Institute of Communication Networks and Computer Engineering
University of Stuttgart
Pfaffenwaldring 47, D-70569 Stuttgart, Germany
Phone: ++49-711-685-68026, Fax: ++49-711-685-67983
Email: mail@ikr.uni-stuttgart.de, http://www.ikr.uni-stuttgart.de

Evaluation of Centralized Solution Methods for the Dynamic Optical Bypassing Problem

Frank Feller
Institute of Communication Networks and Computer Engineering (IKR)
Universität Stuttgart, Pfaffenwaldring 47, 70569 Stuttgart, Germany
Email: frank.feller@ikr.uni-stuttgart.de

***Abstract*—Due to increasing traffic volumes and access bandwidths, the power consumption of core networks will grow considerably. Adapting network configuration to traffic load is one counter-measure. Dynamic optical bypassing is a promising approach to reconfigure multi-layer networks: it adapts the virtual topology while keeping traffic on fixed paths in the physical topology. So far, research focused on distributed bypassing schemes. In this paper, we evaluate three centralized solution methods for the bypassing problem: one based on linear programming, one based on heuristic optimization, and a greedy heuristic. We find that all methods can achieve similar energy savings while limiting changes to the network configuration.**

I. Introduction

Environmental concerns and cost pressure oblige network operators to limit the energy consumption of their transport networks. Currently, the energy consumption of core networks is small compared to other parts of the network. However, this is likely to change due to the exponential growth of core network traffic and the deployment of energy-efficient technologies like FTTx in the access network [1].

In order to meet quality of service (QoS) requirements, operators dimension network resources for estimated worst-case traffic scenarios. Due to significant traffic fluctuations, this results in typically low resource utilization. Predictable diurnal traffic profiles include night periods with traffic values as low as 25 % of the peak-hour traffic [2]. Moving from the current static operation to activating and deactivating network resources according to the load thus promises substantial energy savings.

Transport networks are generally multi-layer networks. They comprise a circuit-switched lower layer, e. g. wavelength switched optical network (WSON), which enables energy-efficient switching of traffic in the coarse granularity of optical circuits. The topmost of the upper electrical layers is packet-switched, implementing e. g. Internet protocol / multi-protocol label switching (IP/MPLS), achieving a fine switching granularity but consuming significantly more energy than optical switching. For an energy-optimal configuration, we have to trade the energy savings by switching traffic in the optical layer off against the energetic cost of operating additional (potentially lowly-loaded) optical circuits. Adapting the configuration of optical circuits to the traffic load therefore promises highly energy-efficient transport network operation. Dynamic optical bypassing (DOB) [3] is one such approach.

The speed and frequency of network resource reconfiguration is limited: For current technology, set-up times of optical circuits range in the order of minutes. In addition, network operators hesitate to move away from the static operation mode which traditionally guaranteed the high reliability of transport networks. For this reason, network reconfiguration should be limited in extent and frequency. These technological and operational constraints impede the adaptation to fast traffic fluctuations and thus reduce possible energy savings. On the upside, the limited reconfiguration frequency enables the signaling of network state information to a central entity making reconfiguration decisions based on a global view of the network domain. Such a centralized approach likely finds better configurations than a distributed scheme.

A number of publications report on centralized multi-layer network reconfiguration schemes aiming at saving energy. To our knowledge, however, our optimization heuristic [4] is alone in centrally solving the DOB problem addressed by distributed schemes in [3,5]. Efficient centralized DOB solution methods are desirable for online network reconfiguration (if combined with subsequent light-path routing) and to provide a reference for distributed strategies. In this paper, we evaluate and compare the quality of DOB solutions obtained by three different centralized methods: a linear-programming based method, a greedy heuristic, and the optimization heuristic mentioned above.

This paper is structured as follows. Section II discusses related work. We define the DOB problem in Section III. Section IV presents the three solution methods which we evaluate in Section V. We conclude in Section VI.

