Performance of IP over Optical Networks with Dynamic Bandwidth Allocation

Joel W. Gannett, George Clapp, Ronald A. Skoog, and Ann Von Lehmen

Telcordia Technologies, Inc., 331 Newman Springs Road, Red Bank, New Jersey 07701-5699
jgannett@research.telcordia.com

Abstract: IP over optical network performance can be improved with dynamic bandwidth allocation, depending on the reallocation paradigm and the network topology. Under high connectivity, dynamic bandwidth allocation provides a notable boost to the network's traffic-carrying capacity.

OCIS codes: (060.4250) Networks; (060.4510) Optical Communications

1. Introduction

Running IP over optical networks that are capable of assigning bandwidth resources dynamically can deliver improved performance [1]. Such networks shift bandwidth resources from one part of the network to another to follow either time-of-day traffic pattern variations or focused traffic loads caused by exceptional events. This paper addresses the performance of dynamic bandwidth allocation with respect to both the chosen reallocation paradigm and the network's connectivity. Our results indicate that the benefits of dynamic bandwidth allocation are modest except for highly connected networks. In the latter case, dynamic bandwidth allocation provides a notable boost to the network's traffic-carrying capacity.

2. Simulation Results

Our simulation model is based on a network of IP routers with OC-192 ports. The ports incorporate SONET VCAT/LCAS (Virtual Concatenation/Link Capacity Adjustment Scheme); hence, the 192 physical STS-1 channels of each port can be grouped into one or more virtual links that provide one-IP-hop connections to other routers. Although we are modeling the flexible bandwidth afforded by SONET VCAT/LCAS, the general principles we demonstrate apply equally well to a network interconnected by a flexibly-switched WDM optical layer. In the latter case, the switching granularity would be wavelengths rather than STS-1 channels.

To enable the simulation of large-scale traffic flow effects in complex networks, we model the usage of the IP network as a stochastic sequence of abstract *calls* between routers, each with a finite stochastic holding time. During the time it is active, each abstract call represents an increase of one STS-1 in the traffic flow between the two routers at its end points. While it is active, such a call in our model consumes one STS-1 of bandwidth on each link over which it is routed. Our call generation model uses exponential arrivals and exponential holding times. As each abstract call arrives, its two end points are assigned randomly such that the expected time-average number of requests terminating at a router is proportional to its number of ports (i.e., if router *X* has twice as many ports as router *Y*, then *X* terminates twice as many abstract calls as *Y* on average over time). Although our traffic termination model is unconventional, we find it more intuitive and more plausible than the conventional uniform traffic assumption. Note also that no concept of distance enters into our selection of call end points, so routers that are widely separated in the network may have as much traffic between them as routers that are close together. This conforms to modern data networks, such as the Internet, where the interaction paradigm seems to be "everyone talks to everyone." Consequently, we feel that our distance-insensitive paradigm for assigning end points is more appropriate than a conventional "gravity" traffic model.

We assume a centralized or distributed network management and control (M&C) system is used that issues bandwidth reallocation requests to the VCAT/LCAS layer. In a real network, the M&C system would issue these requests in response to changes in monitored performance metrics that indicate inadequate bandwidth (congestion) in some part of the network. These metrics may include the packet delay or packet drop ratio. In our simulations, we use the *call blocking ratio* (the ratio of blocked calls to total calls) as a surrogate for assessing network performance. When an abstract call is blocked, that means there is less bandwidth on a link than there ought to be to accommodate the offered traffic. For best performance, the blocking ratio should be kept as low as possible.

We simulate two static bandwidth allocation schemes and two dynamic bandwidth allocation schemes, denoted as follows: Static Allocation, Static Routing (SA-SR); Dynamic Allocation, Wait-For-Trouble, Static Routing (DA-WFT-SR); Dynamic Allocation, Busybody, Static Routing (DA-BB-SR); and Static Allocation, Dynamic Routing (SA-DR). SA-SR is similar to today's Internet routing, that is, it is based on static routes between routers with static bandwidth on the links. The static routes constructed in these simulations are minimum-hop paths between the routers. DA-WFT-SR uses static routes and performs a bandwidth reallocation when (and only when) a call is blocked. The reallocation algorithm incorporates a heuristic that tries to shift any available (unused) capacity over to 100% utilized (blocked) links, thereby unblocking them. DA-BB-SR also uses static routing, but it checks after each successfully routed call for any link that may have become 100% utilized as a result of that new call; hence, DA-BB-SR reallocates more often than DA-WFT-SR. Once reallocation is indicated, DA-BB-SR uses the same heuristic as DA-WFT-SR to unblock links by reallocating bandwidth. The reallocation paradigm must keep at least one STS-1 allocated on every link at all times to avoid triggering IP reconfiguration. Finally, SA-DR leaves the allocation static but uses Dijkstra's algorithm to find the shortest path that routes around blocked links. SA-DR is the only one of the four schemes that uses dynamic routing.

Our example network consists of 10 routers, numbered 0 to 9, with between 3 and 5 backbone ports for each router and a total of 40 ports. The number of virtual VCAT/LCAS links between the router backbone ports is varied from 20 to 45 in increments of 5. The end points of each link must terminate on ports of distinct routers, and we allow no more than one link connecting two routers. Therefore 45 links is the maximum number of links allowed in our example, as this results in one link between each pair of routers. The links are placed randomly, except the placement algorithm guarantees (if there are at least 20 links) that each of the 40 ports is connected to at least one link. Routers with more ports are likely to get proportionately more links. The 25-link network has all the links of the 20-link network, plus five more, while the 30-link network has all the links of the 25-link network, plus five more, etc. See Fig. 1 for a graph of the 30-link network.

