Universität Stuttgart

INSTITUT FÜR KOMMUNIKATIONSNETZE UND RECHNERSYSTEME
Prof. Dr.-Ing. Andreas Kirstädter

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 a Centralized Method for One-Step Multi-Layer Network Reconfiguration

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. Technological and operational limitations thereby require transitioning between subsequent configurations in a single step. In this paper, we define the according reconfiguration problem with resource-preoccupation constraints, and we propose and evaluate a meta-heuristics based solution method for this problem. We obtain energy savings of up to 35 % compared to a simple resource adaptation scheme and only observe rare events of traffic blocking in particular situations.**

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 access technologies like FTTx [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 typically results in 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 off the energy savings by switching traffic in the optical layer and the energetic cost of operating additional (potentially lowly-loaded) optical circuits. Network reconfiguration adapting active optical circuits to the traffic load therefore promises highly energy-efficient transport network operation.

The reconfiguration process is subject to two types of time constraints. On the one hand, a network configuration has to accommodate the traffic occurring while it is active. Approaches aiming at maximal energy savings will rely on traffic forecasts based on the current load in order to closely follow actual needs. While this forecasting is a non-trivial problem beyond the scope of this paper, it arguably becomes more difficult with increasing prediction horizon. On the other hand, technological constraints limit the speed of adapting the configuration of optical networks. In particular, a sudden change of the number of wavelength division multiplex (WDM) channels in an optical fiber impairs the transmission quality on persisting channels due to transient effects in widely-deployed erbium doped fiber amplifiers (EDFA) [3]. One remedy is to gradually adapt the signal power of WDM channels, resulting in reconfiguration times in the order of several minutes.

Most reconfiguration schemes ensure that resulting network configurations respect hardware resource constraints. Many of these schemes compute a transition path, i. e. a sequence of adaptation steps to reach the new configuration [4]. Given the oppositional time constraints, we cannot afford consecutive steps in reconfiguring optical circuits, but we have to reach the new configuration in one step. Since quality of service (QoS) requirements disallow service interruptions in core networks, new circuits need to be established to accommodate rerouted traffic before old circuits are torn down. Consequently, resources occupied by discontinued circuits in the previous configuration cannot be used in the next configuration. Such a *hitless reconfiguration* concept is also applied in [5].

A number of publications report on multi-layer network reconfiguration aiming at saving energy. To our knowledge, however, there is no energy-oriented hitless multi-layer network reconfiguration scheme. In this paper, we propose and evaluate a one-step reconfiguration method derived from a dynamic optical bypassing (DOB) method presented in [6,7].

This paper is structured as follows. Section II discusses related work and details the one-step reconfiguration principle. We define the resulting network reconfiguration problem in Section III. Section IV presents our solution method, which we evaluate in Section V. We conclude in Section VI.

II. Multi-Layer Network Reconfiguration

A. Related Work

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 [8,9]: *(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.

Traditionally, network reconfiguration aims at meeting QoS requirements and balancing resource utilization [8]. The additional goal of limiting the amount of modifications has been pursued by different approaches: Only choosing one optical circuit to modify per reconfiguration event [10], selecting the solution closest to the previous setting from a set of optimal solutions [11], or adding a reconfiguration penalty to the cost function of the optimization [12]. Additionally following the make-before-break principle, the latter work is of particular relevance for this paper.

Several publications address energy-efficient dynamic operation of multi-layer networks under centralized control. By computing independent optimal settings, the energy-saving potential of different scopes of periodic multi-layer network reconfiguration is estimated in [13]. A reconfiguration penalty in the cost function relates consecutive configurations in [14] as well as in our previous work [6,7]. In this paper, we extend these schemes to achieve hitless reconfiguration.