II. Related Work

A. Multi-Layer Network Reconfiguration

Multi-layer network reconfiguration essentially means repeatedly solving the multi-layer network optimization problem (or parts thereof) under varying conditions (e. g. time-dependent traffic load). This problem has four dimensions, which are either considered jointly to find globally optimal configurations or sequentially to limit complexity [6,7]: *(i)* definition of the virtual topology, *(ii)* routing of traffic into this topology (i. e. in the upper layer), *(iii)* routing of the light paths in the lower layer to implement the virtual topology, and *(iv)* wavelength assignment to the light paths, possibly under continuity constraints.

Reconfiguration traditionally aims at meeting QoS requirements and balancing resource utilization [6]. An additional goal is to limit the amount of modifications to the network configuration. This may be achieved by selecting the solution closest to the previous setting from a set of optimal solutions [8] or by adding a reconfiguration term to the cost function of the optimization [9]. Heuristic approaches likewise try to limit changes (e. g. [3]).

B. *Energy-Efficient Network Operation*

Since the initial work by Gupta and Singh [10] on power saving in network nodes, a lot of research has targeted the energy-efficient operation of multi-layer networks. This ranges from appropriate dimensioning [11] over routing [12] and traffic engineering [13] (the latter two focus on traffic in the upper layer) to routing and wavelength assignment (RWA) in the lower layer [14]. While [13] directly defines load-dependent routing weights, the other contributions express the problem as a mixed integer linear program (MILP) and solve it either exactly [12] or using heuristics [11,14].

Idzikowski *et al.* [15] investigate the dependency of energy savings achieved by switching off idle line cards on different degrees of freedom in reconfiguring IP-over-WDM networks. Their evaluation approach is a model for the present paper: Simulating periodic network reconfiguration to match traffic demands according to measured traces available for reference networks. The study does however disregard virtual topology reconfiguration without IP traffic rerouting as technically infeasible. Ruffini [5] argues in favor of precisely this constellation (and exploits it for a distributed reconfiguration scheme): Hiding topology changes from the IP routing mechanisms prevents instability due to the convergence of routing protocols. The decision to route certain traffic into a newly established bypass circuit can be taken locally by a network node. It only requires some additional information in its routing table. Scharf [3] evaluates several distributed DOB heuristics basing on this principle. In this paper, we complement these studies by evaluating centralized DOB methods.

III. Problem Description

DOB constitutes the optimization problem of simultaneously minimizing the energy consumption of the network for a given traffic demand matrix and the number of optical circuits set up or torn down compared to a previous network configuration. Its solution defines the virtual topology by a set of bypasses and the number of circuits on each link. By *link*, we refer to a link of the virtual topology of the upper network layer, which we characterize by its start and end nodes. If these nodes are directly connected in the physical topology, we speak of a *physical link*, otherwise of a *bypass (link)*. A *circuit on a link* is a light-path connecting the start and end nodes of that link. We currently disregard the realization of circuits in the lower layer, i. e. the RWA problem, as well as resource constraints. We enforce at least one circuit on each physical link.

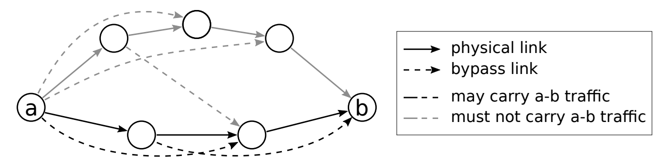

Fig. 1. Illustration of routing options along fixed path in the physical topology

One essential characteristic of the DOB problem is that traffic always follows fixed paths in the physical topology. I. e., a certain traffic demand may only be processed by nodes being on its fixed path (in their sequence on this path). The traffic can however omit some nodes by using bypasses interconnecting distant nodes on the path (cf. Fig. 1). By *routing*, we refer to the remaining degree of freedom of choosing a combination of the admissible links. This routing may be determined by optimization. We currently define the fixed traffic paths as the shortest paths in the physical topology in terms of hop count. Ties are broken in favor of the geographically shorter path.

