

Power-Aware Logical Topology Design Heuristics in Wavelength-Routing Networks

A. Ahmad*, A. Bianco*, E. Bonetto*, D. Cuda[†], G. Gavilanes Castillo*, F. Neri*

*Dipartimento di Elettronica, Politecnico di Torino, Torino, Italy

E-mail: {arsalan.ahmad, andrea.bianco, edoardo.bonetto, guido.gavilanescastillo, fabio.neri}@polito.it

[†]IEIIT, National Research Council (CNR), Torino, Italy

E-mail: davide.cuda@polito.it

Abstract—Wavelength-Routing (WR) networks are the most common solution for core networks. With the access segment moving from copper to Passive Optical Networks (PON), core networks will become one of the major culprits of Internet power consumption. However, WR networks offer some design flexibility which can be exploited to mitigate their energy requirements. One of the main steps which has to be faced in designing WR networks is the planning of the Logical Topology (LT) starting from the matrix of traffic requests. In this paper, we propose a Mixed Integer Linear Programming (MILP) formulation to find power-wise optimal LTs. In addition, due to the complexity of the MILP approach we propose a greedy heuristic and a genetic algorithm (GA) ensuring performance close to the one achieved by the MILP formulation.

I. INTRODUCTION

Energy efficiency is one of the most challenging issues which science must face in the near future. Indeed, as the reduction of CO₂ emissions and the preservation of the environment against the climate change have started to be supported by specific policies of international organizations and governments, energy-efficiency has become a new driver for technological improvements across several research fields. In addition, as the energy demand is steadily increasing, energy costs are increasing as well and becoming a significant percentage of the OPERational EXpenditures (OPEX) of private companies. As such, companies, and in particular network operators, are paying particular attention to energy efficiency issues.

Telecommunication networks, and more in general the ICT sector, are already responsible for consuming a significant part of the world energy budget (estimated from 2% to 10%) and, if current trends are sustained, the Internet will consume 50% of the world electricity in few years, reaching 1.43 Gt of CO₂ worldwide by 2020 [1]. For instance, in Italy, the incumbent network operator Telecom Italia is the first consumer of electricity of the whole country [2]. Indeed, despite more energy efficient network devices are appearing on the market, leading to a reduction of the joule/bit ratio, the increasing amount of bandwidth ensured to an always increasing amount of users makes the joule/user ratio dramatically increasing.

In this context, backbone networks are estimated to consume today 20% of the total energy of the Internet, and their power consumption is expected to become the dominant part of the overall Internet energy requirements in a close future[4].

A more energy-conscious telecommunication network design can therefore significantly reduce global energy consumptions and costs. Optical technologies can help reducing the power requirements thanks to the low power needed to transmit data through the optical media [4]. Indeed, reduction of power consumption was a driver even in the pioneering days of optical technologies, when one of the first applications of optical fibers was in submarine systems, allowing to significantly reduce the number of regenerators required to cover the large oceanic distances. Nowadays, Wavelength Division Multiplexing (WDM) networks are the most common solution to design core networks because they offer large aggregate capacities by exploiting the optical wavelength multiplicity in the transmission of several data channels.

In this paper we consider Wavelength Routing (WR) networks, the most popular instantiation of the WDM concept. In a WR network, information is transmitted using optical circuits, called lightpaths, physically drawn over the fiber network. An important step in designing WR networks is to find a suitable Virtual Topology (VT): given a node-to-node traffic matrix and a physical topology, find which nodes should be connected directly, i.e., through lightpaths, satisfying some optimality criteria. The VT design of WR networks usually goes through two different phases: i) Logical Topology Design (LTD), i.e., the choice of the set of lightpaths and ii) Routing and Wavelength Assignment (RWA), i.e., lightpath routing over the physical topology and wavelengths assignments to each lightpath. The LTD and the RWA problems are usually solved separately because i) the complexity required to jointly solve both problems makes the solution hard and ii) these two design steps are usually performed by two different entities. Indeed, the company in charge of providing services (e.g., an ISP) usually solves the LTD according to its traffic requests, whereas the owner of the physical infrastructure (often a telecom operator) usually faces the RWA problem.

