Optical Switching and Networking 8 (2011) 249–258

Contents lists available at SciVerse ScienceDirect

Optical Switching and Networking

journal homepage: www.elsevier.com/locate/osn

Sparse power equalization placement for limiting jamming attack propagation in transparent optical networks

Amornrat Jirattigalachote [a,*], Nina Skorin-Kapov [b], Marija Furdek [b], Jiajia Chen [a], Paolo Monti [a], Lena Wosinska [a]

[a] Royal Institute of Technology KTH, School of Information and Communication Technology, Isafjordsgatan 22, Electrum 229, 164 40 Kista, Sweden
[b] Department of Telecommunications, Faculty of Electrical Engineering and Computing, University of Zagreb, Unska 3, 10000 Zagreb, Croatia

ARTICLE INFO

Article history:
Available online 6 July 2011

Keywords:
Physical-layer attacks
Power equalizers
Greedy randomized adaptive search procedure (GRASP)
Transparent optical networks

ABSTRACT

The latest advances in Wavelength Division Multiplexing (WDM) technology are making it possible to build all-optical transparent WDM networks, which are expected to be able to satisfy the rapid growth of today's capacity demand. However, the transparency of such networks makes them highly vulnerable to deliberate attacks, specifically targeting the physical layer. Physical-layer attacks, such as high-power jamming, can cause severe service disruption or even service denial, enhanced by their capability to propagate through a transparent optical network. Several attack-aware routing and wavelength assignment algorithms have been proposed to reduce the possible disruption caused by high-power jamming attacks. However, even with network planning approaches which take network security, specifically physical-layer attacks, into account, resilience to deliberate attacks in such scenarios remains an issue.

In this paper, we propose the use of wavelength-selective attenuators as power equalizers inside network nodes to limit the propagation of high-power jamming attacks. Due to the increased cost of optical switching nodes associated with the addition of power equalizers, we aim at minimizing their number through sparse power equalization placement. We developed a set of greedy algorithms to solve what we call the Power Equalization Placement (PEP) problem with the objective of minimizing the number of power equalizers needed to reduce, to a desired level, the propagation of high-power jamming attacks for a given routing scheme. We further improved upon these results by proposing a GRASP (Greedy Randomized Adaptive Search Procedure) heuristic with a somewhat longer execution time, but with significantly superior results. The performance evaluation results indicate that the proposed GRASP heuristic can achieve the same attack propagation reduction as can be obtained by equipping all nodes with power equalizers by placing them at less than 50% of the nodes on average, potentially yielding significant cost savings.

1. Introduction

Transparent optical networks (TONs) are considered to be the most promising solution for satisfying the rapid growth of bandwidth demand in next generation networks and the future Internet. In such networks, the signal is transported from source to destination entirely in the optical domain via all-optical channels called *lightpaths*, and each lightpath can be modulated at very high data rates of up to 100 Gbit/s [1]. The set of lightpaths forms a virtual topology over the physical optical network. These lightpaths do not require any intermediate opto-electronic processing, thus reducing the number of costly optical-to-electrical-to-optical (O/E/O) converters and the operational costs. This type of transmission provides transparency to signal bit rate, format and protocol.

* Corresponding author. Tel.: +46 87904083; fax: +46 87904090.
E-mail addresses: amornrat@kth.se (A. Jirattigalachote), nina.skorin-kapov@fer.hr (N. Skorin-Kapov), marija.furdek@fer.hr (M. Furdek), jiajiac@kth.se (J. Chen), pmonti@kth.se (P. Monti), wosinska@kth.se (L. Wosinska).

1573-4277/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.osn.2011.06.008

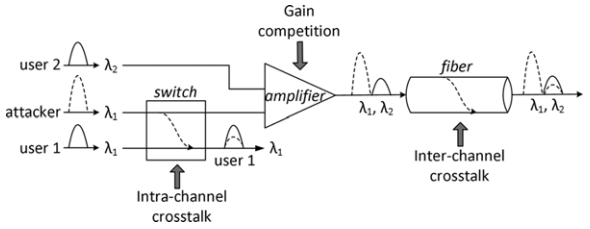

Fig. 1. An example of the potential consequences of a high-power jamming attack.