We assume that the optimization is periodically executed to reconfigure the network according to varying traffic demands. The input parameters comprise the physical network topology (i. e. the graph of all nodes and the physical links), the number of active circuits on each link prior to the reconfiguration, and the directed traffic demands between each disjoint pair of nodes. The demands indicate the (estimated) maximum traffic rates in the time interval the new network configuration applies to. Like [15], we assume that these values are known. We further assume that we can arbitrarily split demands when routing them onto different links. This is justified since demands between core network nodes are aggregates of many transport connections which we can control by traffic engineering.

We use an abstract equipment energy model, expressing the network configuration-dependent part of energy consumption by the number of active optical circuits n_C and the amount of electrically switched transit traffic t_T. We disregard static energy consumption as well as contributions of tributary interfaces and add/drop traffic, since they are unaffected by reconfiguration. Such a model applies e. g. to IP/MPLS routers with line cards terminating single optical circuits, where line cards of unneeded circuits are switched off along with the respective transponders. If active, these components have a power consumption of α per circuit. In addition, the energy consumption of packet processors can scale with the load due to mechanisms like frequency scaling. For simplicity, we assume a linear dependency, denoting the energy consumption per switched traffic unit by β. The model disregards the power consumption of the remaining optical equipment, which is generally comparatively small. One benefit of this model is that the ratio of the energy consumption values (α/β) is its only parameter which depends on the technology.

Besides minimizing energy consumption, we want to limit the extent of modifications. Like [9], we do so by adding a reconfiguration penalty to the cost function. This penalty is the number of newly established or torn-down circuits r compared to the previous network configuration, which we weight by γ $(0 \leq \gamma < \alpha)$. A penalty weight greater or equal to the energetic cost of a circuit would prevent the teardown of unused circuits and thus contradict the idea of DOB.

Altogether, we aim at minimizing the following cost function:

$$\alpha \cdot n_C + \beta \cdot t_T + \gamma \cdot r \tag{1}$$

IV. Solution Methods

A. Mixed Integer Linear Programming

The MILP formulation of the DOB problem formalizes the definition in Section III. It jointly optimizes the number of active circuits on all links and the routing of traffic demands along their fixed paths. We adapted the formulation as a splittable multi-commodity flow problem from [15]. To reduce the problem complexity, we also aggregate all demands originating from one node into one commodity.

1) Parameters: Let V be the set of nodes of the given network topology. We define the set of all possible directed links $L = \{(i,j) \in V \times V \mid i \neq j\}$. Let $E \subseteq L$ be the set of directed physical links. Let further $d_{ij} \in \mathbb{R}_0^+$ be the current value of the directed traffic demand from node i to node j and let $p_{ij} \in \mathbb{Z}_0^+$ be the number of active circuits on link $(i,j) \in L$ in the previous configuration. Let $K \subseteq V$ designate the set of commodities, i. e. of nodes being the source of at least one demand.

Let further $L_{ij} \subset L$ be the set of all links that traffic from node i to node j may use along its fixed path (cf. Fig. 1). We define the set of links traffic of commodity $k \in K$ may use:

$$T_k = \bigcup_{j \in V \setminus \{k\}} L_{kj} \tag{2}$$

As we route traffic along the shortest path in the physical topology, the paths of all demands originating from node k form a source tree rooted at this node. Hence, routing traffic from nodes k to j on links from T_k equals routing on L_{kj}. We further define the net demand values for each commodity $k \in K$ and each node $n \in V$:

$$d_n^k = \begin{cases} \sum_{j \in V} d_{nj} & \text{for } n = k \\ -d_{kn} & \text{for } n \neq k \end{cases} \tag{3}$$

Let finally $C \in \mathbb{R}^+$ be the capacity of one circuit and let $\alpha, \beta, \gamma \in \mathbb{R}_0^+$ denote the cost coefficients as in Eq. (1).

2) Variables: The variables $f_{ij}^k \in \mathbb{R}_0^+, (i,j) \in T_k$ describe the flow of commodity $k \in K$ on the directed link from node i to node j. Variables $c_{ij} \in \mathbb{Z}_0^+, (i,j) \in L$ represent the number of active circuits on the respective link, and $r_{ij} \in \mathbb{Z}_0^+$ give the number of circuits torn down or newly set up on this link.