A power-aware RWA approach for WR networks was proposed in [5]. Results show that significant power savings can be obtained, even for small-size networks. In this paper we focus on the Power-Aware LTD (PA-LTD) problem, namely, defining a set of lightpaths to support a given traffic matrix taking into account power reduction as an optimization objective. The remainder of the paper is organized as follows: in Sec. II we present the LTD problem and, in particular, we focus on its

Power-Aware flavor. In Sec. III we propose a Mixed Integer Linear Programming formulation of the PA-LTD. In Sec. IV we present a greedy heuristic and a meta-heuristic based on a Genetic Algorithm (GA). Results are presented in Sec. V whereas conclusions are drawn in Sec. VI.

II. THE LOGICAL TOPOLOGY DESIGN PROBLEM

The LTD problem is agnostic of the network physical topology and it exploits only the knowledge of the traffic matrix, i.e. the amount of data that each pair of nodes in the network is willing to exchange. Indeed, the LTD output is the lightpath set which specifies a directed graph connecting nodes. Data can be sent from a source to a destination node either directly, in the optical domain, using one lightpath in a single-hop fashion, or following a multi-hop path in which more lightpaths are used. In the latter case, the traffic is switched and processed electronically between two lightpaths at each intermediate node of the multi-hop path. Electronic switching is assumed also due to buffering needs. Therefore, several traffic flows can share the same lightpath, possibly avoiding the deployment of new transmitter-receiver (TX-RX) pairs.

The LTD output comprises the set of lightpaths that satisfies the traffic requests while optimizing a given target. Classical cost functions proposed for the LTD are the minimization of i) the link congestion, ii) the end-to-end latency or iii) the CAPital EXpenditures (CAPEX). In this paper, we consider the LTD problem when power consumption minimization criteria are adopted.

A. Proposed PA-LTD

We assume that the main sources of power consumption are i) optical transceivers (TX-RX pairs) and ii) electronic switching (performed at the intermediate and endpoint nodes). Transceivers perform the electronic-optical conversion (and vice versa) and electronic switching is needed to route through traffic and to perform grooming operations. Therefore, minimizing the power consumption involves finding the best balance between usage of optical transmission and of electronic switching. Our intention is to trade off, from a power consumption perspective, the amount of electronic switching/processing with the number of optical TX-RX pairs which each node must be equipped with. Obviously, if the power consumed by optical TX-RX pairs is negligible with respect to the power required to switch data in electronics, an energy-efficient solution would lead to a fully connected topology, in which electronic switching is almost completely avoided. On the contrary, if the power consumed by the optical TX-RX pairs is considerably larger than the power needed to process data in the electronic domain, a topology with a small number of lightpaths like a star minimizes the power consumption.

III. PA-LTD FORMULATION

We model the PA-LTD problem via a MILP formulation. We consider a network with N nodes, where each node can be both source and destination of data traffic. The input parameter

is the traffic matrix $\underline{T} = [\lambda^{sd}]$, where λ^{sd} specifies the traffic requirements from source node s to destination node d ($s, d \in N$). Output variables are λ_{ij}^s , which indicate the traffic units of source s that travel on the lightpath starting from node i and ending in node j ($s, i, j \in N$). Since routing information is not an essential outcome for our purposes, we use an aggregated formulation which takes into account the overall traffic generated at each node and not each source-destination combination. Although the routing information is lost, the complexity of the MILP solution is reduced by one order of magnitude. Indeed, to solve the problem, we are interested in the number of TX-RX pairs required to transmit the traffic flowing between each pair of nodes and the aggregated amount of traffic each node is switching electronically. Another set of output variables of the PA-LTD problem is n_{ij} which represents the number of lightpaths, i.e. the number of TX-RX pairs, needed between nodes i and j ($i, j \in N$). The summation $\sum_j n_{ij}$ represents the number of transmitters node i is equipped with and, also, the total number of receivers needed in all other nodes to receive traffic generated by i .

We assume synchronous transmission, because they are the systems most commonly deployed today in core networks (e.g., SONET or X-Gigabit Ethernet). Thus, the power consumed by a TX-RX pair, namely P^{TX} , does not depend on the transmission load, but only on their nominal bitrate B^{TX} .

