Energy Efficient Optical Backbone Networks: A Dynamic Threshold Approach

P.Melidis
Dept. of Informatics
Aristotle University of
Thessaloniki 54124, Greece
melidis@csd.auth.gr

P.Nicopolitidis*
Dept. of Informatics
Aristotle University of
Thessaloniki 54124, Greece
petros@csd.auth.gr

G.Papadimitriou*
Dept. of Informatics
Aristotle University of
Thessaloniki 54124, Greece
gp@csd.auth.gr

E.Varvarigos
Dept. of Comp. Eng.
University of Patras
Rion 26500, Greece
manos@ceid.upatras.gr

*Senior member IEEE

***Abstract*—A new approach about virtual topology reconfiguration in IP over WDM optical networks is introduced. Traditional reconfiguration algorithms adjust the virtual topology using two thresholds as a control mechanism. When the load of a lightpath exceeds the upper threshold, a new lightpath is activated to avoid congestion and possible data loss upon the arrival of traffic bursts. Similarly, when the load of a lightpath decreases below the low threshold, a lightpath is deactivated to release the underutilized network equipment. Reconfiguration algorithms use these thresholds to adjust the virtual topology to the constantly fluctuating IP traffic flows. While they save large amounts of energy in comparison to a static virtual topology scenario, they define the upper threshold arbitrary. This is not considered efficient, since setting a very low upper threshold as a precaution is a waste of lightpath capacity. Similarly, setting a high upper threshold arbitrary, may indeed lead to data loss if a sudden data burst arrives to the network. This work proposes a new method, Virtual Topology Reconfiguration with Dynamic Thresholds (DT-VTR), which adaptively calculates the upper threshold by taking into consideration the load variance of the IP traffic flows. The upper threshold is set up based on the worst case scenario excluding any chance of data loss due to burst arrivals, while keeping the upper threshold close to 100%. Furthermore, the adaptive approach of DT-VTR differentiates from previous studies by using an individual dynamic upper threshold for each lightpath, which alternates during the network operation along with the fluctuating traffic flows. As shown by the simulation results, DT-VTR manages to save from 6% up to 18% in comparison to traditional reconfiguration algorithms for various static upper thresholds.**

***Index Terms*—Dynamic traffic, IP over WDM, optical mesh network, virtual topology reconfiguration, dynamic thresholds.**

I. INTRODUCTION

As the bandwidth demand per user is raising WDM mesh networks appear to be a solid solution as the backbone of national networks. However, as traffic demand increases, so does the energy consumption of the equipment needed to support the elevated traffic rates. Information and Communication Technology is responsible for a significant percentage of the energy consumption in global scale. For example, ICT contributes in 10% of UK's total energy consumption and that percentage is expected to grow much further until 2020 [1].

The maintenance of a static virtual topology implies waste of energy during the network operation. Virtual topology is the set of lightpaths in the network topology. On the other hand, an adaptive virtual topology can be adjusted with the main purpose of energy saving. A series of different energy aware reconfiguration algorithms have been developed that adjust the virtual topology to the dynamic IP traffic.

In [2] a virtual topology reconfiguration algorithm is presented that uses load thresholds to define overloaded and underutilized lightpaths. The algorithm attempts to keep the virtual topology in load balance by deactivating low loaded lightpaths and activating new lightpaths when traffic congestion is detected. In [3], the same algorithm is presented with the addition of a power model in order to estimate the exact power consumption of the network. By deactivating underutilized lightpaths and rerouting their traffic from alternative paths, the reconfiguration algorithm manages to save significant amounts of energy. Also, in this work the effect of the chosen threshold is studied. In [4], a routing strategy is proposed in conjunction with a reconfiguration algorithm. The proposed routing strategy assigns significant weights on a lightpath when its load is close to the load thresholds. This way it manages to save up additional energy by offloading underutilized lightpaths and attempting to keep the virtual topology in a load balanced state.