B. One-Step Reconfiguration

We assume a strictly periodic and sequential regime for our hitless one-step network reconfiguration, where each configuration is valid for a time period of ΔT. Fig. 1 illustrates the according time flow. The upper part describes the actual network configuration: once the next configuration is determined, optical circuits required in addition to the active configuration are set up (cf. striped triangles). This requires a non-negligible amount of time. Upon completion, the traffic is rerouted as provided by the new configuration, freeing circuits unused in this configuration, which are then torn down slowly. While rerouting may not be instantaneous throughout a network domain, the time taken by this procedure is small compared to circuit reconfiguration times and hence neglected here. Due to the sequential reconfiguration procedure, the reconfiguration interval ΔT is bounded below by the sum of circuit setup and teardown times.

The lower part of Fig. 1 addresses the computation of new configurations. Computation can start as soon as the input values are available, i. e. the traffic forecast and the previous configuration. The latter bounds the computation time above to ΔT. To exploit this maximum time, traffic forecasts need to be available $1.5\Delta T$ before the respective configuration turns active, corresponding to a total forecast horizon of $2.5\Delta T$. One could reduce the required forecast horizon at the expense of more restricted reconfiguration by allowing consecutive reconfiguration steps to overlap. However, we leave the investigation of this aspect to future work.

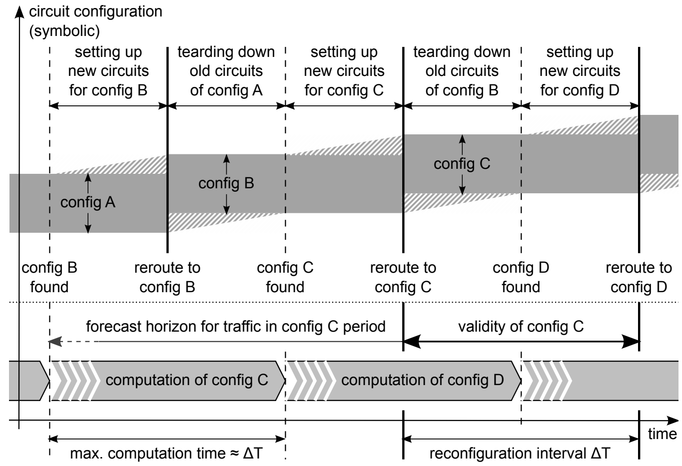

Fig. 1. Time flow of one-step reconfiguration procedure

III. Reconfiguration Problem Description

We derive our problem from DOB [7] by omitting routing and minimal circuit set constraints mimicking distributed approaches. In exchange, we add resource and pre-occupation constraints and consider circuit realization.

We assume an IP/MPLS-over-WSON network. Within the transparent optical reach, any two IP/MPLS nodes are connectable by optical circuits. Each circuit is routed over a series of links in the physical topology and terminated in line card ports of the IP/MPLS nodes. Along this path, we assume full wavelength conversion capability. The set of optical circuits defines the virtual topology of the upper IP/MPLS layer. We consider traffic demands between pairs of IP/MPLS nodes and route them in this virtual topology without path length restriction. We assume that we can arbitrarily split demands to route them on different paths. This is justified since the demands are aggregates of many flows which we can control by traffic engineering. We disregard protection.

In accordance with current network node architecture, we assume that line card interfaces consist of pairs of one input port and one output port to maintain bi-directional circuits. While we require one such pair of ports to connect to *one* opposite port pair of another node, we allow the activation of only one of the circuits to efficiently accommodate asymmetric traffic. We allow a port pair to connect to any other node.

We consider constraints on the number of port pairs in each node and on the fiber capacity of each physical link.

For each network reconfiguration interval, we solve the optimization problem of simultaneously minimizing *(i)* the energy consumption of the network for a given traffic demand matrix, *(ii)* the number of optical circuits set up or torn down compared to the previous network configuration, and *(iii)* traffic blocking under resource availability and pre-occupation constraints. A solution defines a set of active optical circuits along with the resources they occupy and a set of demand routes specified as sequences of virtual links supported by optical circuits.

The input parameters for the optimization comprise the dimensioned physical network, the set of active circuits in the previous configuration (including the resources they occupy), and the directed traffic demands between each disjoint pair of nodes. The demands indicate maximum traffic rates in the time interval the new network configuration is valid. Like [13], we assume that these values are known.