Regarding power consumption due to switching and grooming operations, we analyzed the data sheets of several switching fabric cards by different vendors. Considering the models presented in [4], [6], we assume that switching fabrics can be characterized by two main power consumption contributions, a fixed part needed to supply the fabric and a variable power consumption which depends on the transmitted traffic. Indeed, for simplicity, we assume that every node is equipped with the same switching fabric. Thus, the fixed contribution is equal for all nodes, and does not influence the LTD solution. The variable power consumption contribution depends on the amount of traffic switched at node i (namely λ_i) [4]. More precisely, we assume that the power consumed by the switching fabric linearly depends on λ_i and that it can be computed as $P^{SW}(\lambda_i) = \frac{\lambda_i}{\Delta}$, where $\Delta = \frac{B^{TX}}{P^{SW}(\lambda_i=B^{TX})}$ is a normalization factor measured in $bps/Watt$, that indicates the amount of traffic switched by a switching fabric consuming 1 Watt.

To assess the energy trade off between optical transmission and electronic switching, we introduce $\nu^O = \frac{P^{SW}(B^{TX})}{P^{TX}(B^{TX})}$ as the ratio between the power required to switch electronically a bandwidth equal to B^{TX} and the power consumption of a TX-RX pair (operating at bitrate B^{TX}). Thus, $\nu^O > 1$ ($\nu^O < 1$) implies that transmitting data in the optical domain is power-wise convenient (not convenient) with respect to switching the same amount of information in the electronic domain. Finally, introducing ν^O , $\Delta = \frac{B^{TX}}{\nu^O P^{TX}}$ and $P^{SW}(\lambda_i) = \nu^O \frac{\lambda_i P^{TX}}{B^{TX}}$. By varying ν^O , we can trade the power efficiency relationship between optical and electronic technologies.

In the PA-LTD formulation the objective function is the

minimization of the total power consumption

$$OF = P^O + P^E \quad (1)$$

in which the first term P^O represents the consumption of the optical TX-RX pairs, computed as $P^O = P^{TX} \sum_{i,j} n_{ij}$, and the second term $P^E = \sum_i P^{SW}(\lambda_i)$ is the total power consumption due to switching data in electronics. The traffic each node must process is equal to $\lambda_i = \sum_d \lambda^{id} + \sum_{j,s,i \neq s} \lambda_{ij}^s + \sum_s \lambda^{si}$, equivalent respectively, to the sum of the total traffic produced by the node, the traffic being forwarded by the node and the traffic received by node i . Furthermore, the following constraints must be satisfied to return feasible LTs.

$$\sum_i (\lambda_{ij}^s - \lambda_{ji}^s) = \begin{cases} -\sum_d \lambda^{sd}, & j = s, \forall (j, s) \\ \lambda^{sj}, & j \neq s, \forall (j, s) \end{cases} \quad (2)$$

Eq. (2) expresses the flow conservation constraints which guarantees that traffic addressed to each node j is dropped while all the remaining traffic leaves node j .

$$\sum_s \lambda_{ij}^s \leq n_{ij} B^{TX}, \forall (i, j) \quad (3)$$

Eq. (3) ensures that the number of TX-RX pairs between nodes i and j satisfies the amount of traffic flowing between each pair of nodes.

$$\sum_{j,s} \lambda_{ij}^s \leq B^{SW}, \forall i \quad (4)$$

The maximum aggregate bandwidth that a node can electronically switch is limited to B^{SW} in Eq. (4).

$$\sum_j n_{ij} \leq \delta^{TX}, \forall i; \quad \sum_i n_{ij} \leq \delta^{RX}, \forall j \quad (5)$$

Eqs. (5) fix to δ^{TX} (δ^{RX}) the maximum number of transmitters (receivers) that each node can use.

A. Reasons for heuristic approach

The previously formalized PA-LTD may be solved using any optimization software available on the market. These tools usually can find solutions with a certain degree of optimality, but the inherent complexity of the problem makes the time required to compute the solution impractical. In PA-LTD, the complexity scales rapidly with the network size because the required number of variables scales as $O(N^3)$. Indeed, the optimizer (AMPL+CPLEX software in our case) obtains solutions with 99% of optimality for a network with 24 nodes after running for more than 24 hours on a state of the art computer. Hence, as the number of nodes increases to several tens, it becomes hard to obtain a good solution in a reasonable time. Since real networks comprise from several tens up to hundreds of nodes, it becomes impossible to use any optimizer to solve the LTD problem on standard hardware.