In [5] a different load balancing approach is presented that focuses on avoiding traffic disruption and minimizing the network delay. The algorithm presented in [6] also aims to diminish the traffic disruption. A genetic multi-objective algorithm is demonstrated in [7] aiming to minimize network congestion and energy consumption. The genetic algorithm uses cognitive techniques and includes fault tolerant virtual topology design with shared path protection. The algorithm in [8] also focuses in fault tolerance presenting a virtual topology adaptation algorithm including back-up paths. In [9] a virtual topology reconfiguration algorithm is proposed in combination with traffic prediction.

This paper is organized as follows: Section II briefly describes the virtual topology reconfiguration problem, the constraints and the power model used to examine the energy consumption of the network. In Section III, the key ideas about adaptively defining the upper threshold are presented along with the proposed reconfiguration algorithm. Section IV presents results that reveal the increased energy savings of our adaptive approach. Finally, Section V concludes this work.

978-1-4799-8030-7/14/$31.00 ©2014 IEEE

II. Virtual Topology Adaptation Problem

A. Network Operation

The network topology is modeled as a graph structure. The vertexes and the edges represent network nodes and bidirectional fiber links respectively. The total of network nodes and fiber links compose the physical topology. In order to establish a lightpath from a source node to destination, a route must be found on the physical topology. Then it is examined if there is at least one available wavelength in common across the path, since it is assumed that wavelength conversion is not used. If there is one available, then the lightpath can be established, else the next route is examined. Upon the creation of the lightpath, an optical transmitter is activated on the source node, an optical receiver is activated on the destination node and the wavelength is occupied across the path on each fiber link. Similarly, to deactivate a lightpath, a transmitter and a receiver on the source and destination node must be deactivated respectively, and the wavelength of the lightpath must be released across the path. The set of all available lightpaths on the network forms the virtual topology.

The traffic of the network is modeled as a set of IP-traffic flows between node pairs. Each traffic flow's load fluctuates during network operation. There are two ways of routing IP traffic flows over the virtual topology. The first one is to create a direct lightpath from the source to the destination node of the flow. These two nodes might be many hops away in the physical topology, but their distance in the virtual topology is considered to be one hop. By creating a direct lightpath, signal processing on the IP layer is conducted only on the starting and ending nodes, while the traffic flow is never processed above the optical layer in the intermediate nodes. That provides a performance advantage, since electronic signal processing is quite time consuming in comparison to optical signal processing. The second way of routing a flow is multihopping. In this case a path is created using multiple lightpaths in order to route a flow. Traffic is switched electronically not only on the ending points of the path, but also on intermediate nodes of the virtual topology. Additional electronic switching might be time consuming, but multihopping widens the possibility of reusing already existing lightpaths, which can be beneficial on an energy consumption aspect. The capacity of a lightpath is usually far greater than the load of a single flow. For that reason multiple flows must be assigned to a single lightpath. The technique of grouping up all these IP traffic flows efficiently in one lightpath is called traffic grooming.

Finally, the network operates under a series of constraints that must be fulfilled:

- The number of active transmitters/receivers on a node must never exceed the total number of available transmitters/receivers per node.
- The number of active lightpaths crossing through a fiber must never exceed the total number of available wavelengths on a link.
- The total number of incoming lightpaths on a node must be equal to the number of active optical receivers. Likewise, the total number of outgoing lightpaths on a node must be equal to the number of active optical transmitters.
- Total load of a lightpath must be equal to the load sum of all flows, which are routed through this lightpath.
- Total load of a lightpath must never exceed the lightpath capacity.

