Efficiency Analysis of Distributed Dynamic Optical Bypassing Heuristics

Joachim Scharf
Institute of Communication Networks and Computer Engineering
University of Stuttgart, Pfaffenwaldring 47, Germany
Email: joachim.scharf@ikr.uni-stuttgart.de

***Abstract*—The energy consumption of transport networks gets more and more attention as it is expected to become a problem with future traffic increase. Dynamic optical bypassing is a promising approach to tackle increasing energy consumption as well as growing resource usage. Performance evaluations are necessary in order to find suitable bypassing approaches. Recently, distributed approaches for bypass establishment and teardown have been proposed. There is so far no comprehensive performance evaluation available. In this paper, we present four different distributed heuristics for bypassing. We implemented a simulation framework for the performance evaluation of these heuristics. Two approaches without dynamic bypassing serve as reference. The results show benefits of dynamic bypassing with respect to required energy and resources. We relate these benefits to the necessary switching operations.**

I. INTRODUCTION

Network operators pay more and more attention to the energy consumption of their transport networks. Currently the energy consumed in the core is low compared to that in the access. However, this will likely change due to the combination of traffic growth rates of around 50% per year and an overproportional increase of energy consumption in the core [1]. As the overall energy consumption significantly contributes to the operational expenditures (OPEX), mechanisms which help to reduce the energy consumption are necessary.

Besides the gradual increase in traffic there is also a diurnal traffic profile [2]. Maximum traffic peak rates occur in the evening when most people are at home and surf the Internet. In the early morning most people are sleeping. Here, traffic peak rates are only about 25% of those in the evening. In order to provide a good quality of service (QoS) at all times, network operators have to dimension their networks according to maximum peak rates. Thus, there are a lot of underutilized resources in times without the maximum traffic load. Powering down at least some of these resources helps to save energy without influencing the QoS observed by the users. In order to achieve this, mechanisms are necessary that decide on powering resources on and off.

A further observation is the emergence of large data centers. These data centers are the origin and destination of an enormous amount of traffic. If a service moves from one data center to another (e.g. due to restructuring of responsibility, maintenance or outage), this may severely impact the traffic inside transport networks. Some links may become heavily overloaded whilst others carry only a fraction of the previous traffic. In principle, it is possible to prevent overload purely by overdimensioning all network resources. However, this is far from being cost-efficient – not to mention the excessive energy consumption. Thus, mechanisms are needed that support network operations in the presence of highly dynamic traffic. Due to the high dynamics, these mechanisms should be embedded into the control plane of the network.

Dynamic optical bypassing can solve the above addressed issues in a multilayer network with an upper IP/MPLS layer and a lower wavelength switched optical network (WSON) layer. In section II, we explain the concept of dynamic optical bypassing and present related work. Section III addresses challenges of dynamic optical bypassing and heuristics. We introduce our simulation framework as well as the simulation scenario in section IV. In section V, we analyze the performance of the heuristics. Section VI concludes our work.

II. DYNAMIC OPTICAL BYPASSING

A. Concept

Networks which are not excessively overdimensioned need to adapt to traffic demands. In case of static traffic demands (i.e. demands do not vary over time), this adaptation is done only once during the network dimensioning process itself. In case of dynamic traffic demands (i.e. demands are a function of time), such an adaptation has to be carried out during operation and is better known under the term *network reconfiguration*.

In IP/MPLS over WSON multilayer networks, each layer offers possibilities for a reconfiguration to dynamic traffic demands. Firstly, changes in the routing of the IP/MPLS layer lead to a modified usage of existing optical circuits. Secondly, network operators can establish and tear down optical circuits according to the needs of the traffic in the IP/MPLS layer.

Dynamic optical bypassing adapts the circuits in the WSON layer with the objective to offload IP/MPLS processing from nodes with a high amount of transit traffic in the IP/MPLS layer. New optical circuits called *bypasses* can achieve this objective. Possible energy savings due to less energy consumed for packet processing in the IP/MPLS layer reason this approach. Lowly utilized bypasses may require more energy than they save. Therefore, it is necessary to tear down bypasses when indicated.

