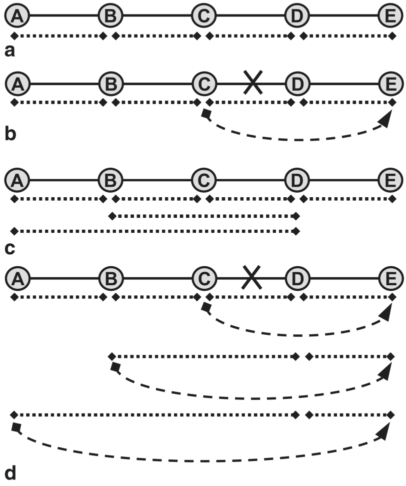

Figure: a IP links, shown by the dotted lines, are established only between neighboring routers using grooming. b When physical Link CD fails, only one IP link fails. The restoration path (in the A to E direction) is represented by the dashed line. c Two express IP links are created, both of which bypass the router at Node C. d When physical Link CD fails, three IP links fail, requiring three different restoration paths.

Grooming Strategies

Initial Bundling and Routing

1. Group the new subrate demands into bundles that contain at most one wavelength's worth of traffic using first fit decreasing bin packing, where all of the demands in a given bundle have the same source and destination nodes.
2. The bundled demands should be compatible in terms of QoS requirements, protection level, ...
3. Route the bundles using the standard routing techniques, e.g., alternative-path routing.
4. Network load can be used as one criterion for selecting which alternative path to select for a given bundle.
5. If one or both of the bundle endpoints do not have grooming capability, then the path chosen must pass through the appropriate "parent nodes" that do have grooming switches.
6. No further grooming operations need to be done with the bundles that may contain a full wavelength, or very close to a full wavelength, of traffic.
7. At this point, a set of grooming connections is obtained, where a grooming connection is a path terminated at both ends on a grooming switch.

Grooming Operations

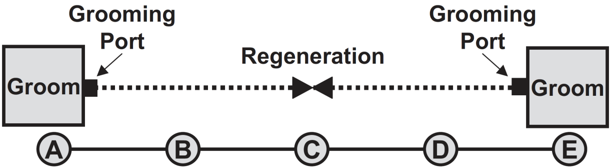

Figure: The grooming connection (GC) represented here by the dotted line terminates on the grooming switches at Nodes A and E. A GC may need to be regenerated, as shown at Node C. The intermediate nodes, B, C, and D, may contain grooming switches as well; however, this GC is not processed by them. There are two types of GCs. First, the existing GCs encompass those GCs that have already been established in the network. Second, there are new GCs, formed from routing the bundles containing the new subrate demands. There is more flexibility with the new GCs. If the grooming switch supports a "make-before-break" feature, then moving or rearranging existing GCs may be possible without bringing down the corresponding demands.

Grooming Operations

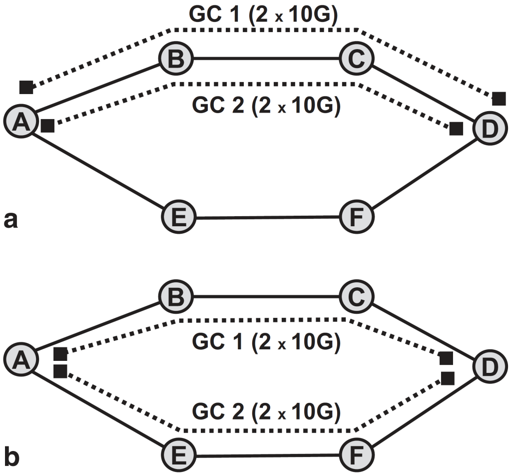

Figure: Combining GCs with the same endpoints to allow one of the GCs to be removed. a The two GCs have the same path. The two 10 Gs from GC 1 can be moved to GC 2. b The two GCs have different paths. The two 10 Gs from GC 1 can be moved to GC 2, assuming this new path is satisfactory for the demands in GC 1.