B. Energy Consumption model

The formula used to calculate the total energy consumption of the network is based on the energy model presented in [2]. Regarding the node architecture, it is assumed that each node is composed by one optical cross connect on the optical layer and an IP packet router on the upper level. Total energy consumption is calculated as:

$$POW = \sum_{\forall(m,n)\in E} a_{mn}u_{mn} + \sum_{\forall i\in N} \emptyset_i n_i + \rho \sum_{\forall(i,j)\in D} \varepsilon^e \lambda_{ij} + (1-\rho)\sum_{\forall(i,j)\in D} \varepsilon^e l_{ij} + \sum_{\forall(m,n)\in E, m\neq s,d} \varepsilon^s p_{mn} \quad (1)$$

The first term of the equation refers to idle power consumption for physical links, where a_{mn} is the energy consumption of inline amplifiers on a fiber from node m to n (9W*[$distance_{mn}$/80km+2]). u_{mn} is equal to 1 if the fiber from node m to n is used by a lightpath, else it is equal to 0. The second term refers to the idle power consumption of nodes. $\emptyset_i$ is the electronic control power consumption at node i (150W). n_i is equal to 1 if node i is used by any lightpath, else 0. The third term refers to idle power consumption per lightpath and the fourth term to traffic dependent power consumption per lightpath. λ_{ij} is equal to 1 if there is a lightpath between nodes i and j, else 0. l_{ij} is the traffic load of a lightpath between nodes i and j. ε^e is electronic processing consumption per IP port for 40Gbps (667W). The value of parameter ρ is 0.9. The fifth term refers to power consumption of intermediate nodes, where the signal is only processed on the optical layer. ε^s is energy consumption per wavelength (0.107W) and p_{mn} is the number of wavelengths on a fiber from node m to n.

C. Virtual Topology Adaptation

As it is normal in any real life network scenario, traffic load fluctuates during the network operation. In a static case the initial virtual topology is designed based on the peak traffic rates and remains constant. This way the virtual topology is able to serve every IP traffic flow even in the most loaded time of the day. However, a static virtual topology is not effective, since network equipment is underutilized most of the time resulting in energy waste. For that reason adaptive schemes have been deployed, so virtual topology can adapt to the dynamic IP traffic. The main purpose of an adaptive scheme is to maintain a constantly evolving virtual topology at all moments, which is capable of serving the dynamic traffic demands in an energy efficient way without degrading the performance of the network.

The virtual topology reconfiguration algorithm presented in [2] aims to keep the virtual topology in load balance. Load balancing is achieved by activating and deactivating lightpaths, thus the virtual topology can adjust to dynamic traffic alterations during the network operation. At the same time the load of lightpaths is bounded between acceptable thresholds, avoiding traffic congestion and underutilization of resources. Load balancing is based solely on the observed IP-traffic flows without using any prediction mechanisms. Therefore, alterations on the virtual topology may only be triggered in case there is an alteration in the traffic. Observation periods are used to define the time intervals between which traffic is observed and lightpath additions or deletions happen. The key idea is that it is possible to detect overloaded and underutilized lightpaths in the virtual topology. To that end, high and low thresholds are used, which are related to the load threshold of a lightpath expressed in percentages, above which a lightpath is considered congested and below which it is considered underutilized, respectively.

The main purpose of this algorithm is to keep the lightpath loads of the virtual topology within the threshold limits. This way the network operation is secured from data loss, because high loaded lightpaths have a considerable capacity margin to avoid overflow and also it is ensured that low load lightpaths are not being utterly underutilized, because a minimum utilization limit is defined. Whenever a lightpath is out of balance, meaning that its load has exceeded the high threshold or decreased below the low threshold, a lightpath addition or a lightpath deletion may happen respectively, so the virtual topology can be transferred in a more balanced state. At the end of each observation period only one change can happen to the network, in order to avoid unnecessary virtual topology alterations, which may lead to traffic disruption.

III. Virtual Topology Reconfiguration With Dynamic Thresholds

A. Dynamic upper threshold

In the reconfiguration algorithm presented in [2], the threshold values are chosen statically based on the objectives of the network operation. As the low threshold increases the reconfiguration algorithm becomes more aggressive, because more lightpath loads belong below the low threshold resulting in an increased number of deactivated lightpaths and consequently lower energy consumption. However, if the chosen low threshold is set very high the network operation might be disrupted from the frequent changes in the virtual topology. Similarly, total energy consumption of the network decreases as the upper threshold approaches the 100% mark. This can be attributed to the fact that as the high threshold increases there is more available capacity in lightpaths, therefore the same IP-traffic can be accommodated with fewer lightpaths. On the other hand, it is not suggested to use high thresholds near the 100% mark, because any sudden traffic increase or burst may overflow a lightpath and lead to data loss. Finally, when the two thresholds are too close to each other, the number of alterations in the virtual topology is increased leading to an unstable virtual topology.