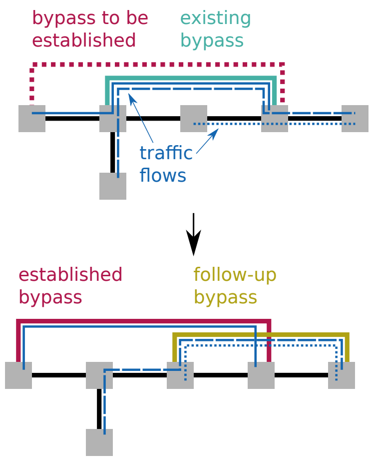

Fig. 1: Mutual influence of bypasses

B. Related Work

The general concept of dynamic optical bypassing, namely reconfiguration of the logical topology in IP/MPLS over WSON networks, is well known and has been extensively studied in literature [3]. Although hybrid networks (e.g. *optical migration capable networks with service guarantees*, OpMiGua [4]) rely on a different network architecture than IP/MPLS over WSON networks, they imply problems similar to those of dynamic bypassing.

The vast majority of work on reconfiguration of logical network topologies concentrates on scenarios in which all traffic demands as well as the network state are available. This information serves for calculation of a better network configuration. Applied methods are (mixed) integer linear problems [5], [6], heuristics [7] and metaheuristics [8]. While the energy consumption is mostly modeled rather abstract, there are few exceptions with a detailed modelling [9].

The central finding of a comparison of reconfiguration in the IP/MPLS layer, in the WSON layer and simultaneous reconfiguration in both layers for current traffic is that reconfiguration in the IP/MPLS layer obtains the largest fraction of achievable energy savings [10]. However, future traffic developments as described above may invalidate this statement.

Few publications consider distributed approaches with less available information [11]. In this paper, we assume such a scenario. The work of Ruffini [12] and the joint work of Milbrandt and Broniecki (contribution to STRONGEST project [13]) are the basis for the used heuristics.

III. Heuristics for Dynamic Optical Bypassing

A. Classification of Approaches

The objective of dynamic optical bypassing is to offload packet switching in the IP/MPLS layer from network nodes. For this objective there are multiple approaches for the bypass decision process. These approaches differ with respect to the required information for the decision as well as in the decision-making (centralized versus distributed).

Omniscient approach: A central instance has complete knowledge of the whole network (i.e. current resource usage and traffic demands). The decision entity has the maximum amount of available information and can therefore determine the most efficient bypasses. However, this solution has two major drawbacks. First, the transport of traffic demand information requires excessive signalling. Second, the central instance is a single point of failure.

Distributed approach: Nodes decide about bypasses by themself without coordination by a central instance. Advantages of this approach are a minimal amount of required signalling (only for bypass establishment and teardown) and the absence of a single point of failure. Less efficient bypasses are the drawback.

Further approaches exist in between these two extremes. Nodes requesting bypass establishments and teardowns combined with a central coordinating instance is one example.

This paper focuses on distributed approaches which have only local node information or at most information originating from neighbor nodes. Information from more distant nodes is not available.

B. Challenges

Mainly due to cost pressure, there is a general trend to more automation in networks (e.g. self-organizing networks in the wireless area). If dynamic reconfigurations enable significant benefits, it is a safe assumption that automation will eventually be used for reconfiguration inside core networks. Therefore, a dynamic bypassing heuristic has to prove two things. First, the heuristic actually saves costs. Second, dynamic bypassing can be integrated into the control plane.

Nowadays, most network reconfigurations (besides protection switching in case of failures) require some human interaction in the control loop. This procedure should guarantee network stability, which is also mandatory in case of automated reconfiguration. Therefore, a dynamic bypassing approach must not have an impact on the stability of network operation and the transported traffic.

While a make-before-break strategy for any reconfiguration limits the impact on current traffic, distributed heuristics for dynamic bypasses may critically influence network stability. The bypasses require switching operations and the mutual influence between bypasses further aggravates this. Fig. 1 shows an example for such a mutual influence. The establishment of a new bypass (dotted red line) takes away significant parts of traffic from an existing bypass (green solid line). In the example, we assume that the load gets too low, which leads to a tear down of the green bypass. The teardown itself impacts the traffic in further nodes, which again may trigger bypasses (yellow bypass). Single actions may thus propagate through the network and trigger follow-up actions. Obviously, this is not desired and therefore dynamic bypass heuristics need to keep switching operations at a tolerable limit.

