Dynamic Cooperative Spectrum Sharing in Elastic Networks

E. Palkopoulou[1], I. Stiakogiannakis[1], D. Klonidis[1], K. Christodoulopoulos[2], E. Varvarigos[3], O. Gerstel[4], and I. Tomkos[1]

[1] *Athens Information Technology Center, Peania, Greece, Tel: (+30)210 6682778, Email: elenip@ait.gr*
[2] *School of Computer Science and Statistics, Trinity College Dublin, Ireland*
[3] *Computer Engineering and Informatics Department, University of Patras, Greece*
[4] *Cisco, Israel*

Abstract: Spectrum expansion, contraction, and re-allocation policies are proposed for flexible networks. Reductions in blocking probability can be achieved by introducing "neighbor avoidance" mechanisms. Trade-offs between blocking probability and the number of re-allocated connections are quantified.

OCIS codes: (060.0060) Fiber optics and optical communications; (060.4256) Networks, network optimization

1. Introduction

The elastic optical networking paradigm sets the stage for envisioning highly adaptive and spectrally efficient networks [1]. It enables traditionally rigid optical connections to adapt to the network conditions and to dynamic traffic demand patterns. However, as these elastic connections fluctuate over time, the issue of bandwidth fragmentation arises, which limits the potential spectral efficiency gains. To resolve this issue, various spectrum defragmentation approaches - requiring the periodic or on-demand re-allocation of spectrum - have been proposed [2,3]. While unused spectrum slots between connections may seem to be a problem (as fragmentation metrics are derived based on spectrum usage patterns), it is not always the case. In fact, as long as there is a need for superchannels to grow over time – contiguous and continuous free spectrum slots are necessary in order for the additional traffic demand to be accommodated without requiring the re-allocation of established connections. A different approach is to deliberately leave space between connections. In [4] spectrum expansion/contraction (SEC) policies are proposed that enable the sharing of free spectrum slots among connections. Case study results indicate that significant savings can be achieved in terms of blocking probability.

In this work, a novel SEC policy, which takes into account the spectrum allocation of the neighboring connections, is proposed and compared with policies from [4]. In addition, we propose different joint SEC and spectrum re-allocation policies. As neighboring connections compete for the same spectral resources, we introduce for the first time cooperative spectrum sharing mechanisms in flexible optical networks. We examine occurring trade-offs with respect to the reduction in blocking probability and the number of connections that have to be reallocated. We consider hitless re-allocation techniques, based on channel re-tuning, that do not require the deployment of additional network resources (such as costly additional transponders). The re-tuning operation is based on capabilities available in the coherent receiver [3]. Our contribution in this respect is twofold: (i) we reduce the probability that a connection will be blocked by neighboring connections (proactive approach); (ii) we examine techniques to resolve this blocking when it arises (reactive approach).

2. Spectrum Expansion, Contraction, and Re-Allocation Policies

In Fig. 1 an example network is presented which carries 3 demands. We proceed to examine the limitations that are imposed on connection *B* by its "lower" and "upper" neighboring connections *A* and *C*, respectively. In all of the SEC and spectrum re-allocation policies each connection is assigned a path and a reference frequency F_p. Additionally, a maximum allowable expansion region is defined.

We consider that the total number of requested spectrum slots varies over time. When a request for a spectrum slot arrives, the relevant spectrum expansion procedure is invoked. If there are no available slots in the maximum allowable expansion region, then the spectrum re-allocation procedure is invoked. If spectrum re-allocation is not allowed or if the re-allocation procedure fails, then the demand is blocked. Re-allocation procedures are based on previously discussed hitless re-allocation techniques.

In the following we describe the different SEC and spectrum re-allocation policies.

- **Constant Spectrum Allocation (CSA):** In CSA a fixed number of spectrum slots are assigned to each connection. No spectrum sharing is permitted among neighboring connections. Spectrum re-allocation is not allowed. This policy is defined in [4].
- **Dynamic Alternate Direction (DAD):** In DAD spectrum sharing is allowed within the region defined by the reference frequencies of the neighboring connections. Hence, the maximum allowable expansion region of connection *B* is defined by F_A and F_C. DAD aims at the symmetrical use of spectrum around the reference

978-1-55752-962-6/13/$31.00

frequency. If this is not possible (due to the occupation of the desired slots by a neighboring connection), then asymmetrical expansion is allowed. Spectrum re-allocation is not allowed. This policy is defined in [4].