The main objective of virtual topology adaptation is energy efficiency and in order to achieve that the reconfiguration algorithm should aim to maximize lightpath capacity and utilization. The upper threshold defines the actual capacity of a lightpath, because the space above that load limit can be considered unused capacity. As it is normal, a higher upper threshold implies that traffic load can be served with fewer lightpaths and less network equipment in comparison to a lower one. However, the static reconfiguration algorithm sets the upper threshold arbitrary, which cannot be considered efficient for two reasons. Firstly, setting a very low upper threshold as a precaution in order to avoid overflowing lightpaths and data loss is a waste of lightpath capacity. Secondly, setting a high upper threshold arbitrary, may indeed lead to data loss if a sudden data burst arrives to the network, since traffic loads or burstiness are not taken into consideration.

Input:
- L_i - i^{th} lightpath
- L_i.load - load percentage of i^{th} lightpath
- L_i.src - source node of i^{th} lightpath
- L_i.dest - destination node of i^{th} lightpath
- L_{sd} - lightpath from node s to node d
- L_{cap} - total lightpath capacity
- $flow_{sd}$ - a traffic flow from node s to node d
- $burst_{max(sd)}$ - the maximum observed load burst for $flow_{sd}$
- SLT - static low threshold
- DHT_i - dynamic high threshold of i^{th} lightpath

Output:
- VT_{new} - the new virtual topology.

Algorithm:
```
At the end of each observation period
For each L_i
  sum = 0
  For each flow_sd in L_i
     sum = sum + burst_max(sd)
  end for
  DHT_i = (L_cap – sum) / L_cap
end for
if there is flow_sd that is not routed then
   route flow_sd
   if flow_sd cannot be routed then
      create L_sd
      route flow_sd
else
   For each L_i in descending order
      if no lightpath added yet and L_i.load > DHT_i then
         For each flow_sd in L_i in descending order
            create L_sd
            if L_sd created successfully then
               break
         end for
   end for
   if no lightpath added then
      For each L_i in ascending order
         if L_i.load < SLT then
            if there is alternative path from L_i.src to L_i.dest then
               reroute all flows of L_i
               deactivate L_i
               break
      end for
end if else
return VT_new
```

Fig.1. Virtual topology reconfiguration algorithm with dynamic thresholds.

The proposed DT-VTR reconfiguration scheme presents a novel method to dynamically define the upper threshold. Every lightpath maintains an individual dynamic upper threshold that is based on the traffic flows that pass through it. The dynamic upper threshold of a lightpath alternates through time along with the respective fluctuating traffic flows. Each traffic flow is observed for a long period in order to collect information about its burstiness. In the worst case scenario, it is assumed that during the next observation period each traffic flow that passes through a lightpath will increase its load by the maximum observed burst. The total load burst for a lightpath can be calculated by aggregating the maximum observed burst of each traffic flow. The dynamic upper threshold is defined as:

$$Dynamic\ Threshold = \frac{Lcap - LBtotal}{Lcap} \qquad (2)$$

$$LBtotal = \sum_{i=1}^{n} LBmax_i \qquad (3)$$

$Lcap$ is the total lightpath capacity, $LBtotal$ is the total load burst of the lightpath and $LBmax_i$ is the maximum observed load burst of the i^{th} traffic flow routed through the corresponding lightpath. By setting the dynamic upper threshold this way, the possibility of data loss upon the arrival of traffic bursts is excluded, while keeping the upper threshold on high levels so most of the lightpath capacity can be utilized properly.