C. One-hop Bypasses

One-hop bypasses are based on the idea of [12]. A set of permanent optical circuits interconnects all nodes and

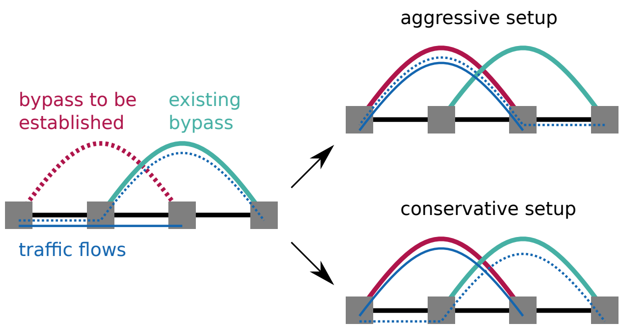

Fig. 2: Alternative approaches for one-hop bypass setup

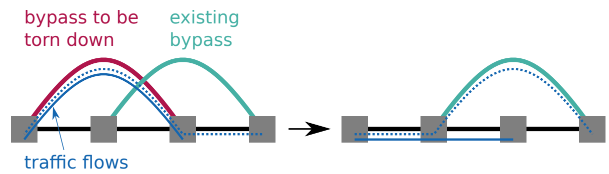

Fig. 3: Tear down of one-hop bypass

ensures reachability. We denote the resulting topology as basic topology. Each node tracks for any node tuple (n_1, n_2) of its neighbor nodes in the basic topology the current transit traffic in the IP/MPLS layer from n_1 to n_2. The bypassing heuristic monitors the resulting demand matrix, which has y very limited scope. As soon as an entry exceeds a threshold t_1, it triggers a bypass between the according nodes. The resulting bypass circuit interconnects nodes that have a distance of two hops in the physical topology (although it has two hops in the physical topology we call it one-hop bypass since it reduces the distance by only one hop). There is no announcement of the bypass to any nodes besides the end points. Thus, the bypass has only influence on the routing decisions in the bypass start node.

We consider two approaches for traffic handling in the start node of the new bypass (cf. Fig. 2). *One-hop/Conservative Setup* (OH/CS) is the first approach. OH/CS puts only the fraction of traffic into the bypass that makes up the actual transit traffic in the node triggering the bypass (solid blue traffic flow). Consequently, the bypass does not affect the depicted dotted traffic flow. While OH/CS is in line with [12], this is not the case for our new second approach *One-hop/Aggressive Setup* (OH/AS). OH/AS additionally puts the traffic into the bypass that would be forwarded to the bypass end node if there were no other bypasses. In this case, the dotted traffic flow in the example uses the new bypass instead of the existing one. While OH/AS is a good example for the mutual influence of bypasses described above, there is no such influence in case of OH/CS.

At some point in time the traffic originally causing the bypass may vanish. In order to prevent underutilized bypasses, both heuristics tear down bypasses whose traffic falls below a threshold t_2. Fig. 3 shows this process. The traffic from the former bypass start point is routed again in the basic topology. The next node decides on its own whether to put the traffic into an existing bypass or routing it along the basic topology.

We assume fixed values for the thresholds t_1 and t_2. In principle, it is possible to change the thresholds t_1 and t_2 dynamically. However, this is subject to future work.

D. Multi-hop Bypasses

One-hop bypasses span exactly two hops in the physical topology. Abandoning this restriction leads to multi-hop bypasses. As most networks have node pairs with a distance greater than two hops, such multi-hop bypasses enable the connection of further node pairs.

The successive extension of bypasses at the start or end node by one physical hop is one possibility [12]. Our approach has slightly less constraints as we admit bypasses between any node pair with a distance of two hops in the virtual topology. Consequently, successive bypass establishments may lead to bypasses spanning multiple hops in the physical topology. This heuristic establishes bypasses analogously to OH/AS. We denote it as *Multi-hop/Successive One-hop* (MH/SOH).

A very recent proposal for multi-hop dynamic optical bypassing is STRIPOFF (self-triggered dynamic sink-tree bypass provisioning for offloading packet switches). This approach by Milbrandt and Broniecki was presented within the STRONGEST project [13], but is not yet published. Therefore, we explain the basic operation mode as far as necessary for understanding this paper. Traffic is forwarded hop-by-hop within a spanning sink tree as basic topology (other basic topologies are also feasible). Each node keeps track of the accumulated traffic to each destination node. If in node n_1 the traffic to node n_2 exceeds a certain threshold t_1, node n_1 requests a bypass directly to node n_2. Either a central unit or all nodes involved in the path to n_2 decide on the bypass establishment and may decline it. In case of bypass establishment, node n_1 puts solely traffic destined to node n_2 into this bypass. Bypasses with too little traffic (threshold t_2) trigger a teardown. The traffic from these bypasses is routed again within the basic toplogy.