We use an abstract equipment power 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. Such a model assumes future IP/MPLS routers able to scale part of their power consumption with the number of active interfaces. We denote the resulting power consumption per circuit by α. In addition, mechanisms like frequency scaling let the energy consumption of packet processors scale linearly with the load. We denote the energy consumption per switched traffic unit by β. We disregard static energy consumption as well as contributions of tributary interfaces and add/drop traffic, since they are unaffected by reconfiguration. Due to its comparatively small power consumption [15], we likewise disregard the remaining optical equipment.

During gradual setup and teardown, circuits consume energy. We account for this by including the number of respective circuits r in the cost function, weighted by δ ($0 \leq \delta < \alpha$). The factor has to be less than the energetic cost of active circuits to maintain an incentive for circuit adaptation. In the cost function, this term constitutes a reconfiguration penalty as used in [7,12].

Due to QoS requirements, traffic blocking is hardly tolerable in core networks. We can however not strictly prohibit traffic loss by a constraint since the pre-configuration of resources may prevent one-step reconfiguration from reaching a setting satisfying all traffic demands although the installed resources would suffice in a more favorable configuration. We hence add a blocking penalty to the cost function targeting two objectives: *(i)* prevent solutions blocking even minimal amounts of traffic despite available resources and *(ii)* if blocking is unavoidable, favor the configuration resulting in the least amount of lost traffic. To achieve the former, we apply a penalty of $\mu \gg \alpha$ for each of the n_B virtual links not featuring enough circuits. For the latter, we weight the blocked traffic volume t_B by $\nu \gg \alpha$. Altogether, we aim at minimizing the following cost function:

$$\alpha \cdot n_C + \beta \cdot t_T + \delta \cdot r + \mu \cdot n_B + \nu \cdot t_B \qquad (1)$$

IV. Solution Method

Since light-path modification times are large compared to signaling delays even in large networks, we assume that network state information is communicated to a central entity which determines the next configuration. Having a global view of the network domain, optimization by such central entity is generally able to find better configurations than distributed schemes.

Except for wavelength assignment, the problem in Section III covers all dimensions of multi-layer network optimization. The energetic terms of the cost function, however, are governed by the virtual topology and the demand routing of the upper layer (since these define the number of required optical circuits). Port pair constraints are likewise considered in the upper-layer optimization. Light-path routing is of relevance only for fiber capacity constraints. Due to the rather high capacity of a single fiber, these constraints proved unproblematic in relevant scenarios. We therefore only optimize the configuration of the upper layer and apply a simple heuristic for light-path routing.

We solve the upper-layer centric optimization problem using a modified version of the Simulated Annealing-based method detailed in [6] along with a post-processing heuristic similar to [7]. For completeness, we outline these procedures and highlight the modifications yielding our *virtual topology centric reconfiguration* (VTCR) method.

A. Meta-Heuristic Optimization Procedure

Optimization meta-heuristics perform a randomized search of the solution space. Simulated Annealing (SA) [16] does so by iteratively modifying, or *perturbating*, one solution. It controls this search process by a *temperature* parameter defining the probability to accept solutions of higher cost. If a candidate solution is rejected, the next perturbation starts from the previously accepted solution.

Compared to [6], we introduce two features to accelerate computation. First, the temperature is not only reduced after a fixed number of perturbations, but already after a certain (smaller) number of solutions have been accepted (as done in [17]). Second, the termination condition does no longer require a given number of perturbations without any accepted solution, but it is satisfied if the maximum relative cost difference of accepted solutions in this period is sufficiently small.