B. Algorithm description

The algorithm is illustrated in Fig. 1. Initially, it works using a static high threshold until the traffic flows are observed for a period of time capable to evaluate their burstiness. The evaluation of the traffic flows is a continuous process and carries on even after the employment of the dynamic threshold reconfiguration algorithm. As soon as the value of maximum load burst is evaluated for each traffic flow, the dynamic high thresholds are calculated. Since the traffic loads constantly alternate and also the routing in the virtual topology changes when a lightpath is activated or deactivated, the dynamic high thresholds must be recalculated at the end of each observation period. The maximum load bursts of all flows that pass through a lightpath are aggregated in order to calculate the dynamic high threshold.

After the calculation of the dynamic high thresholds, the algorithm searches for flows that are not served. In case there are flows that must be routed, the algorithm checks for paths to route through. If there are not available paths then a new lightpath is created to serve the flow. As soon as all flows are served, then the algorithm proceeds with the virtual topology load balancing. Lightpath addition is prioritized over lightpath deletion, because it is more important to accommodate increased load, rather than deactivating an underutilized lightpath. After a virtual topology alteration, all IP-traffic flows are rerouted using shortest path routing.

Whenever one or more lightpaths are detected with load above the dynamic upper threshold, then all of them are examined in a descending order based on their load until a new lightpath is added or all cases have been examined. Assuming a lightpath L_{sd} between nodes N_s and N_d has load above the dynamic upper threshold, then a new lightpath L_{new} must be established to unload some of its traffic. L_{new} is established between the nodes whose multihop traffic $flow_{max}$ contributes the most to the load of L_{sd}. In case there are not available resources for the lightpath creation the next multihop traffic flow is examined, etc. $Flow_{max}$ and probably other flows will be rerouted through L_{new}, resulting in a significant decrease of L_{sd}'s load below the dynamic high threshold. That way the network returns in a balanced state. On the other hand, whenever one or more lightpaths are detected with load below the low threshold, then all of them are examined in ascending order based on their load until an already existing lightpath is deleted or all cases have been examined. Assuming a lightpath L_{sd} between nodes N_s and N_d has load below low threshold, then it is tested if all traffic flows routed through it, can be rerouted through an alternative path. If there is an alternative path then the traffic is rerouted and the lightpath is deactivated releasing resources from N_s and N_d, otherwise the next in order underutilized lightpath is examined, etc.

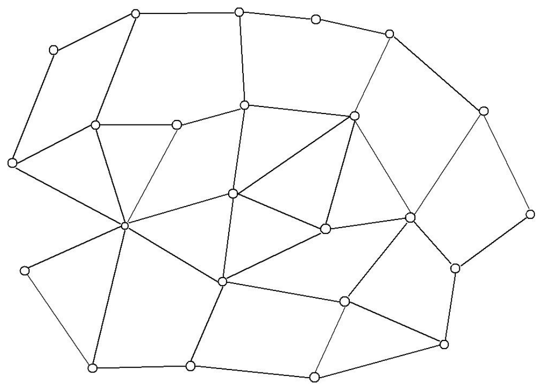

Fig.2. Network topology with 24 nodes and 42 bidirectional physical links.

IV. Comparative Performance Evaluation

Simulations were conducted in order to evaluate the performance of the proposed reconfiguration scheme. Results were compared between Energy Aware Dynamic Virtual Topology Reconfiguration (EAR) method presented in [2] which sets the upper threshold arbitrary without taking into consideration the IP traffic and Dynamic Threshold Virtual Topology Reconfiguration (DT-VTR).

A. Simulation Setup

The network topology used in the simulation consists of 24 nodes connected with 42 bidirectional physical links, as shown in Fig. 2. Each network node can deploy up to 8 transmitters and 8 receivers at a time, while each physical link can support up to 16 different wavelengths. Lightpath capacity is assumed to be the same for all lightpaths at 40 Gbps. Also, it is assumed there is no wavelength conversion in intermediate nodes. The average distance between nodes is set at 400 km based on the network topology demonstrated in [10].