So far, there are no proposals for the bypass decision process of STRIPOFF available in literature. Therefore, we implemented the simplest option in which a bypass request is always granted. We denote the according heuristic as *Multi-hop/Sinktree* (MH/ST).

IV. Simulation Setup

A. Scenario

For the evaluation of the heuristics, we use the Atlanta reference network [14] as physical topology. This network has 15 nodes and 22 links, resulting in a nodal degree of 2.93. For simplicity, we assume uniform traffic that follows a diurnal traffic pattern between all nodes. The pattern is approximated by a shifted sinusoidal curve. The low of this curve is one fourth of its peak value. We generate the traffic by a superposition of traffic flows with variable negative-exponentially distributed interarrival time following the diurnal pattern and negative-exponential holding time.

We assume unlimited resources and full wavelength conversion capability. Thus, there is no blocking. We monitor the actually used resources, which indicate the necessary network dimensioning. In case the demand on a link in the virtual topology exceeds the link's capacity, an instantaneous capacity adaption via a newly established parallel optical circuits takes place. Optical circuits have identical capacity, i.e. there is only one granularity. Routing changes occur only as direct consequence of bypass establishments and teardowns in the respective bypass start node. Otherwise, routing is fixed.

The choice of the thresholds for a specific approach is an optimization problem. In general, a low bypass establishment threshold t_1 will result in optical circuits with low utilization. High thresholds lead to better utilized optical circuits but increase the amount of transit traffic in the IP/MPLS layer. For comparability, we choose identical thresholds for all heuristics. The threshold for bypass establishment is 60% of the capacity of one optical circuit. The teardown threshold is 40%.

B. Simulation Framework

We implemented an event-driven simulator for dynamic optical bypasses using the Java Edition of the IKR Simulation Library [15]. This simulator is flexible with respect to the scenario and allows for integration of further heuristics. Within the simulator, we use a rather abstract model for the multilayer network. There is a notion of an optical circuit in the lower layer and a packet switching unit in the upper layer, but we neither consider the details of optical circuits (e.g. amplifiers) nor network nodes (e.g. racks, shelves, and line cards).

All simulations start with a transient phase. During this phase we do not collect statistics. Afterwards there are 10 batches each equivalent to a period of 10 days. This enables the calculation of 95% confidence intervals for mean values.

C. Reference Approaches

In addition to the already introduced dynamic bypassing heuristics, we use two further approaches as references. These approaches instantaneously adapt the link capacities but do not change the virtual topology. The first reference approach is *hop-by-hop routing* (HBH). HBH routes traffic on the shortest path in the physical topology. The second approach constructs a full mesh as virtual topology and routes all traffic on direct circuits. This approach has only add and drop traffic but no transit traffic in the IP/MPLS layer. We denote it as *full mesh* (FM).

D. Considered Metrics

The focus of dynamic bypassing is on reduction of required resources and consumed energy. In order to investigate the behavior of the heuristics with respect to these targets, we consider both layers of the multilayer network separately.

In the WSON layer, we collect statistics about the mean number of active optical circuits. We assume that it is feasible to dynamically power on and off resources in future (or to have at least some kind of power saving mode). In this case, the mean number of active optical circuits is related to the energy consumption. A further metric in this layer is the total number of used transponders in all nodes. This number indicates how much resources need to be installed for a certain traffic scenario and impacts the capital expenditures.

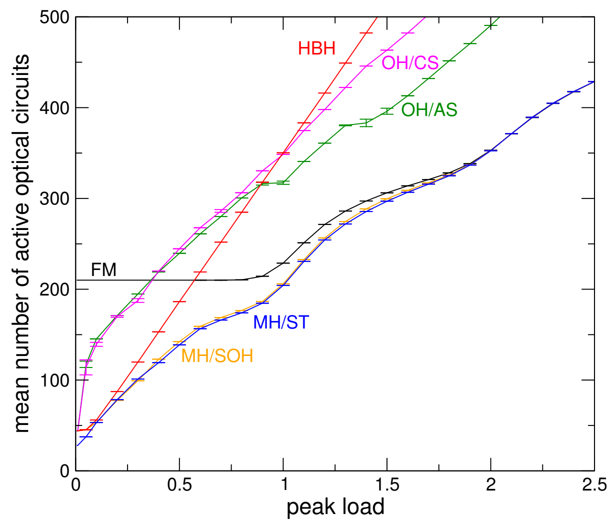

Fig. 4: Mean number of active optical circuits in the network

In the IP/MPLS layer, we concentrate on the mean transit traffic of all nodes. This transit traffic requires energy intensive processing. Avoiding the processing of such traffic comes along with energy savings. Add and drop traffic is disregarded since it is identical for all approaches.