However, due to the lack of signal regeneration functionality, transparency makes TONs more vulnerable to physical-layer attacks. Namely, intermediate nodes in TONs do not interpret or regenerate the signals they carry making it possible for a malicious attacking signal to propagate uncontrollably. In this way, deliberate physical-layer attacks can spread and propagate unhindered through various parts of the network, causing damage to a large set of connections and making attack localization and source identification more difficult. In contrast, component malfunctions, for which most protection and restoration mechanisms are developed, affect only the connections directly passing through them. Various physical-layer attack scenarios are described in [2,3]. One of the most typical attacking scenarios is realized by injecting a high-power jamming signal within the passband of optical components, e.g., optical amplifiers, switches, and fibers. Such a signal, injected on a legitimate lightpath, can lead to the so-called gain competition in erbium-doped fiber amplifiers (EDFAs). In such a situation, weaker legitimate signals on different wavelengths traversing the same amplifier are "robbed" of gain, while the attacking signal receives additional amplification. The same jamming signal can also enhance crosstalk effects among lightpaths at different wavelengths inside optical fibers, giving rise to inter-channel crosstalk. Furthermore, a powerful jamming signal can interfere with other channels routed on the same wavelength via intra-channel crosstalk effects inside their common optical switches. Some consequences of jamming attacks are shown in Fig. 1.

Recently, the concept of attack-aware network planning as a prevention mechanism has been proposed, together with Routing and Wavelength Assignment (RWA) algorithms aimed at minimizing the potential damage caused by physical-layer attacks [2,4]. RWA is the process of finding physical routes and assigning wavelengths to all lightpaths in the virtual topology subject to a set of constraints. This problem in general is NP-complete [5] and many heuristic approaches considering various objectives and constraints have been proposed. The work in [2] presents a new objective for the RWA problem, called the maximum Lightpath Attack Radius (*maxLAR*) which is minimized and whose aim is to reduce the disruption caused by high-power jamming attacks. The *maxLAR* is defined as the maximum number of *unique* lightpaths any one lightpath shares a common directed physical link with and it is equal to the maximum number of lightpaths any one jamming signal injected on any lightpath in the network can attack via gain competition in optical amplifiers and inter-channel crosstalk in fibers. The authors propose an integer linear program (ILP) formulation for the routing sub-problem with the objective to minimize the *maxLAR*. Due to the complexity of the ILP formulation, the work in [2] also proposes a tabu-search heuristic algorithm to solve the attack-aware lightpath routing sub-problem for larger networks. In [4], the authors propose bin-packing-based heuristics for the wavelength assignment sub-problem aimed at minimizing the maximum potential damage caused by jamming attacks exploiting intra-channel crosstalk in switches.

The propagation of high-power jamming attacks to all lightpaths downstream of the point of attack can be stopped by installing power equalizers at all network nodes. In such a scenario, a jamming attack would be terminated at the first node downstream of the point of attack, limiting the damage to only those lightpaths traversing the same link as the attacking signal itself. However, due to the high cost of such equipment, the number of nodes equipped with power equalizers should be minimized and their respective placement in the network should aim at reducing the propagation of potential attacks.

This paper, as an extension of our work in [6], presents a new attack-aware planning approach aimed at reducing the propagation of high-power jamming attacks using power equalizers based on wavelength-selective attenuators [7] placed at a subset of the network nodes. The objective is to minimize the number of power equalizers subject to a desired level of attack vulnerability. We propose fast greedy algorithms, along with a Greedy Randomized Adaptive Search Procedure (GRASP), to perform power equalization placement for a given routing scheme.

The rest of this paper is organized as follows. Section 2 gives an overview of physical-layer attacks propagation, and presents the applied node architecture with integrated power equalizers. Section 3 presents a metric for measuring the propagation of a high-power jamming attack and an illustrative example of how this new metric is computed. In Section 4, the power equalization placement problem is defined and the proposed heuristics to limit jamming attack propagation making use of the described node architecture are presented. The numerical results are analyzed in Section 5 and, finally, Section 6 provides concluding remarks.