3) Model: The MILP consists of the objective function in Eq. (4) and the following constraints: Eq. (5) stipulates flow conservation; Eq. (6) ensures a sufficient number of circuits to carry all traffic on each link; Eq. (7) enforces at least one circuit on each physical link; Eqs. (8) and (9) finally define the number of circuits established or torn down.

$$\min \alpha \sum_{(i,j) \in L} c_{ij} + \beta \sum_{k \in K} \sum_{(i,j) \in T_k : i \neq k} f_{ij}^k + \gamma \sum_{(i,j) \in L} r_{ij} \tag{4}$$

$$\sum_{j \in V : (n,j) \in T_k} f_{nj}^k - \sum_{i \in V : (i,n) \in T_k} f_{in}^k = d_n^k \qquad \forall n \in V, k \in K \tag{5}$$

$$C c_{ij} - \sum_{k \in K : (i,j) \in T_k} f_{ij}^k \geq 0 \qquad \forall (i,j) \in L \tag{6}$$

$$c_{ij} \geq 1 \qquad \forall (i,j) \in E \tag{7}$$

$$r_{ij} - c_{ij} \geq -p_{ij} \qquad \forall (i,j) \in L \tag{8}$$

$$r_{ij} + c_{ij} \geq p_{ij} \qquad \forall (i,j) \in L \tag{9}$$

B. Optimization Meta-Heuristic

The simulated-annealing (SA) based method we proposed in [4] solves a variant of the DOB problem that differs from the MILP by not optimizing demand routing. It routes each demand as a whole into a sequence of links yielding the least number of hops with electrical processing. This may result in suboptimal configurations with lowly utilized circuits when the traffic on some bypass slightly exceeds the capacity of one or several circuits.

For this study, we added an optional post-processing step addressing this issue. It iterates over all bypasses in descending order of their length in terms of hops in the physical topology. For each bypass, it tries to reroute the share of traffic exceeding an integer number of circuits (the *excess traffic*): We exclude the considered bypass and select one of the shortest sequences of remaining links along the shortest physical path from its start node to its end node. If the energetic cost of electrically processing the excess traffic in the nodes along this sequence exceeds the cost of operating one circuit, we do not reroute. Otherwise we continue by verifying whether each link of the sequence can accommodate the excess traffic without using an additional circuit, which would offset the benefit of avoiding the original bypass circuit. If so, we reroute the excess traffic over this link sequence. We otherwise exclude the links with insufficient spare capacity and repeat the procedure (from the selection of the shortest link sequence on). If no other termination condition is met earlier, we stop when the exclusion of links leaves the start and end nodes disconnected.

C. Greedy Elimination Heuristic

We extended the idea of the post-processing procedure above into a greedy stand-alone DOB heuristic. It starts from a full mesh of links, i. e. each demand is initially routed over a direct link from its source to its destination. It then removes uneconomical bypass circuits similarly to the post-processing. We modified this procedure to factor in the reconfiguration penalty: *(i)* the cost of the circuit to be removed by rerouting now depends on its existence in the previous configuration; *(ii)* additional circuits on the alternative path are tolerated if pre-existent and favorable in terms of overall cost. Fig. 2 describes the algorithm in pseudo-code (using symbols defined in Section IV-A). It returns the network configuration in terms of the number of circuits on each link.

