

Using optimal ILP techniques and approximating heuristics to examine the influence of relocation on network dimension in optical grids

Jens Buysse, Marc De Leenheer, Chris Develder, Bart Dhoedt

January 7, 2010

Abstract

We present a study on resiliency against single link failures in optical grids and demonstrate how the Grid-specific anycast principle can be exploited to reduce network capacity requirements when adopting shared path protection. Since grid users usually do not care about the location their jobs end up being executed, we take advantage of relocation to alternate backup sites in case of failures. We solve the resulting routing-and-wavelength-assignment (RWA) problem optimally using Integer Linear Programming (ILP), extending earlier preliminary work. We also present a new, scalable heuristic to demonstrate that in a realistic case study (29 node European network, up to 150 connections) exploiting relocation results in about 20% less wavelength capacity compared to traditional shared path protection.

An important aspect of network deployment is the ability to survive from certain network failures. To deal with those, a network operator can opt to provide so-called network protection. We compare two protection schemes where, to protect a connection path from one point to another, a back-up path is reserved (for use when the primary fails). In case of so-called shared path protection, each primary path has a backup path where wavelengths can be shared between several back-up paths, as long as their corresponding primary light paths do not overlap. We have extended this shared path protection algorithm in a Grid context, where the anycast principle prevails: in general, multiple processing locations exist in a Grid network and any of those is an equal candidate for processing users' jobs (hence there is no a priori fixed destination of a particular job). Thus, instead of reserving a back-up path to the original destination determined by the Grid scheduler, we may relocate the job to another, possibly closer resource and thereby reduce overall network capacity.

We evaluate the classical shared path protection (CSP) method and the shared path protection with relocation (SPR) scheme from a network dimensioning perspective by an Integer Linear Program (ILP). We start from a demand vector stating for each source node in the network a number of connections to be established between it and some resource. Hence we are free to determine not only the resource to serve as a backup, but also the best primary resource. We model the protection schemes as a minimum multi-commodity flow problem with according demand- and flow conservation constraints. We extend the CSP formulation to a SPR formulation by allowing the primary- and backup resources to differ. As demonstrated in [1] we show that these ILPs can only render us solutions for rather small network topologies where the demand has to stay relatively small. In these limited-scale cases, we found that SPR attains a network load reduction (NLR) (the number of necessary wavelengths to establish the desired protected connections) of about 20% compared to CSP.

The ILP formulation discussed above allows finding optimal solutions, but it is well-known not to be scalable for large problem instances (high complexity in terms of memory utilization and execution time). Hence, in order to evaluate the relocation strategy on a larger scale, we also propose a heuristic solution. These heuristics work in 4 stages.

1. For every connection we calculate two initial edge disjoint paths from the source node to one of the available resources in the network. This is achieved by an adaptation of Suurballe's algorithm which will find a pair of edge-disjoint paths from vertex s to vertex d such that the total cost of the two paths is minimal among all such paths. We take the resource which minimizes its sum of links of the link-disjoint paths.
2. We choose which one of the pair of those link-disjoint paths to use as a primary path and which one as a backup path.
3. Then, for every connection try to maximize the sharing among the backup paths by rerouting its backup path.
4. And the final step includes minimizing the resources needed for the primary paths by rerouting its primary path.

The heuristic proves to be a good approximation of the optimum solution (ILP). On average, the total number of wavelengths for the shared path protection heuristic only differs 4.5% from the number of wavelengths for the ILP and this is even less for the relocation heuristic (3,7%).

In step 3 and 4 of the heuristic we constantly pick one connection and we try to reroute one of its paths. The order in which these connections are chosen could be an important choice because rerouting one path could mean we cannot reroute another path in the future. Therefore we compared different strategies, which turned out to all strand at about the same solution (in terms of wavelength capacity).

In earlier limited case studies with ILPs [1] [2] we found slightly different savings achieved by relocation. The main differences between these papers were the considered topologies, which lead us to believe that these savings are topology dependent. As a result we performed the study on three different topologies namely our European reference network, a much sparser ring network and an almost completely meshed network. Our results indicate that savings are strongly topology dependent (but nearly independent on demand vector size): the NLRs for ring, reference and mesh network were respectively 69%, 83% and 87%.

Summarizing our conclusions, in this work we have described an alternative protection method applicable in optical grids. While in classical shared path protection, a primary path from a source to a destination is protected by a back-up path to the same destination, our relocation scheme allows a back-up path to another resource. This is acceptable by the anycast principle (typical of grids), implying the user does not care about the location of job execution. We formulated an exact ILP solution which did not prove to be scalable and hence devised a scalable heuristic. Results in a realistic European network scenario show that exploiting relocation achieves considerable bandwidth savings compared to traditional shared protection (depending on the topology, ranging from 13% to 31%).

References

- [1] Jens Buysse, Marc De Leenheer, Chris Develder, and Bart Dhoedt. Exploiting relocation to reduce network dimensions of resilient optical grids. In *Design of Reliable Communication Networks (DRCN)*, 2009.
- [2] Jens Buysse, Marc De Leenheer, Chris Develder, and Bart Dhoedt. On the impact of relocation on network dimensions in resilient optical grids. In *14th Conference on Optical Network Design and Modeling*, 2010.