Therefore, heuristic approaches are mandatory. We developed an iterative greedy heuristic and a meta-heuristic based on a Genetic Algorithm, which are able to find good solutions in a limited amount of time, as shown later.

IV. HEURISTIC TECHNIQUES

We set the power consumption of the network as target, but the proposed algorithms can be easily adapted to other minimization problems, as, for instance, CAPEX minimization. All heuristics use a routing algorithm based on the Dijkstra algorithm, and the routing information is at the base of the heuristic calculation. Since the routing is applied incrementally (that is, traffic requests are routed one at a time using the current status of the network) the routing algorithm has to keep into account the current lightpaths traffic load to calculate further routes. To support multi-path routing without further increasing complexity, each traffic demand is divided in $h+1$ different traffic requests such that $\lambda^{sd} = h \times B^{TX} + x$, where $h = \lfloor \lambda^{sd} / B^{TX} \rfloor$ and $x = \lambda^{sd} \bmod B^{TX}$. As such, we are able to mimic the MILP formulation which supports multi-path routing, i.e, traffic λ^{sd} can be divided over several paths, subject to the constraints of Eqs 2-4.

A. Less Energy Incremental Heuristic

We first propose an iterative greedy heuristic, aiming at minimizing the LT power consumption by performing local optimal choices. The heuristic proposed is named Less Energy Incremental heuristic (LE-I); Algorithm 1 describes how lightpaths between nodes are added into the LT and how flows are allocated on lightpaths. Starting from a void LT, each traffic demand is satisfied by choosing the less power consuming alternative between adding a new direct link or by using an already available path, provided that there exists a path with enough bandwidth.

Note that the order in which traffic demands are considered may affect the performance of the heuristic. Indeed, by allocating first larger traffic requests, smaller traffic requests may later use some spare capacity, while, when working in ascending order, it may be more difficult to allocate larger traffic requests on lightpaths which are partially used.

In this paper, we consider three different criteria to order requests, according to their traffic load in bps: increasing, decreasing and random order.

B. Meta-heuristic Techniques

Meta-heuristics are computational methods that implement a repeated search through a large set of possible solutions to optimize a target function over the set. They make no assumptions on the problem to be solved, but make use of random mechanisms to avoid getting stuck in sub-optimal points or regions of the test set. Two examples of these techniques are the GAs (that are inspired by natural evolution mechanisms) and simulating annealing (inspired by a metallurgical technique used to reduce material defects). We propose to exploit a meta-heuristic based on a GA.

1) *Brief Recap of Genetic Algorithms:* GAs adopt dynamics typical of the natural evolution to converge to better solutions. GAs are based on a population of individuals, each one representing a solution to the problem, that evolves through generations and in which only the best individuals survive, i.e., pass to the next generations. At algorithm start-up, an initial

Algorithm 1 LE-I heuristic

```
1: start from a void LT with  $N$  nodes
2: for all traffic demands  $\lambda^{sd}$  do
3:   calculate all the possible paths between  $s$  and  $d$ 
4:   if not  $\exists$  a path with the required capacity between  $s$ 
   and  $d$  then
5:     add a direct link between  $s$  and  $d$ 
6:   else
7:     compute the additional power consumed by the short-
     est path carrying  $\lambda^{sd}$  ( $p_{\text{path}}^{sd}$ )
8:     if  $p_{\text{path}}^{sd} \leq P^{TX}$  then
9:       use the path to carry  $\lambda^{sd}$ 
10:    else
11:      add a direct lightpath between  $s$  and  $d$ 
12:    end if
13:  end if
14:  update lightpath traffic loads with  $\lambda_{sd}$ 
15:  update total power consumption of the network
16: end for
```

population is created, randomly generating new individuals. Since it may happen that some individuals do not represent a feasible solution to the problem, before adding a new individual to the population, it is necessary to check for individual feasibility. If the new individual is feasible it is added to the population, otherwise it is dropped. The process iterates until the complete initial population is obtained. At each generation, a certain number of new individuals are obtained from the reproduction of individuals belonging to the current population. The reproduction consists in the following steps:

- *Selection*: selection of two individuals (the parents) from the previous generation (a certain priority criterion applies);
- *Crossover*: the parents' descriptive parameters are split into two sections of random size; then a new individual is obtained using a different section of each parent;
- *Mutation*: each element of the child solution array is mutated (changed) with a given probability;
- *Acceptance*: the feasibility of the individual is checked; if the resulting individual is unfeasible, discard it; otherwise add the new individual to the population.