2. A node architecture for limiting the propagation of high-power jamming attacks

The propagation of high-power jamming attacks can be efficiently limited by installing power equalizers at network nodes. To obtain power equalization, wavelength-selective attenuators [7] can be used. A wavelength-selective attenuator [7] consists of a set of variable optical attenuators (VOAs), an array of photodetectors, and a control unit, which is able to provide the dynamic control of signal powers at any level. A jamming signal, once injected in the network where nodes are not equipped with power equalizers, can disrupt not only the lightpath it is injected on, but potentially all other lightpaths sharing common links downstream of the attacking point, as a result of gain competition in amplifiers and inter-channel crosstalk in fibers. The true damage extent

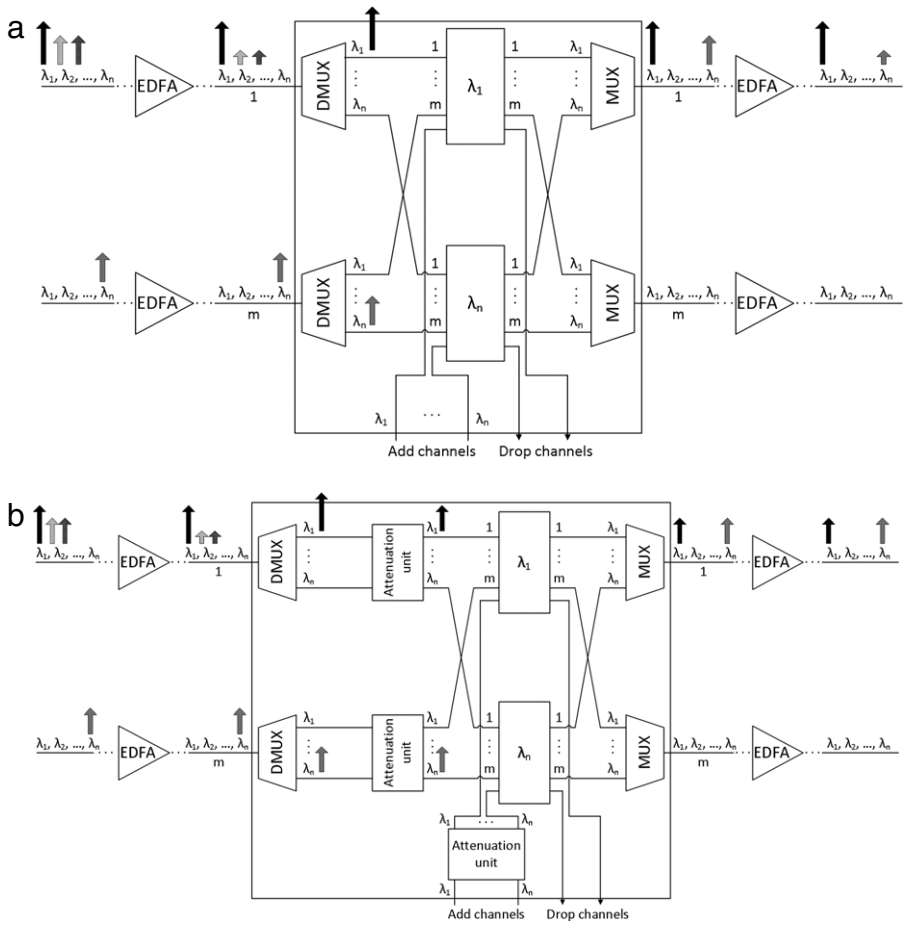

Fig. 2. An example of the propagation of a high-power jamming attack using a node architecture (a) without and (b) with power equalization.