TABLE I. SIMULATION PARAMETERS

Network topology	
Number of nodes	24
Number of fiber links	42
Transmitters per node	8
Receivers per node	8
Wavelengths per node	16
Wavelength conversion	none
Lightpath capacity	40 Gbps
Average node distance	400 km
Network traffic	
Non-zero load node pairs	60%
Traffic pattern values	20 - 110 Mbps
Initial scale factor	uniform [0.2, 1.2]
Adjustment scale factor	45
Simulation	
High threshold	60%, 70%, 80%
Low threshold	20%
Observation period	5 minutes
Warm up period	2 days
Simulation duration	2 days

Network traffic is modeled as in [3]. Total network traffic is the set of IP-traffic flows from a source node to a destination node. Each flow's load fluctuates during the operation time of the network and starts repeating its values every 24 hours. Load changes can happen to a flow every 5 minutes. A flow operates as one direction link, so for example a $flow_{AB}$ is different from $flow_{BA}$. Also, there is no correlation between these two flows, since one can exist without the other.

The initial step in traffic creation is to choose node pairs that have non-zero traffic between them. In these simulations 60% of the node pairs have been chosen. For each one of these node pairs a traffic pattern is chosen uniformly between a set of 15 different patterns. A traffic pattern consists of 288 load values, one for each 5 minute period of the day. Traffic patterns can be categorized as calm, normal and bursty. At this 5 minute period, the discrepancy of the traffic flow load after the final scaling can be up to 100 Mbps, 500 Mbps and 1Gbps approximately, depending on the chosen pattern category. As soon as the traffic pattern is chosen, the next step is to obtain a scale factor in order to multiply the load values of the pattern. The scale factor is chosen randomly using uniform distribution within a range of values [$factor_{low}$, $factor_{high}$] and ensures that even if the same traffic pattern is chosen between two different flows, the load values will probably still differ. A second scale factor can be used to adjust the flow load to the right proportions of the lightpath capacity. In this simulation experiment, the initial scale factor obtained values from 0.2 to 1.2, while a second scale factor with value 45 is used to adjust the traffic patterns, whose load values fluctuate between 20 and 110 Mbps. Adjustment scale factor is crucial in order to produce enough traffic load for a 40Gbps per lightpath network.

A 5 minute time interval is used as observation period, while three different values of high thresholds (80%, 70% and 60%) and a standard value for low threshold (20%) were chosen. When the simulation starts, each traffic flow is examined and a lightpath is created to accommodate it. If there are not enough resources for a new lightpath, the flow is routed through existing lightpaths. As soon as all flows are routed, virtual topology reconfiguration algorithm is used until virtual topology reaches a balanced state. Regarding EAR, the network operates for a 2-days period before capturing data, so a stable virtual topology structure is created before observing the algorithm operation. In the case of DT-VTR, network operates with a static high threshold (80%) for a 2-days period to collect enough information about the burstiness of the IP traffic flows. After that the simulation lasts for 48 hours and the results are recorded.

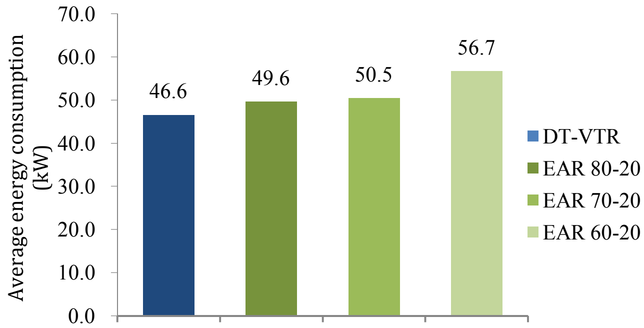

Fig.3. Comparison of average energy consumption between DT-VTR and EAR for various upper thresholds.