The reproduction phase repeats a given number of times, and at the end of the process, the population comprises old individuals and their offspring. Since the population size has to be stable, only a part of the population can be selected for the next generation. Thus, the GA defines a *fitness function* that identifies the more suitable individuals that should be selected to survive. The GA usually ends under certain convergence conditions: *i*) a maximum number of generations has been reached, *ii*) the best suited individual does not show any improvement for more than a certain number of generations. The first case provides a precise ending condition; the second one guarantees that a good individual is obtained, considering

a priori that no further improvement is possible.

2) *Genetic Algorithm for PA-LTD*: In the proposed GA for PA-LTD, each element of the population represents a possible LT of our networks. Thus, each element is represented through a set of lightpaths corresponding to a feasible solution (if the individual does not represent a feasible solution it is immediately discarded). We implement each individual through an array of size N^2 in which element $l = \{0, 1, \dots, N^2 - 1\}$ of the array represents the number of TX-RX pairs used between node i and node j , with $(i = \lfloor N^2/l \rfloor)$ and $j = l \bmod N$. A LT is considered feasible only if it satisfies all the traffic requests, i.e., all the traffic request can be routed over it. We consider as the fitness function the overall power consumption of the network as defined in Eq.(1): the lower the power consumption the more fit the individuals are. Therefore, on each generation, only the less power-consuming LTs survive. According to the classical methodology, we consider a constant population size of 30 LTs, reproducing an offspring of 20 LTs on each generation. In the mutation process each individual randomly changes one element of the array, i.e., after reproduction, each element of each new generated individual is mutated with probability $1/N^2$. Finally, the algorithm stops after N^2 iterations if no power improvement is obtained, this limit being heuristically set.

V. PERFORMANCE RESULTS

The performance of the heuristics is evaluated comparing their results with optimal results obtained using an optimization software which in our case is AMPL+CPLEX. The heuristics were implemented in the C programming language using an ad-hoc library for networking applications. After having analyzed several data sheets, we assumed to use transmitters (and consequently receivers) with capacity $B^{TX} = 10$ Gbps and a nominal power consumption equal to $P^{TX} = 8$ Watt [7]. We consider several optimization scenarios in which the parameter ν^O assumes values in the range from 1 to 20, to test how the relative energy efficiency between optical and electronic technologies can affect the final LT. Values smaller than 1 or larger than 20 would not provide new insights, as shown by the graphs. The comparison is performed under uniform traffic in a network of $N = 16$ nodes and under an unbalanced traffic matrix derived from traffic on the US backbone (retrieved from [8]). In the greedy heuristic, we use the suffix -rand to indicate random ordering, -asc for ascending and -desc for descending order in processing traffic requests.

For the uniform traffic case, we defined two scenarios: *low* traffic and *high* traffic. In the first scenario, all traffic demands are equal to $\lambda^{sd} = 0.6$ Gbps ($\forall s, d$), while in the second one they are set to $\lambda^{sd} = 5$ Gbps ($\forall s, d$). The first scenario represents the case in which all traffic flows of a node fit in a single lightpath, while the latter scenario is representative of the node to node traffic comparable to B^{TX} . All ordering criteria degenerate in this uniform traffic matrix scenario.

Fig. 1 shows the average number of transmitters per node both in the *low* and *high* uniform traffic cases. Different algorithms are identified by different marker shapes. In addition,

white (black) filled markers identify the *low* (*high*) traffic.

Considering the *low* uniform traffic, the optimal LT is a star for $\nu^O < 16$, while it becomes a full mesh for larger ν^O . The LE-I heuristic results are optimal for any ν^O , while GA only approximates the star topology and converges to the optimal solution only when the optimal LT becomes a full mesh. In the *high* traffic the optimal solution is a mesh topology (becoming a full mesh for $\nu^O \geq 2$) and both algorithms catch this transition ensuring a topology close to the optimal one. Small differences are present for low values of ν^O , while the same topologies are found as ν^O increases. Note that in the *low* uniform traffic scenario, there is just one transition from a low-connected to a highly-connected LT. Indeed, there exists a threshold equal to $\bar{\nu}^O = B^{TX}/\lambda^{sd}$, under which it is convenient to perform electronic switching (thus the star LT is the optimal solution), while optical transmission becomes more energy efficient, above $\bar{\nu}^O$, hence leading to a full mesh LT.