- **Avoid Close Neighbors (ACN):** In ACN spectrum sharing is allowed within the region defined by the reference frequencies of the neighboring connections – as in DAD. As its name suggests, ACN is based on the principle of cooperative "neighbor-avoidance". Cooperation in this respect is derived from the notion that each connection performs SEC in a manner that avoids consuming spectral resources from connections that have potentially less available resources. More specifically, each connection expands towards the opposite direction of its closest neighbor on any of the links along its path. Note that there are at most two direct neighbors in each of the links of the established path of the connection. Following the same principle, contraction is conducted from the direction of the closest neighboring connection. Spectrum re-allocation is not allowed.
- **Shift-ACN:** Shift-ACN is an enhanced version of ACN that allows spectrum re-allocation. If there are no available slots in the maximum allowable expansion region, then Shift-ACN tries to shift its neighbors. It first tries to shift the spectrum occupied by its neighbors towards the direction that maximizes the minimum available free slots among all of its neighbors. If this procedure fails, then it tries to shift the spectrum occupied by its neighbors from the other side. In Shift-ACN we impose a restriction on how connections can be shifted. The assigned reference frequencies F_p act as anchors for the connections – meaning that connections cannot be shifted beyond their reference frequency. As a result a maximum allowed expansion region is defined. Note that we only allow shifting of direct neighboring connections of the connection requesting a spectrum expansion.
- **Float-ACN:** Float-ACN is an enhanced version of Shift-ACN that does not impose a restriction on the shifting of connections. In other words, connections are free to "float" in the spectrum as they are "pushed" by their neighbors. Note that the reference frequencies are shifted only as a last resort. The first step of Float-ACN is the invocation of the procedures described in Shift-ACN. If spectrum allocation fails, it means that the connection is limited by at least two reference frequencies. In this case Float-ACN first tries to shift the reference frequencies towards the direction that maximizes the minimum available free slots among all of its neighbors. If this procedure fails, then it tries from the other direction.

In Fig.1 we additionally show how spectrum allocation is conducted for the consecutive arrival of three requests for connection *B*. Each request corresponds to one spectrum slot. In CSA the request at t_1 is blocked, because there are no free slots in its allowed expansion region. The request at t_2 is accommodated only by Shift-ACN and Float-ACN, because spectrum re-allocation is performed for connection *C*. The request at t_3 is accommodated only by Float-ACN, because it is allowed to move the reference frequency of connection *C*.

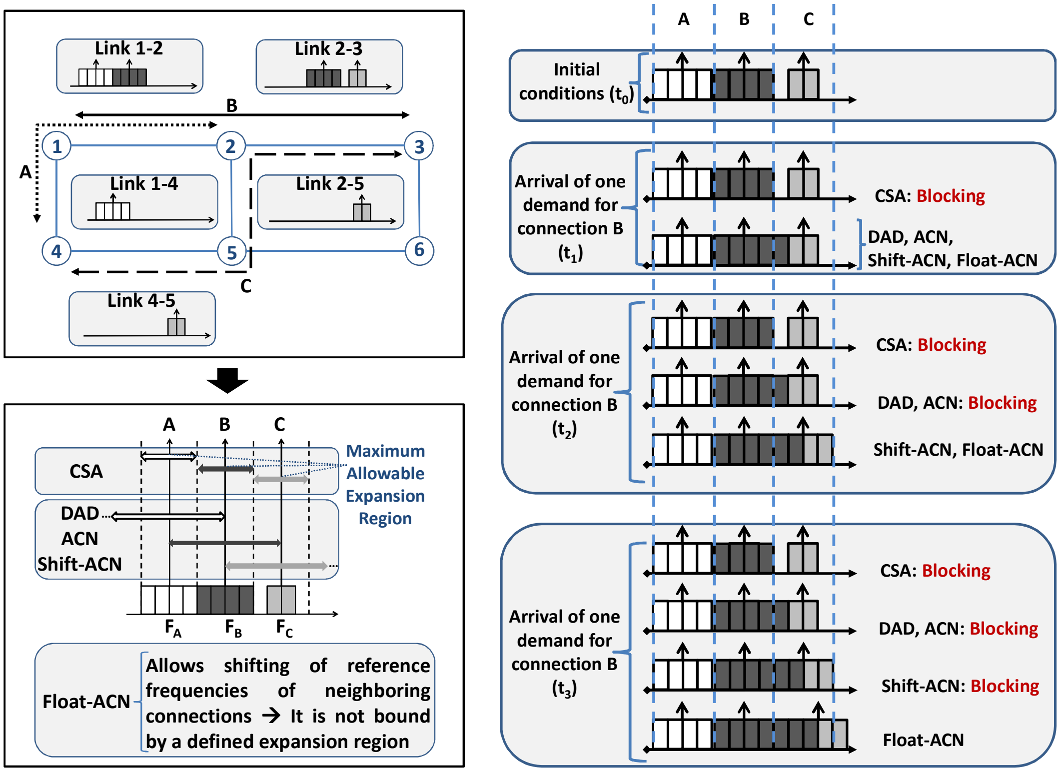

Fig. 1. Spectrum expansion, contraction, and re-allocation policies. Guard-bands among connections are not depicted for simplicity reasons.

3. Case Studies

In this section we evaluate the performance of the proposed SEC and spectrum re-allocation schemes. Case studies are conducted using the Deutsche Telekom reference network (14-nodes, 23 links) and assuming that 250 spectrum slots are available per fiber. Spectrum slot requests are generated according to a Poisson process having a rate of λ_p and an exponentially distributed duration with mean $1/\mu_p=1$. The arrival rate λ_p of the slot requests for each connection p is given by an exponential distribution with mean λ. The initial routing and spectrum allocation (RSA) is performed based on the Erlang-B formula, as described in [4]. For each traffic load, simulations for 10^7 slot requests are conducted.