of such attacks depends on the quality and on the speed of network detection, localization and protection mechanisms. However, due to the high data rates of TONs, even the instantaneous damage of the "bits in flight" affected before protection mechanisms can react can cause significant service degradation. In networks where nodes are equipped with power equalizers, the high-power jamming signal will be attenuated to an acceptable level after passing through a power equalizing node, without causing further disruption to other lightpaths traversing common physical links downstream of that node. In that case, the maximum number of lightpaths that could be attacked would be equal to the maximum number of lightpaths traversing any one physical link, i.e., the lightpath congestion. Another option to limit the propagation of high-power jamming attacks would be to upgrade/replace EDFAs along the link with Optical Limiting Amplifiers (OLAs) [8,9], which provide a constant output power for a dynamic range of input power. However, since a jamming attack is assumed to be able to be injected anywhere along a fiber link, it might happen that the power difference between a jamming attack and the other lightpaths on the fiber link exceeds the dynamic range of the OLAs. For this reason, while in most cases to limit the propagation of a jamming attack it is enough to replace the first and last EDFAs on a fiber link with OLAs, dealing with jamming signals inserted on particularly long fiber links would require some extra replacement of EDFAs with OLAs, resulting in a higher investment cost.

An example of high-power jamming attack propagation in networks with and without nodes equipped with power equalizers is shown in Fig. 2. Let us assume that a high-power jamming signal is injected on lightpath 1 (wavelength λ_1) on fiber link 1. As a result of gain competition in optical amplifiers and inter-channel crosstalk in fibers, other lightpaths sharing the same fiber link with the attacked lightpath could also be affected, i.e., their gain in the common EDFAs before entering the optical switch would be decreased (Fig. 2(a)). Lightpath n (wavelength λ_n) on fiber link m is safe up to that point. However lightpath n from input fiber m, after going through the same network node as lightpath 1 from input fiber 1, can also be degraded via gain competition at the EDFAs at their common output fiber, i.e., fiber 1. Furthermore, lightpath 1 receives additional amplification at the EDFA on the input side of the switch, which increases intra-channel crosstalk inside the switch.

Fig. 2(b) shows an alternative case where network nodes are equipped with power equalizers based on wavelength-selective attenuators. In this case, the power of attacking lightpath 1 on fiber link 1 is attenuated to an acceptable level after passing through the node. Thus,

lightpath n on fiber link m is not disrupted at their common output fiber. Furthermore, using attenuators decreases the intra-channel crosstalk inside the switching fabric.

Network nodes equipped with VOAs placed after the wavelength division multiplexers and de-multiplexers, which can be used for power equalization, are currently available on the market [10]. In the past, the most commonly used optical node architecture comprised Fixed Optical Add-Drop-Multiplexers (FOADMs) where signal power settings are determined in the system commissioning phase and cannot be changed afterward. Due to the fact that FOADMs do not have the capability of dynamically managing power level fluctuations of incoming signals, these nodes offer no protection against high-power attacks. FOADMs are still used in approximately 80% of nodes in currently deployed networks, but are gradually being replaced by Reconfigurable OADMs (ROADMs) and Colorless/Directionless (C/D) ROADMs which are usually equipped with VOAs. These reconfigurable nodes can be dynamically tuned via the network management system and are, thus, able to react to attack occurrences according to a pre-defined attack protection scenario. However, VOAs [11,12] are still quite expensive and (C/D) ROADMs are not yet widely deployed. Current market trends show a tendency of ROADM and C/D ROADM usage increasing to 50% of network nodes, while the remaining nodes will still consist of FOADMs [13]. Therefore, careful placement of power equalizing capabilities at a subset of nodes in the network (either via VOA-equipped ROADMs, C/D ROADMs or the proposed node architecture) becomes an important network planning problem which can help limit the attack propagation and in this way increase network security. Moreover, assuming (C/D) ROADMs are available in some network nodes, it would also be possible to address the problem of attack-awareness in a dynamic lightpath provisioning scenario.

3. Calculation of the maximum propagation of high-power jamming attacks

In order to formally describe the power equalization placement problem, a metric for measuring the propagation of a jamming attack needs to be properly defined. This section first introduces the concept of the lightpath attack radius from [2], modified to account for the presence of power equalizers in the switching nodes. An illustrative example of how this new metric is computed is then presented.

3.1. The maximum Lightpath Attack Radius (maxLAR)

The work in [2] assumes the worst-case scenario in which a high-power jamming signal can be injected anywhere along a legitimate lightpath, after which it can propagate unhindered along the rest of the lightpaths attacking all connections sharing common links with it. Based on these assumptions, the maximum Lightpath Attack Radius (*maxLAR*) is defined as the maximum number of lightpaths that any one lightpath is link-sharing with, representing the maximum number of lightpaths that could be disrupted.