The average number of transmitters per node of the 24 nodes US network is depicted in Fig. 2. Also in this scenario, as ν^O increases the LT becomes more and more connected because optical transmission becomes more energy efficient with respect to electronic switching. Regarding heuristics, performance are closer to the optimal one as ν^O increases. For low ν^O , heuristics present irregular behaviors, i.e. they do not show a monotonic trend, due to the fact that they get stuck in local minima. Note that almost independently of the traffic scenarios (more scenarios have been tested and results are not reported here due to space limits), the LE-I-desc tends to achieve a lower average number of transmitters than the other traffic ordering. Indeed, processing the larger requests first through a single hop path permits to reuse some spare capacity for the smaller traffic requests which will be routed later.

Fig 3 shows the average number of hops per traffic request as a function of ν^O . Recall that as ν^O increases, the LT becomes more connected. Hence, the average number of hops steadily decreases. The average number of hops for CPLEX is not available because in the MILP formulation the routing information is not taken into account for simplicity.

In addition, Fig. 4 shows the relative power consumption normalized to the power of the optimal LT, i.e., $\Delta P_{\%} = (OF^x - OF)/OF$, where OF^x represents the total power consumption for the heuristics x ($x \in \{GA, LE-I\text{-rand}, LE-I\text{-asc}, LE-I\text{-desc}\}$) evaluated as the summation of the power consumption due to electronic routing and optical transmission for the heuristic x . OF is evaluated according to Eq. (1). GA usually ensures lower power consumption than the LE-I heuristic, independently of the ordering used to accommodate traffic requests. Regarding the three different ordering methods of LE-I, no ordering presents an obvious advantage from the power consumption perspective, because the best performance ordering changes depending on the value of ν^O . However, differences are mostly negligible, and they decrease as ν^O increases.

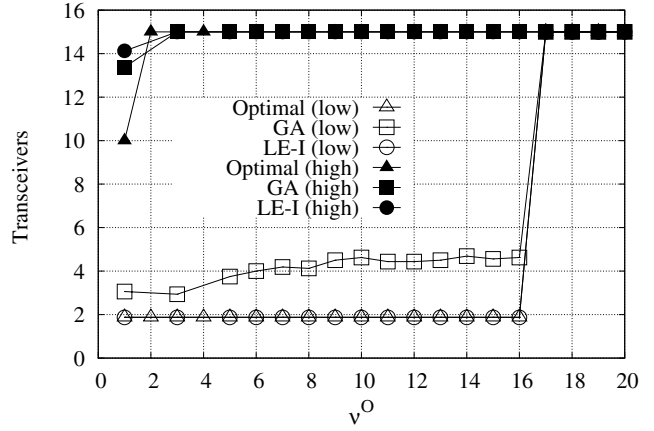


Fig. 1. Average number of transmitters per node for uniform traffic cases in the 16-nodes network.

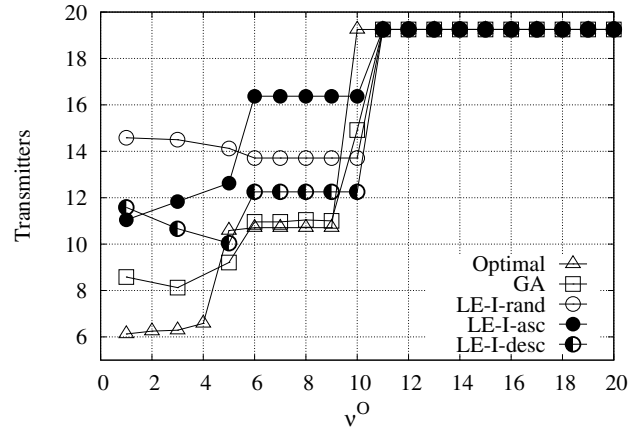


Fig. 2. Average number of transmitters per node in the US backbone 24-nodes network.