B. Resource States

In order to keep track of the resource occupation in the previous configuration while incrementally modifying the new setting in perturbations, we assign states according to Fig. 2 to circuits and associated resources. Dependent on their use in the previous configuration, resources are initially either *on* or *off*. Perturbations may deactivate pre-existing circuits (putting their resources in *teardown* state) or allocate unused resources to create new circuits (putting them into *setup* state). Subsequent perturbations may undo such changes, returning *teardown* circuits / resources to *on* and releasing the allocation for *setup* circuits whose resources turn *off* again. We refer to the transitions between these pairs of states by *activation* resp. *deactivation*. When a new configuration is finally applied to the network, the *setup* and *teardown* states transition to their definite counterparts to prepare for the next optimization run.

C. Perturbation and Cost Computation

Our SA based method essentially optimizes the virtual topology while routing demands deterministically in this topology. A perturbation consists in adding or removing one directed virtual link. Such a link can connect any pair of nodes, i. e. we omit minimum active link constraints from [6]. We however exclude virtual links whose shortest light-path realizations would exceed the transparent optical reach.

In the perturbated virtual topology, we route all traffic demands along one of their shortest paths in terms of hop count. We thereby determine the traffic load on each virtual link and count the number d_B of demands unroutable due to partitioning of the virtual topology.

Next, we iterate over all virtual links in increasing order of the length of their shortest light-path realizations and adjust the circuit configuration to the traffic load. If more circuits are active than required, we preferentially eliminate circuits in *setup* state and proceed to deactivating *on* circuits if necessary. If the circuit capacity is insufficient, we retain the virtual link. After releasing all unneeded resources, we re-iterate over the retained links and first reactivate circuits from *teardown* state. If the aggregate capacity on the link remains insufficient, we try to allocate resources and set up new circuits. If this fails, n_B and t_B are increased accordingly.

To set up a new circuit, we first allocate ports on the source and target node, preferably in port pairs sustaining a circuit in the opposite direction. If successful, we route the light path: We determine the shortest path in the physical topology and try to allocate WDM channels on each link. If resources are depleted on some link, we exclude this link and repeat the procedure on the remaining graph. We terminate successfully if we have found a feasible path. We abort if the length of the physical path exceeds the transparent optical reach or the exclusion of links leaves source and target node in different graph partitions.

The cost is finally computed according to Eq. 1, where n_C counts all circuits in *setup* and *on* state and r is obtained as the number of circuits in *setup* and *teardown* state. We add the number of blocked demands d_B weighted by a factor of $\eta \approx 2\mu$ in order to discourage partitioning, which we do not address by post processing.

D. Post-Processing

Not optimizing demand routing in the SA run may entail two issues: First, traffic may be blocked due to insufficient capacity on the shortest path while spare capacity exists on alternative paths. Second, energetically suboptimal configurations with lowly utilized circuits result when the traffic on some virtual links slightly exceeds the capacity of one or several circuits. We address these issues in this order by the following post-processing heuristic.

To be able to use free line card ports in port pairs maintaining unidirectional circuits, we start by adding all missing inverse links to the virtual topology resulting from the optimization. We then iterate (in increasing order of shortest light-path length) over all virtual links with insufficient circuit capacity. In case previous iterations have freed resources, we first try to set up additional circuits on the considered virtual link. If blocking persists, we proceed to reroute the blocked *excess* traffic. For this, we exclude the considered link and compute the shortest alternative path in the virtual topology. For each link of this path, we determine the available spare capacity including additionally activated circuits if feasible and required to absorb the excess traffic. We reroute as much of the excess traffic to the alternative path as allowed by the minimum of the spare capacities, thereby activating additional circuits as required. If we cannot completely reroute the excess traffic, we exclude the limiting link and repeat the procedure from finding the alternative path. We stop with the remaining excess traffic blocked if no further alternative path is found.

In a second step, we iterate over all remaining virtual links (in decreasing order of shortest light-path length) and reroute the fraction of their traffic not filling a complete circuit if energetically advantageous. For this, we proceed similarly to above but only use an alternative path if the cost of the avoided circuit is superior to that of transit traffic processing and required additional circuits. The circuit cost respects modification: if the avoided circuit existed previously, we assign a cost of $\alpha + \delta$, otherwise $\alpha - \delta$. Likewise, a pre-existing additional circuit costs $\alpha - \delta$, a newly established one $\alpha + \delta$ (and will thus never be used). Besides, we only reroute the excess traffic as a whole, i. e. we disregard paths able to absorb parts of it. This second step corresponds to the post-processing in [7].