Recall that with power equalizers installed at the network nodes, a jamming attack could still affect all the lightpaths with which it shares physical links with, but only until the first power equalizer is reached. After this point, the attack propagation is thwarted. To model this behavior, we divide each lightpath into sub-lightpaths, assuming division points at all nodes equipped with power equalizers which the lightpath traverses. Consequently, a lightpath traversing n intermediate nodes equipped with power equalizers is divided into $n+1$ sub-lightpaths, each one with mutually independent attack propagation properties. Therefore, the definition of the *maxLAR* from [2] is modified here as the maximum number of lightpaths that any one *sub*-lightpath shares a common fiber link with. Note that, despite the new definition, the network lightpath congestion, i.e., the maximum number of lightpaths routed over any one physical link in the network, still represents a lower bound on the *maxLAR* as defined here.

3.2. Illustrative examples of maxLAR calculation

Fig. 3 shows an example of two different schemes for power equalization placement in a six-node network with four lightpaths. We denote the load of a node as the sum of all transit and outgoing lightpaths which traverse and originate at that particular node (Fig. 3(a)). Incoming lightpaths are not considered in the load since a power equalizing node can only reduce the attacking capabilities of lightpaths which propagate beyond the node. With no power equalization, the *maxLAR* of the routing scheme in Fig. 3(a) would be equal to 4. Namely, a jamming attack injected at the beginning of, e.g., lightpath LP_1, could potentially disrupt lightpaths LP_2, LP_3, and LP_4, as well as the lightpath it was injected on (i.e., LP_1). Lightpath congestion in the example is equal to 3.

In order to limit attack propagation, it seems reasonable to place power equalizers at the most loaded nodes. Namely, placing power equalizers at nodes 1, 3, and 5 in this example would not affect jamming attack propagation. Here, there are two most loaded nodes, i.e., nodes 2 and 4. Let node 2 be chosen first (Fig. 3(b)), dividing lightpaths LP_1 and LP_2 into two sub-lightpaths. Consequently, if a jamming attack is injected on sub-lightpath $LP_{1,1}$, besides itself, it could only disrupt lightpath LP_2, i.e., sub-lightpath $LP_{2,1}$. However, if sub-lightpath $LP_{1,2}$ is attacked, lightpaths LP_2 (i.e., sub-lightpath $LP_{2,2}$), LP_3 and LP_4 could all be disrupted. Despite the power equalizer placed at node 2, the *maxLAR* is, in this case, still equal to 4. In order to reduce the *maxLAR* value to its minimum, i.e., network congestion, additional power equalizers need to be placed in the network. Suppose we now choose node 4, the most loaded of the remaining nodes. After placing a power equalizer at node 4 (Fig. 3(b)), lightpath LP_1 is divided into 3 sub-lightpaths, while lightpaths LP_2, LP_3, and LP_4 are divided into 2 sub-lightpaths each. If a jamming signal is inserted on sub-lightpath $LP_{1,2}$, only lightpaths LP_2 (sub-lightpath $LP_{2,2}$) and LP_3 (sub-lightpath $LP_{3,1}$) can be disrupted, in addition to LP_1. Thus, by placing power equalizers at nodes 2 and 4, we have successfully reduced