V. Evaluation

A. Simulation Setup

We evaluated the solution quality of the DOB methods by event-driven simulation based on the Java edition of the IKR

```
B ← L \ E                                              {set of all bypass links}
f_ij ← d_ij   ∀(i, j) ∈ L                     {initial traffic flows on all links}
for all (i, j) ∈ B in order of descending bypass length do
  t ← f_ij mod C                                  {excess traffic on bypass}
  if f_ij/C > p_ij then                    {bypass circuit not pre-existent}
    c_l ← α + γ                                        {cost of bypass circuit}
  else
    c_l ← α − γ
  end if
  S ← L_ij \ {(i, j)}            {set of candidate links for rerouting}
  loop alternative_path_search
    P ← links of shortest path from i to j in graph G = (V, S)
    c_r ← (|P| − 1) · βt  {(switching) cost on alternative path}
    if P = ∅ or c_r > c_l then              {no path or higher cost}
      break alternative_path_search                 {do not reroute t}
    end if
    feasible ← true
    for all (m, n) ∈ P in sequence i → j do
      if (f_mn mod C) + t > C then
        if p_mn > f_mn/C and c_r + α − γ < c_l then
          {additional circuit pre-existant and within budget}
          c_r ← c_r + α − γ
        else
          S ← S \ {(m, n)}                       {retry without this link}
          feasible ← false
          break for all (m, n)
        end if
      end if
    end for
    if feasible then                          {reroute excess traffic to P}
      f_ij ← f_ij − t
      for all (m, n) ∈ P do
        f_mn ← f_mn + t
      end for
      break alternative_path_search
    end if
  end loop
end for
return c_ij ← ⌈f_ij/C⌉   ∀(i, j) ∈ L
```

Fig. 2. Pseudo-code of greedy elimination heuristic

SimLib [16]. While both heuristics are directly implemented in Java, the simulator calls SCIP [17] to solve the MILP.

1) Scenario: We present results for the Géant reference network with 22 nodes and 36 links, which is available from SNDlib [18] along with dynamic demand matrices obtained by measurement over four months [19]. From this trace of demands, we selected and concatenated the profiles of 10 different working-days as input for the simulation studies. Based on the assumption that these 10 day profiles show a similar statistical behavior, we compute 95 % confidence intervals (which are thus rather conservatively estimated) for all metrics. In accordance with the assumed technological constraints, we reconfigure the network every 15 minutes, which also corresponds to the granularity of the demand trace.

In order to vary the traffic load, we scale all demand matrices of the trace by one factor. We quantify the scaling based on a *peak demand matrix* containing the maximum values of the traffic demands between every node pair over the course of the 10 days. We characterize the traffic load by the average of these peak demands. Like all traffic values, we express this demand relative to the capacity of one optical circuit (in *circuit equivalents*), since this ratio has a decisive impact on the DOB performance.

In this study, we vary the average peak demand between 0.1 and 4 circuit equivalents. Assuming a circuit capacity of 40 Gbps, this translates into a total peak demand (sum over all demands in the peak matrix) between 1.85 Tbps and 73.2 Tbps. The corresponding time-averaged total demand ranges between 558 Gbps and 22.3 Tbps.

2) Reference Configurations: To provide a rough indication of the energy saving potential of the DOB principle, we include results for two baseline configurations: Firstly, the static operation of all resources dimensioned for the peak demand matrix. We determine the network configuration by our MILP, i. e. optimizing virtual topology and demand routing along fixed paths. We refer to this case by always-on (AO).

Secondly, we consider static bypassing (SB) with dynamic resource operation. For this, we determine a fixed virtual topology by executing the SA-based algorithm (without demand splitting) for the peak demand matrix. During operation, all traffic follows fixed (shortest) paths in this topology, but we deactivate unneeded parallel circuits and we let electrical processing scale with the actual transit traffic. Except for our restrictions on routing during the initial optimization, this strategy corresponds to FUFL in [15].

3) Parameterization: The cost function is the primary common feature of the DOB methods. We assume that the energetic cost of switching one circuit worth of traffic electrically corresponds to the cost of operating one optical circuit, and we normalize the cost to this value: $\alpha = \beta = 1$. Given the chosen traffic unit, the circuit capacity is $C = 1$. We vary the reconfiguration penalty γ between 0 and 0.85.

With the optimization-based schemes, limiting the computation time is essential in view of online application. Solving the MILP on commodity hardware, we set a time limit of 1 minute, which results in an average optimality gap of up to 2 % for $\gamma = 0$ and of less than 0.5 % for $\gamma = 0.5$. For the SA-based strategies, we applied the improvement-based termination condition detailed in [4]. It resulted in computation times of 4 to 8 minutes on the same hardware.