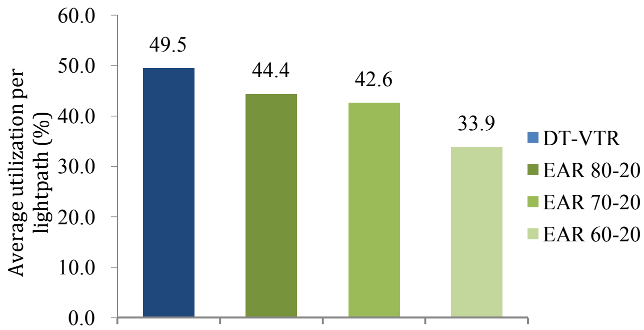

Fig.4. Comparison of average utilization per lightpath between DT-VTR and EAR for various upper thresholds.

B. Energy Consumption results

In Fig. 3, the total energy consumption of the network is shown for both compared methods. The total energy cost depends on the total number of lightpaths in the virtual topology. Consistent with the results presented in [2] and [3], higher upper threshold guarantees greater energy savings due to the fact that greater fraction of the lightpath capacity can be used to serve traffic flows. For that reason, the same traffic load can be accommodated with less lightpaths as the upper threshold increases. It can be seen from Fig. 3, that as DT-VTR has a higher average upper threshold in comparison to the arbitrary chosen high thresholds of EAR, it manages to yield greater energy savings. DT-VTR achieves 6%, 8% and 18% energy saving in comparison to EAR 80-20, EAR 70-20 and EAR 60-20 respectively. Fig. 4 explains the energy savings, as it can be observed that DT-VTR displays greater utilization per

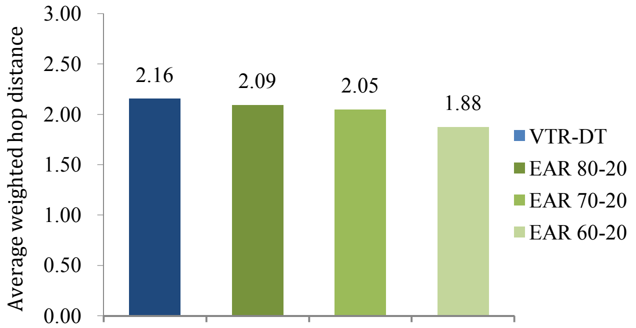

Fig.5. Comparison of average weighted hop distance between DT-VTR and EAR for various upper thresholds.

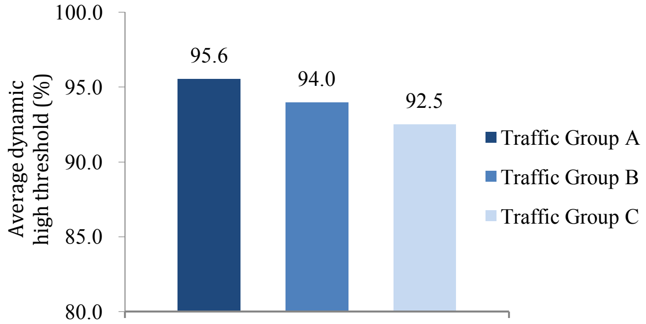

Fig.6. Comparison of average dynamic high threshold for different traffic groups.

lightpath with 49% in comparison to the non-dynamic EAR versions that achieve 44%, 42% and 34%.

C. Average Weighted Hop Distance

Average weighted hop distance is defined as the average hop distance of all flows from source to destination weighted by their load. In Fig. 5 it is shown that DT-VTR has increased hop distance in comparison to EAR. As it is normal in a virtual topology with fewer lightpaths the chance of routing flows through longer paths is increased, since fewer flows will be able to be routed directly. DT-VTR has an average weighted hop distance of 2.16, while EAR has 2.09, 2.05 and 1.88 respectively for the various upper thresholds. This slightly raised delay is a necessary trade-off in order to achieve high energy savings.