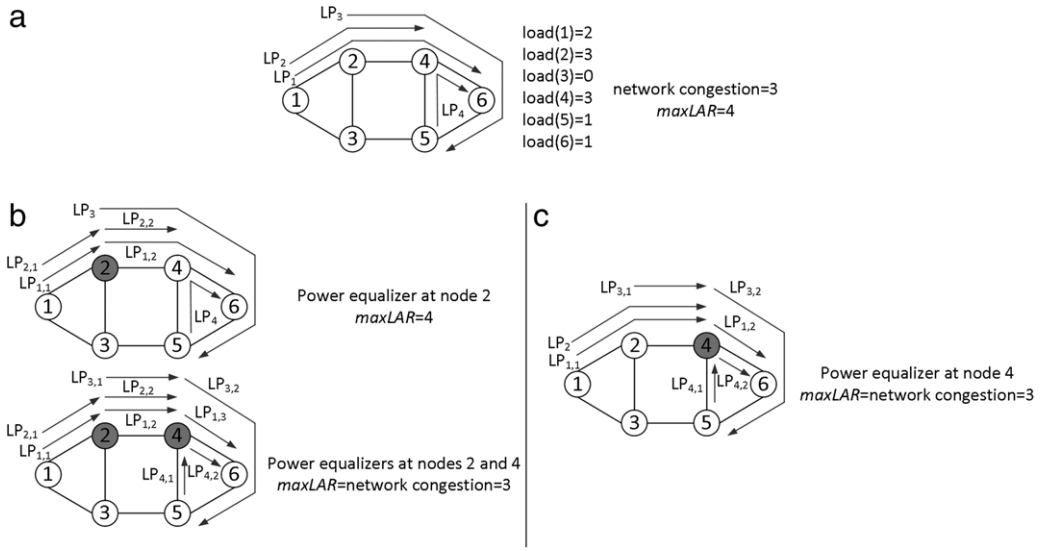

Fig. 3. An example of power equalization placement and the respective *maxLAR* calculation.

the *maxLAR* to its lowest possible value equal to the network congestion.

Alternatively, suppose a new placement of power equalizers, starting this time from node 4 (Fig. 3(c)). In this case, a jamming signal injected on any lightpath or sub-lightpath could at most disrupt two other lightpaths/sub-lightpaths, giving a *maxLAR* already equal to the network congestion. This clearly illustrates how the order in which nodes are chosen for power equalization placement can affect the resulting network *maxLAR* value, as well as the total number of nodes equipped with power equalizers necessary to lower the *maxLAR* to the network congestion.

4. Power equalization placement

This section first introduces the problem of power equalization placement to limit the propagation of high-power jamming attacks, and then presents a series of greedy algorithms, as well as a Greedy Randomized Adaptive Search Procedure (GRASP), to minimize the number of power equalizers needed.

4.1. Problem definition

Given is a physical topology of a network and a routing scheme. The physical topology is represented by a graph $G = (V, E)$, where V is the set of network nodes and E is the set of network edges. It is assumed that every network edge consists of a set of bidirectional fibers, i.e., one fiber per direction. The routing scheme consists of a set of physical paths, each path corresponding to one lightpath demand. The objective of the power equalization placement problem is to place a minimal number of power equalizers at a subset of the network nodes in order to reduce the maximal potential attack propagation, i.e., the *maxLAR* described in Section 3.1, for a given routing scheme to a pre-defined value.

4.2. Greedy algorithms for power equalization placement problem (Greedy–PEP)

We developed a series of greedy algorithms for the PEP problem, with variations on the greedy function. The greedy functions we tested include iteratively choosing: (i) the most loaded node as defined in Section 3.2, (ii) the source node of the link traversed by the highest number of lightpaths, (iii) the node of the lightpaths or sub-lightpaths with attacking capabilities equal to the *maxLAR* whose equipping with power equalization achieves the minimal average *LAR*, and (iv) the most loaded node of lightpath or sub-lightpath with attacking capabilities equal to the *maxLAR*. In this work we present only the best performing algorithm, i.e., the one using greedy function (i), denoted as the Greedy Power Equalization Placement (Greedy–PEP) heuristic. As input, Greedy–PEP takes the physical topology $G = (V, E)$, a set of physical paths corresponding to the set of lightpath demands, the *maxLAR* of the given routing scheme and the desired *target_maxLAR*. The algorithm sorts the nodes in V in descending order of their respective loads and then places a power equalizer at each node, starting from the most loaded one, until the desired *target_maxLAR* is reached. If there are multiple nodes with the same load, one is chosen at random. The pseudo-code of Greedy–PEP follows.