For the network reconfiguration effort, we monitor the number of bypass switching operations. We do not consider capacity adaptations of a link in the virtual topology.

We describe the traffic demand by its peak load. If this peak load is zero, there are no demands at all. A value of one means that during the peak of the diurnal traffic pattern, there is a (mean) demand equivalent to one optical circuit between any node pair.

V. Performance Results

Within this performance evaluation the developing of the metrics for each heuristic and the relation between them is of importance. The exact values are of less interest since they strongly depend on the scenario.

A. Resources in WSON Layer

The mean number of active optical circuits in Fig. 4 is the first metric which we consider here. For very low load values, the optical circuits required for the basic topology define the mean number of active optical circuits. For FM with its full mesh, this is significantly more than for the other approaches. MH/ST has the least number of circuits since its basic topology is a spanning tree.

The mean number of active circuits grows with with increasing load. For HBH this is basically a linear increase. The utilization of once established OH/CS and OH/AS bypasses is very low during periods with little traffic but is still above the tear down threshold. Consequently, the number of optical

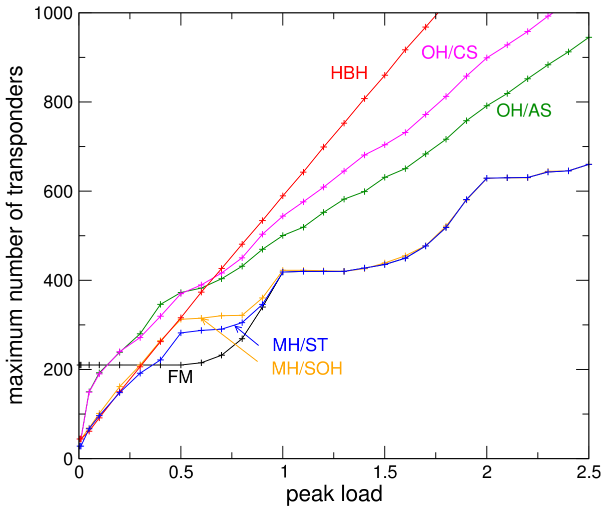

Fig. 5: Total number of used transponders in all nodes

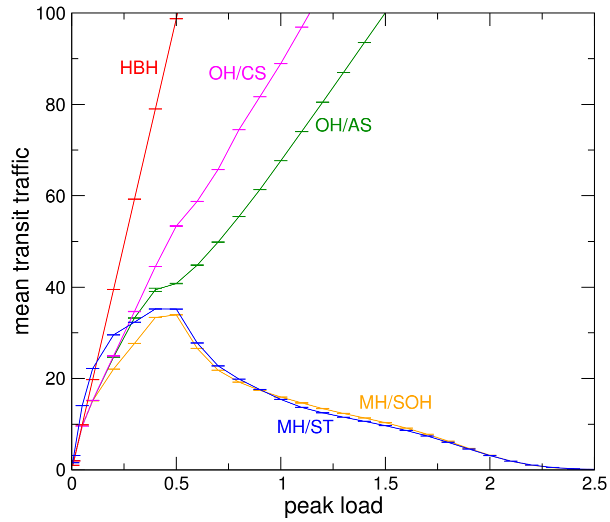

Fig. 6: Mean transit traffic in all nodes

circuits rises above that of HBH. For higher loads, both one-hop heuristics outperform HBH. Thereby, OH/AS is better than OH/CS. We identified two reasons for this. First, the aggressive bypass setup leads to less bypasses with more traffic. Thus, there is a higher multiplexing gain. Second, due to the aggressive setup, it is possible to overcome inefficient bypasses much faster than in the case of OH/CS. Although being closely related to a one-hop heuristic, MH/SOH performs much better. The reason is that this heuristic already enables for low loads longer bypasses which can be torn down again. Both multi-hop heuristics show nearly identical behavior with marginal advantages for MH/ST. These two heuristics have the lowest mean number of active circuits in the considered load range and approach that of FM for larger loads.