Grooming Operations

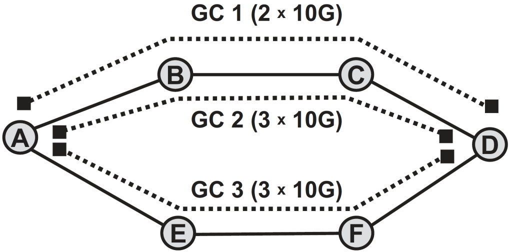

Figure: One 10 G from GC 1 (2 x 10G) grooming connection (GC) 1 can be moved to GC 2 and the other 10 G can be moved to GC 3, allowing GC 1 to be removed.

Grooming Operations

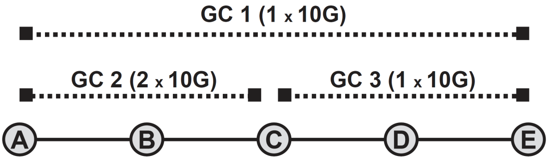

Figure: The demands from grooming connection (GC) 1 can be moved into both GC 2 and GC 3. This adds another grooming point for the demands in GC 1.

Grooming Operations

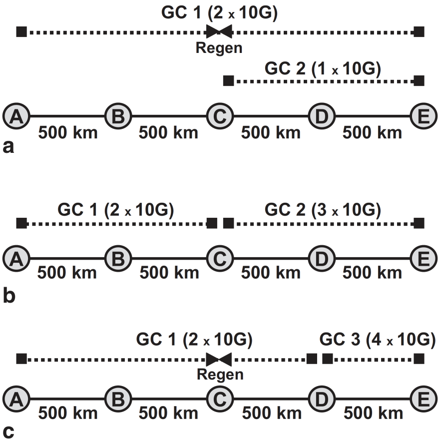

Figure: a In the original setup, two wavelengths are utilized along C-D-E. b After grooming connection (GC) 1 is shortened and its demands also placed in GC 2, only one wavelength is required along C-D-E. Assuming an optical reach of 1,000 km, the regeneration at Node C is removed. c If GC 1 had been terminated at Node D instead of Node C and the demands carried in both GC 1 and GC 3, one regeneration would still be required.

Grooming Operations

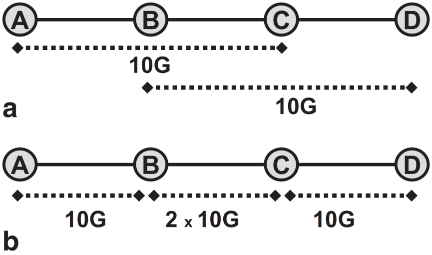

Figure: The line rate is 40 Gb/s. Two 10 G demands are added, one between A and C, and one between B and D. a The two 10 Gs are carried in separate connections between their respective endpoints. This option requires four grooming ports and utilizes two wavelengths on Link BC. b Nodes B and C are used to groom the traffic such that there is a single connection between B and C carrying two 10 Gs. This option occupies just one wavelength on Link BC, but it requires a total of six grooming ports. Here, the extra grooming potentially reduces the reliability of the circuits. Additionally, if the intermediate grooming occurs in an IP router, then there may be additional latency or jitter.

Grooming Performance

Grooming Operations

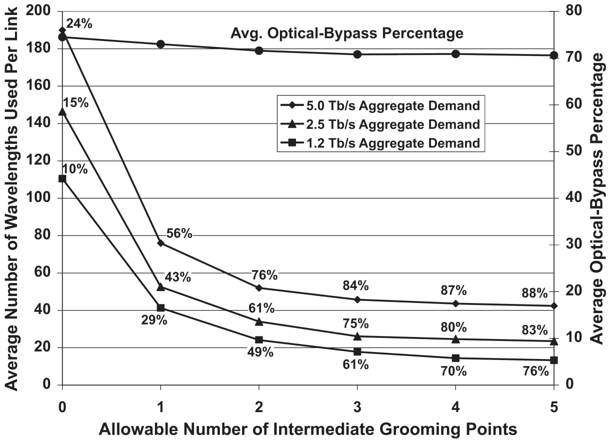

Figure: Average link utilization (with 1.2, 2.5, and 5 Tb/s of aggregate demand) and average optical-bypass percentage as a function of the allowable number of intermediate grooming points per demand when all the nodes have grooming switch. The percentages specified next to the data points are the average fill rates of the resulting GCs. The optical-bypass curve is for the 5 Tb/s scenario; the percentages were slightly lower for the other scenarios. All 60 network nodes were equipped with grooming switches in this scenario.

The End