B. Results

Like [3], we evaluate the quality of the solutions obtained by the DOB methods along the terms of the cost function: For the energy consumption, we separately consider the average number of optical circuits and the average amount of electrically switched transit traffic. Regarding network reconfiguration, we evaluate the average number of circuit modifications (i. e. of optical circuits established and torn down) per reconfiguration event, i. e. per 15 minute interval.

Using these criteria, we compare the MILP-based optimization approach (MILP), the heuristic SA-based optimization with post-processing for demand splitting (SA-PP) and without this step (SA), and the greedy elimination heuristic (GEH) with the non-DOB references AO and SB. We give results for

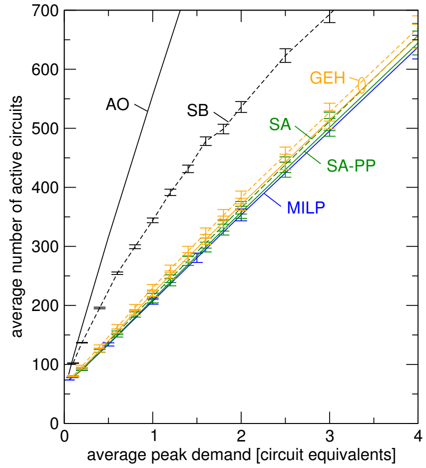

Fig. 3. Mean number of active circuits

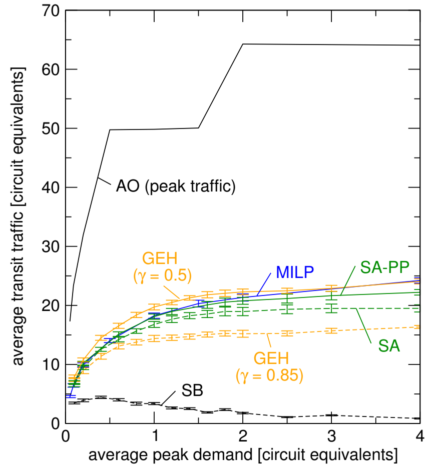

Fig. 4. Mean amount of transit traffic

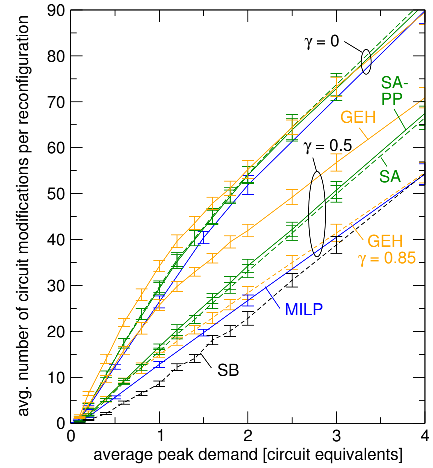

Fig. 5. Mean number of circuit modifications

$\gamma = 0.5$ for all DOB methods and additionally for $\gamma = 0.85$ for GEH.

1) Active Circuits: Fig. 3 plots the time average of the number of active circuits in the network over the load. For low load, all curves converge to the minimum number of 72 implied by the physical topology. With increasing load, the metric grows almost linearly for all DOB methods and for the AO configuration. This implies that the effectiveness of virtual topology optimization does not depend on the load in our scenario – due to a wide variety in the scale of the demand values. Since the AO configuration is governed by potentially singular peak demand values, much caution is required in generalizing the energy savings relative to it. For low load, SB resembles AO since few parallel circuits allow adaptation. For high load, several circuits are needed on many links to carry the peak traffic, allowing their deactivation during off-peak hours.

All DOB methods produce a similar result, the difference being of at most 10 % and at the limit of statistical significance. The order of the average circuit numbers is however instructive: MILP optimizing the routing of splitted demands performs best, closely followed by SA-PP doing splitted demand routing in a non-optimized way. SA (dashed line) is next, suffering from not splitting demands. GEH (with $\gamma = 0.5$, solid line) closely follows SA: while it can split demands, it faces the handicap of not optimizing the bypass setting. The comparatively high circuit number for GEH with $\gamma = 0.85$ (dashed line) confirms our expectation that limiting reconfiguration comes at some cost in terms of energy consumption.