Newly established circuits offer capacity for further load increases. This is the reason for the bends in the developing of all bypassing heuristics and FM. For HBH, no such bends are visible as the superposition of a larger number of traffic flows blurs this effect.

The mean number of active optical circuits does not necessarily correspond to the number of transponders which are required in the nodes. Fig. 5 gives the total number of transponders and shows two substantial differences compared to Fig. 4. First, FM needs significantly less transponders for loads between 0.4 and 1 than the multi-hop heuristics. The reason is the make-before-break approach of the bypassing heuristics which requires additional resources. Second, MH/ST needs slightly less transponders than MH/SOH for low loads.

B. Resources in IP/MPLS Layer

Fig. 6 shows the accumulated mean transit traffic in the IP/MPLS layer for all nodes. By design, FM has no such transit traffic. HBH, OH/CS and OH/AS show an approximately linear increase with different slopes. Again, OH/AS performs better due to the reasons stated earlier. For low traffic loads, the multi-hop heuristics show an increase in the transit traffic comparable to the other approaches. The transit traffic for MH/ST is even a bit higher due to the spanning tree as basic topology. With increasing load, the mean transit traffic of the multi-hop heuristics decreases again. For MH/SOH and MH/ST, the values tend to zero since for very high loads there are bypasses between any node pair and thus no transit traffic at all. MH/SOH outperforms MH/ST for peak loads smaller than 0.4. This is due to the capability of MH/SOH to establish bypasses faster than MH/ST. In case of MH/SOH, any transit traffic can use a bypass. For MH/ST it has to be traffic with identical destination. For peak loads greater than 0.4, both heuristics achieve nearly identical results.

C. Bypass Switching Operations

The mean number of bypass establishments per day, which is approximately also the number of teardowns, is given in Fig. 7. We omit FM and HBH since they do not use bypasses. All bypassing heuristics show similar behavior with two peaks. The first one is around load 0.6 and the second one in the range 1.6 to 2. All heuristics tend to zero bypass switching operations for very large loads (in this case all bypasses are established in the transient phase and remain unchanged afterwards). However, the heuristics differ heavily with respect to the actual number of establishments. In general, OH/AS has the lowest number. OH/CS has at least twice the amount of switching operations in the range from load 0.6 to 2. Both multi-hop heuristics show even larger numbers of bypass establishments. Thereby MH/ST has by far the highest values.

The superposition of traffic flows in the network cause the general developing. The threshold for bypass establishment is exceeded at an increasing number of nodes with increasing load. As long as the traffic falls below the threshold for bypass teardown at the low time, these bypasses are repeatedly established and torn down. However, with further increases the traffic does not fall below this threshold anymore and the bypasses get permanent.

Around load 0.6, the bypassing heuristics can establish a bypass for each single flow. This causes the maximum peak.

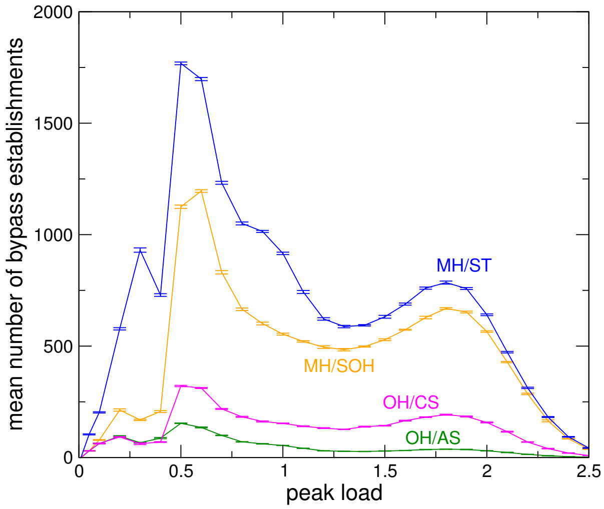

Fig. 7: Mean number of bypass establishments per day

The smaller peak around load 1.8 is due to the fact that traffic at low times is near the teardown threshold t_2.