When rerouting excess traffic from a virtual link, we iterate over the demand routes it sustains in descending order of demand value and assign a new path to the entirety or part of the demand until a volume corresponding to the excess traffic is reached. We thereby eliminate loops from the new demand routes which occur if the rerouted demand previously traversed nodes of the alternative path. Since this may reduce the traffic load on some virtual links, we finally re-iterate over all virtual links and deactivate circuits exceeding demand.

V. Evaluation

A. Simulation Setup

We evaluated energy savings and potential QoS impacts of one-step network reconfiguration using our solution method by event-driven simulation based on the IKR SimLib [18].

1) Scenario: We present results for the Géant reference network with 22 nodes and 36 links, which is available from SNDlib [19] along with dynamic demand matrices obtained by measurement over four months [20]. Due to effort constraints, we selected a period of 14 days from this demand trace as input for the simulation studies. Following the granularity of the demand trace, we set the reconfiguration interval to $\Delta T = 15\,\text{min}$.

In order to vary the traffic load relative to the circuit capacity, 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 14 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*).

In this study, we vary the average peak demand between 0.01 and 2 circuit equivalents. Assuming a circuit capacity

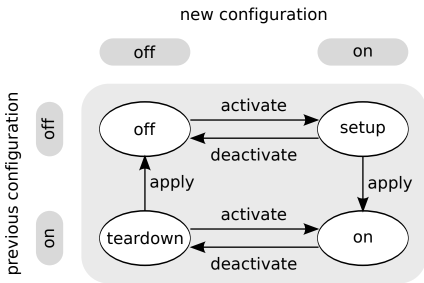

Fig. 2. Resource states and transitions

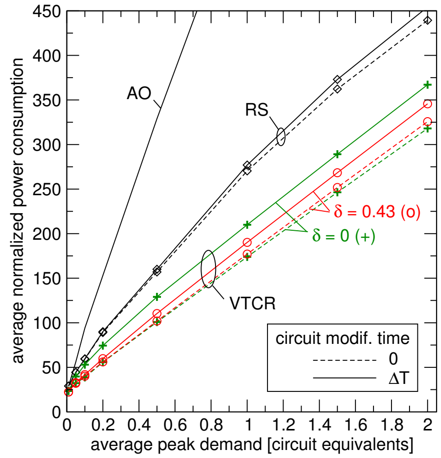

Fig. 3. Average normalized power consumption

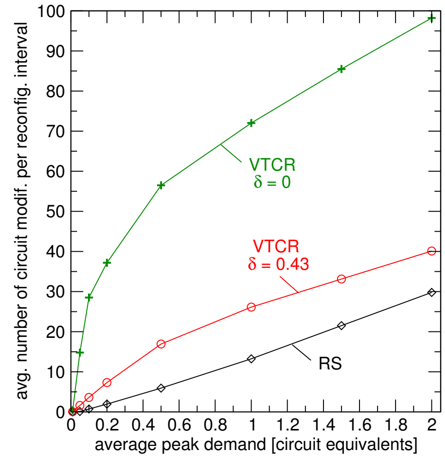

Fig. 4. Average number of circuit modifications

of 40 Gbps, this translates into a total peak demand (sum over all demands in the peak matrix) between 184 Gbps and 36.7 Tbps. The corresponding time-averaged total demand ranges between 50.1 Gbps and 10.0 Tbps.

We dimension network resources individually for each load scenario. For this, we determine an energy-optimized configuration for the peak demand matrix using our solution method with $\delta = 0$ and assume the port pairs and fibers used in this configuration to be installed. While fewer resources are likely to suffice to carry all traffic since demand peaks do generally not occur simultaneously, this dimensioning approach is commonly applied by researchers [13]. It is also in line with over-dimensioning of core networks due to traffic estimation uncertainties.

