Simulations of a Service Velocity Network Employing Regenerator Site Concentration

Mark D. Feuer[1], Sheryl L. Woodward[1], Inwoong Kim[2], Paparao Palacharla[2], Xi Wang[2], Daniel Bihon[3], Balagangadhar G. Bathula[4], Weiyi Zhang[1], Rakesh Sinha[1], Guangzhi Li[1] and Angela L. Chiu[1]
[1]AT&T Labs - Research, Middletown & Florham Park, NJ, USA (mdfeuer@att.com)
[2]Fujitsu Labs of America & [3]Fujitsu Network Communications, Richardson, TX, USA, [4]Columbia Univ., New York, NY, USA

Abstract: We use Monte Carlo simulations of quasi-static traffic growth to study concentration of regenerator sites in dynamic photonic networks supporting service velocity. Idle regenerators can be reduced by >60% compared to unrestricted regeneration sites.

OCIS codes: (060.4265) Networks, wavelength routing; (060.1155) All-optical networks; (060.2330) Fiber optics communications; (060.2360) Fiber optics links and subsystems

1. Introduction

Dynamic photonic networks, in which lightpaths are rapidly set up and torn down on demand, have been widely researched [1,2]. However, the transition from today's static networks, in which manual wavelength setup takes many weeks and circuits remain in place for years, has been stymied by a range of technology and cost issues. Practical ROADM designs that provide colorless and non-directional add/drop capability have been published [3], and hardware is becoming commercially available, but challenges remain in operational practices and business models. Recently, we introduced the concept of service velocity (SV) which ensures rapid setup of lightpaths in quasi-static networks by pre-deploying optical regenerators for all possible circuit requests longer than the system's all-optical reach [4]. Pre-deployed equipment necessarily spends the first part of its life idle (and thus without a supporting revenue stream), so it is critical to minimize the number of idle regenerators, while maintaining sufficient resources to meet service velocity goals. A further complexity is that practical ROADMs may have potential wavelength contention among different sets of regenerators defined by their connection paths to the ROADM core. In [4], we used Monte Carlo methods to simulate traffic growth in a realistic continental-scale backbone network that supports service velocity. With appropriate algorithms to determine the timing and location of regenerator pre-deployment, we observed that the number of idle regenerators remained nearly constant as traffic grew, at about 2 per network node. Furthermore, we found that by using a constraint-aware algorithm for wavelength assignment, blocking due to intra-node contention could be completely avoided, even when traffic demand became so large that double overbuilds of busy fiber routes were required.

In this paper, we study the effect of restricting regeneration to a limited subset of sites within the same continental-scale network. By concentrating the regeneration at a limited number of nodes, we expect to reduce the number of idle regenerators, since each idle regenerator becomes available for use in a larger number of source-destination routes. Limiting regeneration sites could also reduce operating cost, by reducing the number of large offices and reducing the amount of travel needed by operations personnel. On the other hand, the add/drop load at each selected site is greatly increased, so it seems possible that intra-node wavelength contention could become an issue. Here, we quantify these costs and benefits of regenerator site concentration by conducting a Monte Carlo simulation of traffic growth in a quasi-static network, choosing regenerator site sets based on either shortest-path routing or minimum-regenerator routing. We compare the network blocking, the idle regenerator counts, and the active regenerator counts with results from simulations that did not attempt to consolidate regeneration sites.

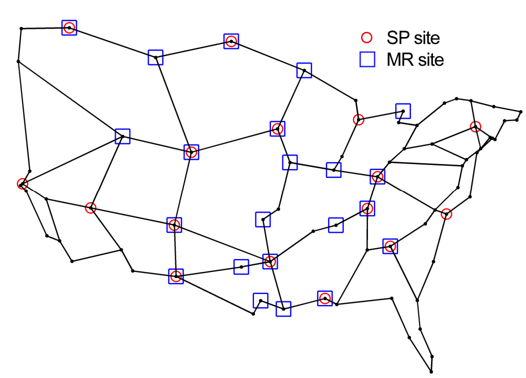

Fig. 1. CONUS network topology. Different markers identify regeneration sites included in the various subsets under study.

2. Methodology

The network is made up of colorless, non-directional ROADM nodes of the design shown in [3], with the addition of dedicated regenerators as shown in [4]. The regenerators lack client interfaces, and therefore cannot be routed through the client-side crossconnect. This ROADM design scales gracefully over a wide range of sizes and can be upgraded while in service. The network topology under study is the continental United States segment of the topology developed for the DARPA

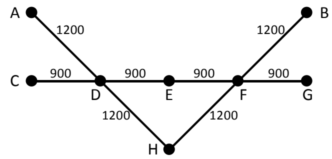

Fig. 2. Simple network example for a 2000 km reach. The SP method will select nodes D & F as regeneration sites, so two regenerators will be needed on path C-G.