Since multi-hop heuristics tend to bypasses with a smaller number of superposed traffic flows, they have a higher number of switching operations. Follow-up actions play another important role. Especially MH/ST causes a large number of such follow-up actions and thus has a much higher number of switching operations. With respect to the one-hop heuristics, OH/AS tends to bypasses with lots of traffic. Consequently, these bypasses have a longer lifetime. OH/CS uses more bypasses which are therefore less utilized and thus are torn down and reestablished more often. Nevertheless, this number is small compared to those of the multi-hop heuristics.

VI. Conclusion and Outlook

We investigated the performance of four different dynamic bypassing heuristics and two reference approaches. The multi-hop dynamic bypassing heuristics outperform the one-hop heuristics with respect to resource usage in almost all cases. With respect to mean active resources in the optical domain, they outperform also both reference approaches. Only the FM approach requires a smaller total number of transponders for a certain load range.

The advantage in efficiency of the multi-hop heuristics comes at the cost of a much higher number of bypass switching operations. Today, it is not foreseeable when the resulting amount will be accepted by network operators.

Among the one-hop heuristics, OH/AS performs better for all considered metrics. The OH/CS approach to leave traffic in existing bypasses is therefore void. Considering the multi-hop heuristics, MH/SOH has slight advantages regarding the required number of switching operations and the transit traffic in the IP/MPLS layer. In return, MH/ST slightly outperforms MH/SOH with respect to the required resources in the optical layer. Thus, MH/SOH seems to be the better choice at the moment.

The investigated variant of MH/ST leaves room for further optimization. A coordination with direct neighbor nodes may further increase the performance. We believe that this should impact the number of switching operations to a great extent and we want to investigate such enhancements in our future work. Further investigations with nonuniform traffic demands between the nodes should complete the picture.

Acknowledgment

The author would like to thank Jens Milbrandt and Frank Feller for valuable discussions and feedback.

The work described in this paper was carried out with the support of STRONGEST, an Integrated Project funded by the European Commission through the 7th ICT-Framework Programme under grant agreement No. 247674.

References

[1] J. Baliga, R. Ayre, K. Hinton, W. Sorin, and R. Tucker, “Energy Consumption in Optical IP Networks,” *Journal of Lightwave Technology*, vol. 27, no. 13, pp. 2391 –2403, July 2009.
[2] DE-CIX traffic statistics. [Online]. Available: http://www.de-cix.net/content/network/Traffic-Statistics.html
[3] B. Mukherjee, *Optical WDM Networks*. Springer, 2006.
[4] S. Bjørnstad, D. R. Hjelme, and N. Stol, “A highly efficient optical packet switching node design supporting guaranteed service,” in *Proceedings of ECOC 2003*, September 2003, pp. 110–111.
[5] M. Pióro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Elsevier/Morgan Kaufmann, 2004.
[6] D. Banerjee and B. Mukherjee, “Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs, and a reconfiguration study,” *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 598–607, 2000.
[7] Y. Zhang, M. Murata, H. Takagi, and Y. Ji, “Traffic-based reconfiguration for logical topologies in large-scale WDM optical networks,” *Journal of Lightwave Technology*, vol. 23, no. 10, pp. 2854–2867, 2005.
[8] R. Durán, I. de Miguel, N. Merayo, P. Fernández, J. Aguado, R. Lorenzo, and E. Abril, “Genetic algorithm to design logical topologies in reconfigurable WDM networks,” *Photonic Network Communications*, vol. 17, pp. 21–33, 2009.
[9] S. Pachnicke, H. Kagba, and P. Krummrich, “Load Adaptive Optical-Bypassing for Reducing Core Network Energy Consumption,” in *ITG-Conference ”Photonic Networks”*, Leipzig, 2011.
[10] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, “Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings,” *Optical Switching and Networking, Special Issue on Green Communications*, no. 3, pp. 181–200, July 2011.
[11] P. Tran and U. Killat, “Distributed algorithm for dynamic logical topology reconfiguration in IP over WDM networks,” in *Proceedings of ISCC 2009*, 2009, pp. 748–756.
[12] M. Ruffini, “Optical IP Switching,” Ph.D. dissertation, Computer Science Department, University of Dublin, Trinity College, 2008.
[13] STRONGEST project. [Online]. Available: http://www.ict-strongest.eu
[14] Atlanta reference network. [Online]. Available: http://sndlib.zib.de/
[15] IKR Simulation Library. [Online]. Available: http://www.ikr.uni-stuttgart.de/Content/IKRSimLib