2) Reference Cases: We evaluate energy savings by one-step reconfiguration in comparison to two baseline cases: First, the static operation of all resources in the configuration determined for dimensioning. We refer to this case by always-on (AO). Second, resource scaling (RS), where we dynamically operate resources in the same static configuration: all traffic follows fixed paths in the fixed virtual topology, but we deactivate unneeded parallel circuits and we let electrical processing scale with the actual transit traffic.

3) Parameterization: We normalize power values to the consumption of one optical circuit, i.e. $\alpha = 1$. In [15], we find that the scalable power share β of processing one circuit worth of packet traffic is significantly smaller. Including shares of line card and chassis power in the circuit cost, we obtain $\beta = 4.3 \cdot 10^{-5}$. We vary the circuit modification cost $\delta \in \{0, 0.43\}$. The blocking penalties are $\mu = \nu = 17$ and $\eta = 34$, respectively. We assume a transparent optical reach of 3000 km, but allow circuits on longer physical links. The SA control parameter setting lets the optimization terminate in 2 to 4 minutes on commodity hardware.

B. Results

We evaluate the network configurations obtained by VTCR in terms of dynamic power consumption and circuit modifications per reconfiguration interval relative to RS and AO. In addition, we discuss traffic blocking.

1) Energy savings: Figure 3 plots the time average of the normalized dynamic power consumption in the network over the traffic volume. Since $\alpha \gg \beta$, this metric approximates the number of active circuits. For minimal volumes, all curves converge to the minimal number of circuits required for connectivity. The metric increases roughly linearly with the traffic volume, albeit at different slopes for the different network operation schemes.

For adaptive schemes, the plot shows two extremal assumptions for light-path setup and teardown: either circuit modification is instantaneous and we see the power consumption of the final configuration (dashed lines), or circuits under modification consume as much power as active circuits for the whole reconfiguration interval (solid lines). For VTCR, the final configurations are slightly more energy-efficient for modification cost $\delta = 0$ than for $\delta = 0.43$, but if we also consider the energy for transient circuits, $\delta = 0.43$ proves energetically advantageous. Applying a reconfiguration penalty is thus generally advisable.

Accounting for transient power consumption, VTCR (with $\delta = 0.43$) saves between 25 % and 35 % of energy compared to RS. This is comparable to the savings we obtained by a similar method with relaxed resource constraints but restricted traffic routing in [7]. Since the AO configuration may be governed by singular peak demand values, much caution is required in generalizing the significant energy savings (50 % to 70 % for VTCR) relative to it.

2) Reconfiguration: Since increasing traffic volumes bring about higher traffic variations in absolute terms, load-dependent network reconfiguration affects more circuits. This is reflected in all curves giving the average numbers of circuit modifications per reconfiguration interval over the traffic volume in Fig. 4.

A positive modification cost significantly reduces circuit reconfigurations for VTCR. The comparison of their number with results for DOB in [7] is instructive: While the additional degrees of freedom VTCR exploits in routing almost double modifications for $\delta = 0$ compared to DOB without a reconfiguration penalty, their number almost drops to the level of DOB with a penalty of 0.5 if we set $\delta = 0.43$. The reconfiguration penalty thus proves particularly effective for VTCR.

3) Traffic Blocking: We observe traffic blocking in eight of the 21,504 configurations computed for this study. Within individual 14 day periods, blocking occurs at most in two 15 minute intervals and concerns up to twice the average peak demand or 0.80 % of the average traffic load in one reconfiguration interval.

An analysis of these cases revealed two causes: In very low traffic scenarios, dimensioning essentially reduces the network topology to a ring with a few shortcuts. While achieving very limited energy savings due to connectivity requirements, dynamic adaptation in low-load situations may produce configurations no longer allowing the setup of the few additional circuits required during peak hours. VTCR is therefore not suitable for such extreme scenarios. Here, reconfiguration schemes rerouting traffic to deactivate circuits in a fixed virtual topology (like DUFL in [13]) should realize comparable energy savings without risking blocking.