CORONET program [5]; we refer to it as the CONUS topology. The CONUS topology, shown in Fig. 1, has 75 nodes and 99 links, using actual city locations. We have simulated transmission systems that support 88 wavelengths with an all-optical reach of 2000 km.

Traffic in the simulation is quasi-static: connections appear at random, but once established, they remain in place through the end of the simulation run. To obtain results that are more realistic than those for uniform traffic, we have adopted a traffic matrix in which the population of each city determines its probability to be used as a source or destination. To meet the needs of our proposed SV mode of operation, two regenerators are pre-deployed at each regeneration-capable node at the beginning of the simulation. As each demand is received, these are put into use as needed and a special SV algorithm checks to see whether more regenerators may be needed to sustain future demands. In the present work, we used a global-constraint aware SV (GC-SV) algorithm, which assures resources for any next-demand possibility. If needed, regenerators are added (instantly, within this simulation) in groups of at least 4 to minimize truck rolls. In our earlier study [4], we found that fiber exhaust could distort our traffic pattern after relatively few demands, because of congestion on a few node-to-node fiber links, so the simulation also includes the overbuild of individual fiber links when they exceed 60% wavelength fill, up to a maximum of three fiber pairs per link.

We have selected two distinct regenerator concentration strategies for study: (a) all connection paths were fixed to the shortest-distance paths and integer linear programming (ILP) was used to find the smallest set of regeneration sites that could support those shortest paths; or (b) a heuristic algorithm searched for the smallest set of sites which allows every connection path to be routed by a minimal–regenerator path, then each connection was routed by the shortest min-regen path available with the chosen subset of regenerator sites. Option (a), which we call shortest-path routing (SP), assures minimum latency, an important consideration for some customers, while option (b), which we call minimal-regeneration routing (MR), is a first-order approximation to least hardware cost. Details of the site selection algorithms will be published separately [6]. For comparison, we also simulated a simple strategy with unrestricted regeneration (any site allowed) and shortest-path routing.

Although limiting regeneration sites reduces the number of *idle* regenerators, it can actually increase the number of *active* regenerators if SP site minimization is used, as shown in the example of Fig. 2. Consider a system reach of 2000km. If regeneration sites are unrestricted, the route C-G will be served with a single regenerator at E. However, if SP site minimization is applied, regeneration will be limited to nodes D and F, and the path C-G will require two regenerators. If MR site minimization is used, the number of regenerators will never increase, though some paths may become longer than they would be with shortest-path routing.

3. Results and Discussion

The SP site subset comprises 16 nodes and the MR site subset comprises 22 nodes, as identified in Fig. 1. The mandatory routing of the SP strategy leads to a rather uniform distribution of sites, spaced at nearly constant intervals. In contrast, the MR algorithm allows multiple alternative routes to be used (so long as no added regenerations are needed), resulting in a set of sites that are centrally clustered along heavily-used links.

The total blocking rate due to all causes is plotted in Fig. 3. No demands are blocked until traffic exceeds 900

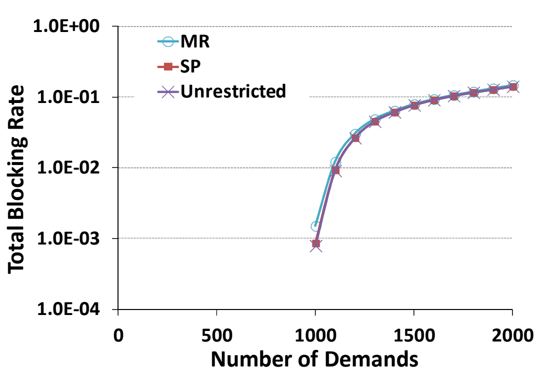

Fig. 3. Total blocking rate vs. traffic demand. Blocking is almost entirely due to wavelength exhaust; intra-node contention is identically zero at all points.

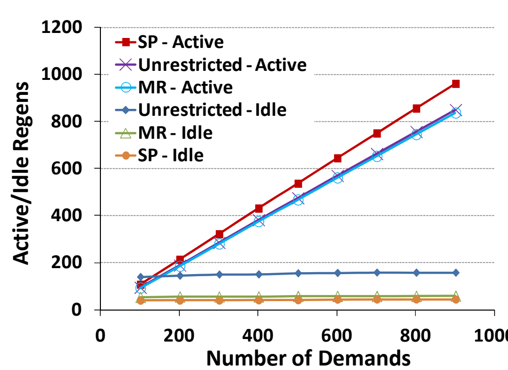

Fig. 4. Counts of idle and active regenerators vs. traffic demand.