2) Transit Traffic: The time-averaged amount of electrically switched transit traffic across all nodes is depicted over the load in Fig. 4. For DOB, it initially increases with the load, but eventually saturates as the additional traffic is absorbed by bypass circuits. The same applies for the processing capacity dimensioned for the peak demand, which is plotted for AO. The high number of bypass links persisting in off-peak hours in the SB configuration results in minimal transit traffic.

The DOB methods again show similar performance, but not as close as for the number of circuits. With the exception of GEH with $\gamma = 0.5$, the order of the methods by transit traffic is inverse to the order by active circuits. This is in line with the tradeoff between circuits and transit traffic when configuring a network.

3) Reconfigurations: Since increasing peak demands bring about higher traffic variations in absolute terms, load-dependent network reconfiguration affects more circuits. This is reflected in all plots of the average numbers of circuit modifications per reconfiguration event over the load in Fig. 5.

Without reconfiguration penalty ($\gamma = 0$), all DOB methods produce similar, high numbers of circuit modifications. The slope of these curves slightly decreases for average peak demands exceeding 1.5 circuit equivalents since certain bypass circuits turn permanent due to increasing off-peak load. The effect of setting a reconfiguration penalty of $\gamma = 0.5$ differs between the DOB methods: For SA and SA-PP, it reduces circuit modifications by 25 % to 50 %. For MILP, it results in reductions by 40 % to 50 % due to rerouting avoiding circuit modifications. Therewith, MILP approaches the number of modifications incurred by simply deactivating unneeded circuits with SB. For GEH, the penalty hardly reduces modifications for low loads, and the gains remain inferior to SA for high loads. As illustrated with $\gamma = 0.85$, applying a higher penalty with GEH can however reduce modifications to the levels of SA and MILP.

C. Discussion

By principle, the effects of network reconfiguration strongly depend on network and traffic properties. However, partial studies with different settings give us reason to believe that our observations are transferable to other realistic scenarios.

Under our assumptions, all DOB solution methods achieve approximately the same energy cost for one set of cost function parameters. Comparing Figs. 3 and 4, we find that the load-dependent part of the power consumption is dominated by the operation of optical circuits[1].

The energy savings of SB are in line with the observations for FUFL in the Géant network in [15] – when considering that their highest traffic load scenario corresponds to an average peak demand of 0.27 circuit equivalents. Idzikowski *et al.* further observe that dynamically optimizing IP traffic routing yields substantial gains in energy efficiency, whereas additionally reconfiguring the virtual topology does hardly bring additional benefit. The energy savings they achieve by rerouting range in the same order as ours obtained by DOB. This suggests that multi-layer network reconfiguration bears a certain potential for energy savings, which is realizable by *either* of rerouting or topology reconfiguration. A combination of both does not enable additional savings. However, systematic studies are needed to validate this statement.

While energy savings by network reconfiguration necessarily involve circuit modifications, it is possible to substantially reduce the number of modifications without compromising energy efficiency: Energy metrics for $\gamma = 0$ proved close to those plotted for $\gamma = 0.5$ in Figs. 3 and 4. The effect of a certain γ value significantly differs for different DOB methods. We thus have to tune the penalty for the respective method.

While all DOB methods can produce solutions of similar quality, it is noteworthy that MILP combines the best energy efficiency with the lowest extent of circuit reconfiguration. This is due to the degree of freedom it exploits in routing. In case this rerouting has a negative effect on network operation, we may need to add a rerouting penalty.

VI. Conclusion

In this paper, we first defined the DOB problem, which optimizes the virtual topology of multi-layer networks while routing traffic along fixed paths in the physical topology. We then presented three centralized solution methods for this problem: one based on linear programming, one based on heuristic optimization, and a greedy heuristic. We finally evaluated these methods in terms of the achieved energy efficiency and the required circuit modifications when periodically adapting the network configuration to the load by means of simulation.

All three solution methods achieve comparable energy efficiency in terms of active optical circuits and electrically switched transit traffic. Compared to load-dependent resource operation in a fixed virtual topology with fixed routing, DOB can reduce the load-dependent share of the power consumption by 20 % to 35 %.