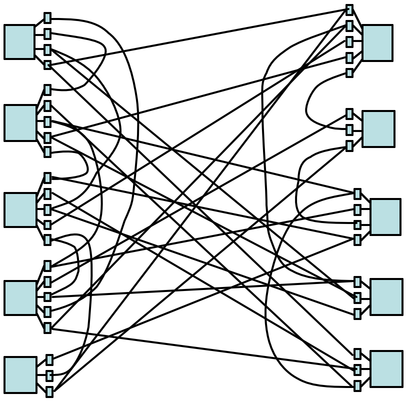

Fig. 1. Network with 10 routers, 40 backbone ports, and 30 virtual links.

Each simulation experiment consists of 3,000,000 abstract call attempts, with the blocking ratio tallied at the end. For the first 20% of the calls in an experiment (first 600,000 calls), the end points are generated randomly as described earlier. For the next 20%, the focus is on router 8, which gets its average terminating traffic doubled from about 10% of the total to about 20%. The focus shifts to router 1 in the next 20% of the calls, then to router 5 in the next 20%. For the last 20% of the calls, the random end point generation paradigm of the first 20% is used again. The dynamic allocation paradigms must track these changes in traffic focus to reduce the blocking ratio. Figure 2 shows the simulation results for the 30-link network. The graph shows the increase in call blocking as the network-wide load in Erlangs (average call holding time) is increased gradually.

Figure 3 summarizes the results for the six network topologies simulated. We define the capacity for each network topology and each allocation paradigm as the offered load in Erlangs that causes a blocking ratio of 0.001.

3. Concluding Remarks

With the Static Allocation, Static Routing (SA-SR) paradigm, which is similar to ordinary Internet routing, varying the number of virtual links between routers implies two phenomena that seem to cancel. On the one hand, more links means fewer multi-hop routes, which implies more efficient bandwidth usage and *less* blocking. On the other hand, more links means less static bandwidth per link, which means *more* blocking. The net result, shown in Fig. 3, is that the network capacity oscillates around 800 plus or minus 200 Erlangs as the number of links is varied.

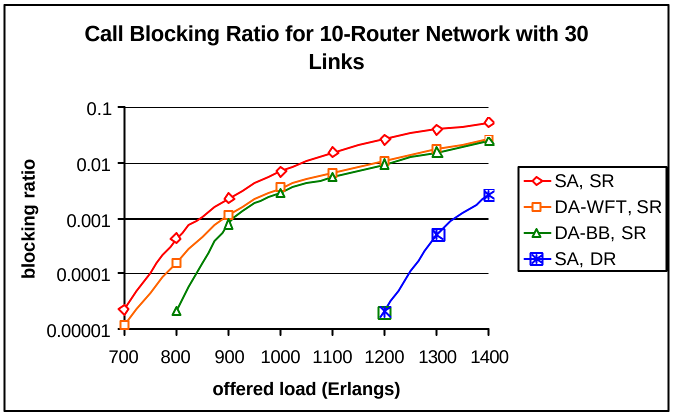

Fig. 2. Call Blocking Ratio for the 30-link case.

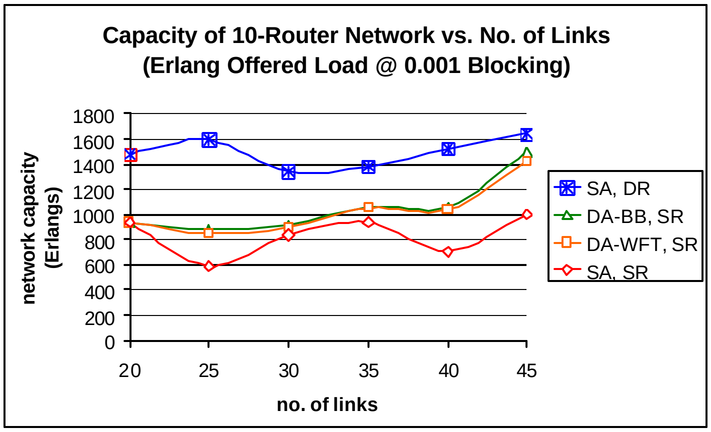

Fig. 3. Network capacity vs. number of links.

Similarly, the Static Allocation, Dynamic Routing (SA-DR) paradigm gives the highest network capacity but also seems fairly invariant to the number of virtual links between the routers. The resulting capacity is almost twice as high as SA-SR, being around 1500 plus or minus 150 Erlangs.

The two dynamic allocation paradigms, DA-WFT-SR and DA-BB-SR, perform similarly, with the more intensive DA-BB-SR having a slight edge. These two paradigms show performance intermediate between the two static allocation paradigms, with modestly rising performance as the number of virtual links is increased. Interestingly, the performance of the dynamic allocation paradigms improves sharply, rivaling the performance of SA-DR, when full virtual connectivity is provided and each router is just one IP hop away from every other router. Such a topology is viable for small networks, but may be impractical when the number of routers is large.

Although dynamic bandwidth allocation may not quite match the performance of dynamic routing, our results show that it can come close under certain conditions. Moreover, dynamic allocation networks may be easier to manage than dynamic routing networks, as the latter are subject to thrashing, congestion spreading, and other instabilities. Further work should explore what types of load variations, topologies, and other network properties make dynamic bandwidth allocation a superior strategy to dynamic routing in terms of both management and performance.

References

[1] R. Skoog and S. Yun, "The value proposition for bandwidth on demand," Optical Fiber Communication Conference 2003 *Technical Digest*, Atlanta, Georgia, March 23-28, paper ThH3, pp. 482-483.