The second cause is suboptimal solutions returned by SA: While available resources would allow routing the traffic on an alternative path, the required links are missing in the virtual topology. We will extend the post-processing to detect and resolve such conditions. The effect of this remedy on energy metrics should be negligible due to the scarcity of traffic blocking events.

VI. Conclusion

In this paper, we first motivated and outlined a one-step reconfiguration procedure for load-adaptive operation of core networks. We then described the related optimization problem of finding the next multi-layer network configuration under constraints on installed resources and their previous occupation. For this problem, we derive a meta-heuristics based solution procedure focusing on the upper network layer. We finally evaluate this method in terms of achieved energy savings and incurred traffic blocking by means of simulation.

Compared to load-dependent resource operation in a fixed virtual topology with fixed routing, our reconfiguration scheme reduces the power consumption related to active optical circuits by 25 % to 35 %. This is comparable to the savings we obtained without considering resource constraints while restricting traffic routing [7]. Using a modification cost significantly reduces the number of set-up and torn-down circuits and has a positive effect on energy savings due to the consumption of transient circuits.

Disadvantageous previous resource configurations resulted in traffic blocking only in extreme scenarios where scarce installed resources severely limit feasible network configurations. In these cases, reconfiguration could be limited to traffic routing and dynamic operation of pre-defined circuits without sacrificing energy savings. Other rare events of blocking revealed a weakness of our method in certain situations.

Future work will improve our heuristic to prevent spurious blocking. The resulting method should then be evaluated in a wider range of scenarios. Moreover, a systematic study could investigate the dependency of blocking on the dimensioning of network resources.

Acknowledgment

The work presented in this paper was partly funded within the SASER project SSEN by the German Bundesministerium für Bildung und Forschung under contract No. 16BP12202.

References

[1] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical IP networks," *J. Lightwave 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] A. Bononi and L. Rusch, "Doped-fiber amplifier dynamics: a system perspective," *J. Lightwave Technol.*, vol. 16, no. 5, pp. 945 – 956, May 1998.

[4] W. Golab and R. Boutaba, "Policy-driven automated reconfiguration for performance management in WDM optical networks," *IEEE Commun. Mag.*, vol. 42, no. 1, pp. 44 – 51, Jan. 2004.

[5] K. Bala, G. Ellinas, M. Post, C.-C. Shen, J. Wei, and N. Antoniades, "Towards hitless reconfiguration in WDM optical networks for ATM transport," in *Proc. Global Telecommun. Conf. (GLOBECOM)*, vol. 1, London, UK, Nov. 1996, pp. 316 – 320.

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

[7] ——, "Evaluation of centralized solution methods for the dynamic optical bypassing problem," in *Proc. 17th Conf. on Optical Netw. Design & Modeling (ONDM)*, Brest, France, Apr. 2013.

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

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

[10] S. Gieselman, N. Singhal, and B. Mukherjee, "Minimum-cost virtual-topology adaptation for optical WDM mesh networks," in *Proc. IEEE Int. Conf. on Commun. (ICC)*, vol. 3, 2005, pp. 1787–1791.

[11] 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.

[12] 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.

[13] 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.

[14] E. Bonetto, L. Chiaraviglio, F. Idzikowski, and E. L. Rouzic, "Algorithms for the multi-period power-aware logical topology design with reconfiguration costs," *J. Opt. Commun. Netw.*, vol. 5, no. 5, pp. 394–410, May 2013.

[15] U. Bauknecht and F. Feller, "Dynamic resource operation and power model for IP-over-WSON networks," in *Proc. 19th EUNICE Workshop on Advances in Commun. Netw.*, Chemnitz, Germany, Aug. 2013.

[16] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671–680, 1983.

[17] J. Sommer, E. Doumith, and Q. Duval, "On link harness optimization of embedded Ethernet networks," in *Proc. IEEE Symp. on Ind. Embedded Syst. (SIES)*, Jul. 2009.

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

[19] 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.

[20] 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.