We define the DOB problem to include a reconfiguration penalty allowing to control the extent of network reconfiguration. We observe that this penalty can reduce the number of circuits established and torn down by up to 50 % without significant impact on energy consumption. However, the effect of a certain penalty value varies between the solution methods.

[1] We may therefore compare our results with studies like [15] which only evaluate the power consumption of active line cards (which correspond to circuits).

Future work will complement the virtual topology definition methods by light-path routing to realize bypass circuits. Such an extension enables the consideration of resource constraints in the optical layer. An evaluation of the resulting solution methods could additionally compare the benefits of DOB to other network reconfiguration schemes.

Acknowledgment

The author would like to thank Joachim Scharf, Jens Milbrand, and David Wagner for valuable discussions.

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. V. Sorin, and R. S. Tucker, “Energy consumption in optical IP networks,” *J. Lightw. Technol.*, vol. 27, no. 13, pp. 2391–2403, Jul. 2009.

[2] “DE-CIX traffic statistics,” http://www.de-cix.net/content/network/Traffic-Statistics.html.

[3] J. Scharf, “Efficiency analysis of distributed dynamic optical bypassing heuristics,” in *Proc. Workshop on Green Commun. and Netw. at ICC*, Jun. 2012.

[4] F. Feller, “An optimization-heuristic approach to dynamic optical bypassing,” in *Proc. 13th ITG Symp. on Photonic Netw.*, Leipzig, Germany, May 2012.

[5] M. Ruffini, “Optical IP Switching,” Ph.D. dissertation, Computer Science Department, University of Dublin, Trinity College, 2008.

[6] B. Mukherjee, *Optical WDM Networks*, ser. Optical Networks. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

[7] M. Pióro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2004.

[8] D. Banerjee and B. Mukherjee, “Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs, and a reconfiguration study,” *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 598–607, Oct. 2000.

[9] P. N. Tran and U. Killat, “Dynamic reconfiguration of logical topology for WDM networks under traffic changes,” in *Proc. IEEE Netw. Operations and Manag. Symp. (NOMS)*, Apr. 2008, pp. 279–286.

[10] M. Gupta and S. Singh, “Greening of the Internet,” in *Proc. Conf. on Applicat., Technol., Architectures, and Protocols for Computer Commun.*, ser. SIGCOMM ’03. New York, NY, USA: ACM, 2003, pp. 19–26.

[11] G. Shen and R. Tucker, “Energy-minimized design for IP over WDM networks,” *Optical Commun. and Netw.*, vol. 1, no. 1, pp. 176–186, Jun. 2009.

[12] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, “Power awareness in network design and routing,” in *Proc. 27th IEEE Conf. on Computer Commun. (INFOCOM)*, Apr. 2008, pp. 457–465.

[13] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Power reduction techniques in multilayer traffic engineering,” in *Proc. 11th Intl. Conf. on Transparent Optical Netw., (ICTON)*, Jul. 2009, pp. 1–4.

[14] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, “Power-aware routing and wavelength assignment in optical networks,” in *Proc. 35th European Conf. on Optical Commun. (ECOC)*, Sep. 2009, pp. 1–2.

[15] 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 Netw.*, vol. 8, no. 3, pp. 181–200, 2011.

[16] IKR Simulation Library. [Online]. Available: http://www.ikr.uni-stuttgart.de/Content/IKRSimLib

[17] T. Achterberg, “SCIP: Solving constraint integer programs,” *Mathematical Programming Computation*, vol. 1, no. 1, pp. 1–41, Jul. 2009.

[18] S. Orlowski, M. Pióro, A. Tomaszewski, and R. Wessäly, “SNDlib 1.0–Survivable Network Design Library,” *Netw.*, vol. 55, no. 3, pp. 276–286, 2010, http://sndlib.zib.de.

[19] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, “Providing public intradomain traffic matrices to the research community,” *SIGCOMM Comput. Commun. Rev.*, vol. 36, pp. 83–86, Jan. 2006.