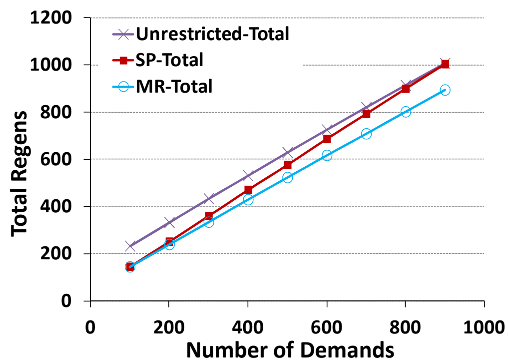

Fig. 5. Counts of total deployed regenerators vs. traffic demand in a service velocity network.

demands. During simulation, the cause of each blocked demand was tabulated, and it was found that virtually all blocking was due to wavelength exhaust (i.e., no contiguous wavelength was available for one of the regenerator-to-regenerator segments of the path). At very high demand levels, a few connections were blocked by transponder exhaust (this is an artifact of the simulation, which used fixed transponder pre-deployment). No blocking events were attributable to intra-node contention or to filled regenerator banks.

Fig. 4 shows the counts of active and idle regenerators as a function of the growing traffic demand for unrestricted regeneration sites as well as for the SP and MR strategies. Results above 900 demands are not plotted as these represent a distorted traffic pattern and an unrealistic mode of operation (above 900 demands, the network operator would be turning away customers whose requests are inconvenient). With unrestricted regeneration sites, the number of idle regenerators ranges from 139 (1.9 per network node) to 158 (2.1 /node). Both the SP and the MR subsets bring a substantial improvement: for SP, the number of idle regenerators ranges from 39 (0.52 /node) to 44 (0.59 /node), while for MR, the number of idle regenerators ranges from 53 (0.71 /node) to 58 (0.77 /node). Although the small SP site set leads to the lowest number of idle regenerators, the number of active regenerators used in SP is higher, as some routes replace a single midpoint regenerator that is not within the SP site set by two regenerators at subset locations. For the MR subset, the number of active regenerators is slightly lower than the unrestricted regeneration site case. As can be seen in Fig. 5, the MR subset achieves the lowest count of total deployed regenerators, offering a net reduction of 88 to 112 regenerators compared to the unrestricted case.

Although the cost of endpoint transponders and intermediate regenerators usually dominates the capital cost of a network, the cost of fiber, line amplifiers and other shared resources can also be significant. These shared costs can be approximately captured by an imputed cost per λ-km. For the MR case, the total λ-km will be increased over the unrestricted case, because some source-destination pairs are assigned paths that are not minimum distance. To quantify this effect, we calculated the excess path length ΔL for the 145 source/destination pairs for which this occurs, then multiplied each ΔL by the traffic probability for that pair, thus arriving at an estimate of the total increase in λ-km across the entire network. Under our population-based traffic model, using the MR strategy instead of unrestricted regeneration results in a 0.59% increase of the net λ-km in the CONUS network.

Other optimizations of the routes and site sets are possible. For example, a joint optimization of site set and routes that incorporates both regenerator cost and λ-km cost may be worth considering.

4. Conclusion

We have studied the pre-deployment of optical regenerators in a 75-node continental-scale network supporting service velocity, by performing Monte Carlo simulations of quasi-static traffic growth under a population-weighted traffic pattern. For a limited set of 22 regeneration sites, the number of idle regenerators in the network is reduced by >60% compared to unrestricted regeneration sites, while the number of active regenerators remains constant. A few routes become longer, leading to a 0.59% increase in the active λ-km across the network. With ~0.75 idle regenerators per network node, the results are very encouraging for service velocity networks, especially considering that the idle units can serve as spares in case of regenerator failure.

We gratefully acknowledge continuing encouragement and advice from Peter Magill and Robert Doverspike of AT&T Labs – Research and Keren Bergman of Columbia University.

5. References

[1] J. Strand and A. Chiu, "Realizing the advantages of optical reconfigurability and restoration with integrated optical cross-connects" J. Lightwave Technol. 21, pp. 2871-2882 (2003).
[2] G. Wellbrock, "Preparing for the future," in Proc. ECOC 2010, Torino, Italy, Symposium Th.9.G.1.
[3] S. L. Woodward, M. D. Feuer, J. L. Jackel, and A. Agarwal, "Massively-Scaleable Highly-Dynamic Optical Node Design", *Technical Digest of OFC/NFOEC 2010*, paper JThA18, 22-25 March 2010.
[4] S.L. Woodward et al., "Service Velocity Strategies in Optical ROADM Networks ", submitted to J. Optical Comm. and Networks (2011).
[5] The DARPA CORONET topology for the Contiguous United States (CONUS): http://monarchna.com/CORONET_CONUS_Topology.xls
[6] B.G. Bathula et al., "On Concentrating Regenerator Sites in Optical ROADM Networks", submitted to OFC/NFOEC2012.