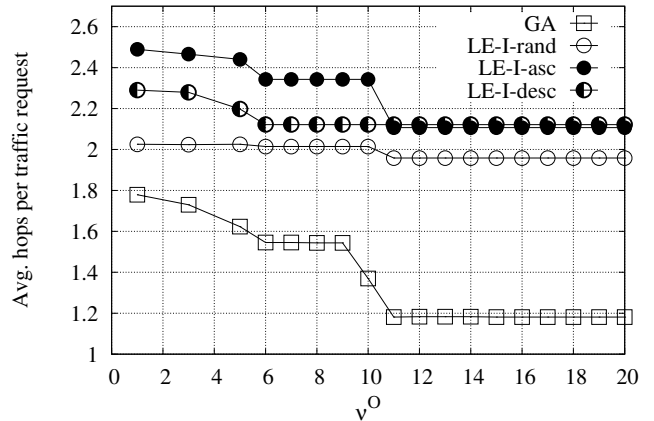


Fig. 3. Average number of hops per traffic request in the US backbone 24-nodes network.

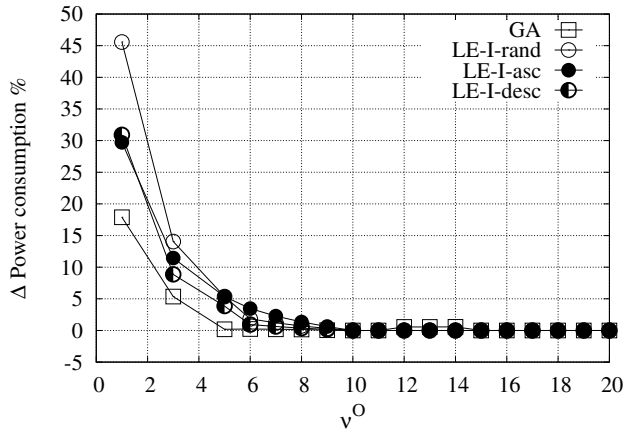


Fig. 4. Δ power consumption percentage in the US backbone 24-nodes network.

VI. CONCLUSIONS

We explored some alternatives in the study of the power awareness in the logical topology design of WR networks. We defined an optimal model and efficient heuristic algorithms.

Among the examined techniques, optimal solutions solved through commercial optimizers permit to infer the behavior of LTs as the power cost between optics and electronics change. Significant differences in the LT design are highlighted that permit to exploit power benefits of the different technologies. Furthermore, the optimal solutions were taken as a benchmark for the performance of a heuristic and a meta-heuristic algorithms. We found that meta-heuristic approaches based on GA can lead to sub-optimal solutions, although close to optimal LTs. Besides complexity, an additional advantage of the heuristic approach is the possibility to identify the routing information or other metrics (such as average hop count per traffic request) to improve configuration algorithms. Finally, thanks to the heuristic techniques, it is possible to solve the LTD problem even for large scale networks, when a general purpose optimization software cannot be used because of the large execution times.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7 (2007-2013) under grant agreement 247674 (STRONGEST).

REFERENCES

- [1] SMART 2020 Report, "Enabling the low carbon economy in the information age," <http://www.theclimategroup.org>, 2008.
- [2] G.Griffa, "NGN2 and energy OPEX savings," Giornata di studio su Efficienza Energetica, <http://www.telecomitalia.com> Torino, Italy, Nov. 2008.
- [3] J.Baliga, R.Ayre, K.Hinton, W.V.Sorin, R.S.Tucker, "Energy Consumption in Optical IP Networks," *Journal of Lightwave Technology*, v.27, n.13, pp.2391-2403.
- [4] E.Bonetto, L.Chiaraviglio, D.Cuda, G.Gavilanes Castillo, F.Neri, "Optical Technologies Can Improve the Energy Efficiency of Networks," ECOC, Vienna, Sep. 2009.
- [5] L.Chiaraviglio, M.Mellia, F.Neri, Y.Wu, "Power-Aware Routing and Wavelength Assignment in Optical Networks," ECOC, Vienna, Sep. 2009.
- [6] J. Cardona Restrepo, C.Gruber, C.Mas Machuca, "Energy Profile Aware Routing," ICC, Dresden, Germany, June 2009.
- [7] Fiber Optics Inventory Store, <http://www.lightwavestore.com>
- [8] A.Nag, M.Tornatore, "Optical network design with mixed line rates", *Optical Switching and Networking*, v.6, n.4, pp.227-234, 2009.