In Fig. 2(a) the average blocking probability is presented as a function of the total load. The CSA and DAD policies from [4] are used as a benchmark. As expected, DAD offers a significant improvement in performance compared to CSA by allowing neighboring connections to share spectrum slots. We show that additional benefits can be reaped via the introduction of “neighbor-avoidance” mechanisms. By comparing DAD with ACN, we find that an average relative reduction in blocking probability of 44% is achieved (note that the y-axis is in log-scale). Shift-ACN provides an average relative reduction in blocking probability of 22% compared to ACN, by allowing connections to be shifted. The best performing policy is Float-ACN, which achieves a 96% average reduction of blocking probability compared to CSA. Thus, a significant reduction in blocking probability can be achieved by allowing the reference frequencies of connections to freely “float”.

We now quantify the number of re-allocations that are performed in order to achieve the relevant reductions in blocking probability for policies Shift-ACN and Float-ACN. Since we only allow shifting of direct neighboring connections of the connection requesting a spectrum expansion, all re-allocations can be performed in parallel. We define C as the number of connections that have to be shifted every time the re-allocation procedure is invoked. In Fig. 2(b) we present the maximum and the average value of C as a function of the total load. As expected, for higher traffic loads, the maximum number of re-allocated connections C increases. Such cases, however, occur rarely - as the average number of re-allocated connections remains fairly constant for increasing traffic load. In Fig. 2(c) the percentage of arrivals that successfully invoke the re-allocation procedure is presented as a function of the total load. Note that the computational overhead introduced by the proposed schemes is expected to be insignificant to the overall connection set-up time as it lies in the µs-range in our implementation.

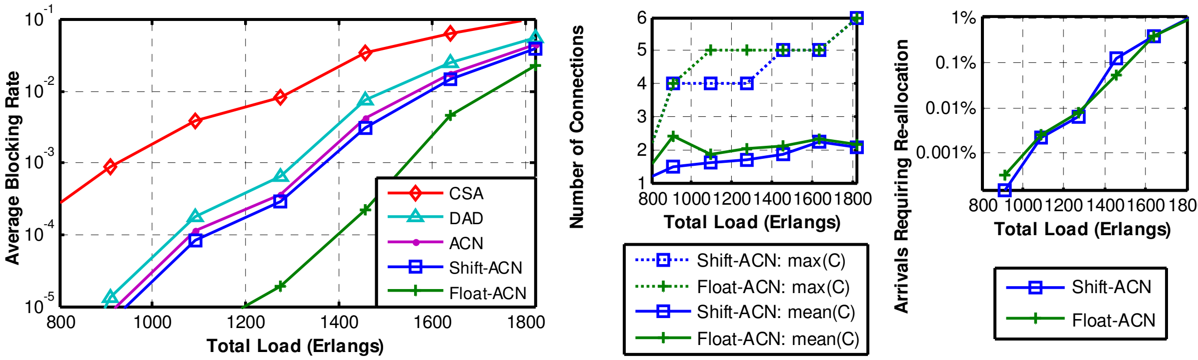

Fig. 2 (a) The average blocking probability as a function of the total load; (b) The maximum and the average number of connections that have to be shifted every time the re-allocation procedure is invoked as a function of the total load; (c) The percentage of arrivals that successfully invoke the re-allocation procedure as a function of the total load. The total load is equal to $\lambda|V|(|V|-1)$, where $|V|=14$ is the number of network nodes.

4. Conclusion

We jointly consider and propose different spectrum expansion, contraction, and re-allocation policies. We find that reductions in blocking probability can be achieved by introducing cooperative “neighbor avoidance” mechanisms. We, additionally, examine the benefits that can be reaped by allowing connections to be dynamically re-allocated based on capabilities available in the coherent receiver. It is shown that in order to gain the maximum benefits in performance, connections should be allowed to freely shift in the spectrum – without imposing restrictions on their reference frequency.

This work is partially supported by the EC FP7 CHRON project n° 258644 and by the CISCO University Research Program.

References

[1] O. Gerstel et al., “Elastic Optical Networking: A New Dawn for the Optical Layer,” IEEE Communications Magazine, February 2012.
[2] Xi Wang et al., “A Hitless Defragmentation Method for Self-optimizing Flexible Grid Optical Networks,” ECOC 2012.
[3] F. Cugini et al, "Push-Pull Technique for Defragmentation in Flexible Optical Networks," OFC/NFOEC 2012.
[4] K. Christodoulopoulos et al., “Time-Varying Spectrum Allocation Policies and Blocking Analysis in Flexible Optical Networks,” IEEE JSAC - Special Issue on Next-Generation Spectrum-Efficient and Elastic Optical Transport Networks.