D. Dynamic High Threshold

In Fig. 6 the average dynamic high threshold of DT-VTR is illustrated for different IP traffic groups. It is demonstrated that as the burstiness of the IP traffic increases the dynamic high threshold is decreased. This behavior is normal because as the IP traffic flows become more burst intensive, the dynamic high threshold must be set lower to avoid overflowing. Three different IP traffic groups were created with different percentages of calm, normal and bursty patterns, {group A: 60%, 30%, 10%}, {group B: 50%, 30%, 20%}, and {group C: 40%, 30%, 30%}. For traffic group A, which was the mildest, dynamic upper threshold was set on average at 95.6%. For group B and C which had a greater percentage of bursty flows the threshold was set lower at 94% and 92.5% respectively. DT-VTR adapts better to variable traffic patterns while not sacrificing lightpath capacity due to transient bursts.

V. CONCLUSION

Virtual Topology Reconfiguration with Dynamic Thresholds (DT-VTR) is proposed as a new reconfiguration scheme. The main contribution of this algorithm is the introduction of a new method to dynamically calculate the high threshold. High and low thresholds are deployed as a control mechanism to activate and deactivate lightpaths. The proposed method was compared to EAR, which activates or deactivates lightpaths on fixed time intervals by observing the fluctuating dynamic IP-traffic flows. While EAR manages to save remarkable amounts of energy compared to a static virtual topology scenario, it defines the high threshold arbitrary. DT-VTR proposes a new way of calculating the upper threshold taking into consideration the load variance of the IP traffic flows. Moreover, it differentiates from EAR since it suggests using a dynamic upper threshold for each lightpath, which alternates during the network operation along with the fluctuating traffic flows. As demonstrated, DT-VTR manages to save from 6% up to 18% in comparison to EAR for various static upper thresholds. Future work will be focused on implementing traffic prediction tools, so the dynamic high threshold can be calculated based on traffic burst estimations.

ACKNOWLEDGMENT

This work has been funded by the NSRF (2007-2013) Synergasia 2011/EPAN-II Program "Energy Efficient Optical Network Planning and Operation", General Secretariat for Research and Technology, Ministry of Education, Religious Affairs, Culture and Sports (contract no. 11SYN_6_1942).

REFERENCES

[1] Global Action Plan Report, “An inefficient truth”, http://www.globalactionplan.org.uk/, 2007

[2] Yayimli, Aysegul, and Cicek Cavdar, "Energy-aware virtual topology reconfiguration under dynamic traffic", in Proceedings of the 14th International Conference on Transparent Optical Networks (ICTON), 2012.

[3] Gençata, Aysegül, and Biswanath Mukherjee, "Virtual-topology adaptation for WDM mesh networks under dynamic traffic", IEEE/ACM Transactions on Networking, vol.11, no.2, 2003, pp.236-247.

[4] Ribera, Carlos, et al, "Multi Layer Energy-Efficiency in IP over WDM Networks", in Proceedings of Asia Communications and Photonics Conference, Optical Society of America, 2012.

[5] Narula-Tam, Aradhana, and Eytan Modiano, "Dynamic load balancing for WDM-based packet networks", in Proceedings of IEEE INFOCOM 2000.

[6] Zhang, Yongbing, et al, "Traffic-based reconfiguration for logical topologies in large-scale WDM optical networks", IEEE Journal of Lightwave Technology, vol.23, no.10, 2005, pp. 2854-2867.

[7] Fernández, Natalia, et al, "Survivable and impairment-aware virtual topologies for reconfigurable optical networks: A cognitive approach", in Proceedings of the 4th International Congress on Ultra Modern Telecommunications and Control Systems (IEEE ICUMT), 2012.

[8] Bouillet, Eric, et al, "Lightpath re-optimization in mesh optical networks", IEEE/ACM Transactions on Networking, vol.13, no..2, 2005, pp. 437-447.

[9] Ohsita, Yuichi, et al, "Gradually reconfiguring virtual network topologies based on estimated traffic matrices", IEEE/ACM Transactions on Networking , vol.18, no.1, 2010, pp. 177-189.

[10] Shen, Gangxiang, and RodneyS Tucker, "Energy-minimized design for IP over WDM networks", IEEE/OSA Journal of Optical Communications and Networking, vol. 1, no.1, 2009, pp. 176-186.