Due to the anomaly which can occur using Greedy–PEP as described in Section 3.2, i.e., the order in which the nodes with the same load are chosen can yield significant differences and some maximally loaded nodes (e.g., node 2 in the example) do not affect the values of the *maxLAR*, a post-processing step to remove superfluous power equalizers could be executed. However, the search would still yield a single local, not necessarily global, optimum. Consequently, we have developed a Greedy Randomized Adaptive Search Procedure (GRASP) for the power equalization placement problem, called GRASP–PEP, to improve upon the results by investigating a series of local optima in an effort to obtain superior sub-optimal solutions.

Begin Greedy–PEP
//Initialization
Input $G = (V, E)$, *set of physical paths*, *maxLAR*, *target_maxLAR*;
$S = \emptyset$; *//Set of nodes equipped with power equalizers*
current_maxLAR = *maxLAR*;
for all $i \in V$ **do**
    load(i) = number of outgoing lightpaths from i + number of transit lightpaths at i;
end for
while *current_maxLAR* > *target_maxLAR* **do**
    Sort nodes $i \in V \backslash S$ in the descending order of their respective *load*(i);
    maxLoad = max{*load*(i)};
    most_loaded_nodes = {Nodes i | *load*(i) = *maxLoad*};
    if |*most_loaded_nodes*| > 1 **then**
        Node j = random (*most_loaded_nodes*);
    else
        Node j = *most_loaded_nodes*(1);
    end if
    $S = S \cup \{j\}$;
    current_maxLAR = *maxLAR* if power equalizers are placed at nodes in S;
end while
End Greedy–PEP

4.3. The Greedy Randomized Adaptive Search Procedure for the power equalization placement problem (GRASP–PEP)

In this section, the general concept of Greedy Randomized Adaptive Search Procedure (GRASP) is explained, followed by the description of the proposed GRASP–PEP algorithm.

4.3.1. Greedy Randomized Adaptive Search Procedure (GRASP)

The Greedy Randomized Adaptive Search Procedure (GRASP) is a multi-start metaheuristic used for combinatorial optimization problems. Each GRASP iteration consists of two main steps: the construction phase and the local search phase. In the construction phase, a feasible solution is built iteratively using a randomized greedy algorithm. This algorithm first creates a list of candidate elements which are not yet included in the solution, but whose contribution to the partial solution, evaluated against the greedy function, is in an upper percentage (according to a candidate list parameter). This list is called the Restricted Candidate List (RCL), and its size can be determined by the number of elements considered or by their respective qualities. In each iteration of the greedy algorithm, an element is randomly selected from the RCL and added to the partial solution. After a feasible solution is built, the construction phase terminates. This solution is not necessarily locally optimal so a local search phase is applied. In the local search phase, the neighborhood of the initial solution is iteratively searched to find the best, or first, improving neighboring solution, which replaces the current solution. The neighborhood of a solution is defined as all those solutions which can be obtained by applying an elementary transformation to the current one. After running a desired number of GRASP iterations, or a desired number of iterations without improvement of the incumbent solution, the best found solution over all iterations is deemed the final result. A detailed explanation of GRASP can be found in [14].

4.3.2. The proposed GRASP–PEP algorithm

In the initialization phase of our GRASP algorithm for the power equalization placement problem (GRASP–PEP) we take as input the physical topology $G = (V, E)$, a set of physical paths corresponding to each set of the lightpath demands, the initial *maxLAR* of the routing scheme (without power equalizers), the desired *target_maxLAR*, the maximum allowed number of GRASP–PEP iterations, the maximum allowed number of iterations without any improvement of the incumbent solution, and the maximum size of the RCL (denoted as *RCL_size*).

The construction phase of GRASP–PEP builds a feasible solution in the following way. First, the load of each node in V not yet equipped with a power equalizer is calculated and the nodes are sorted in the descending order of their corresponding loads. Next, a fixed-sized RCL is built. It consists of (*RCL_size*−1) of the most loaded nodes, along with one randomly selected node from those not yet included in RCL for diversification purposes. Power equalizers are then placed subsequently at randomly selected nodes from the RCL. Each time a new power equalizer is placed, the new *maxLAR* of the routing solution is calculated. For as long as the new *maxLAR* is greater than the desired *target_maxLAR*, the construction phase continues. Finally, when the *target_maxLAR* value is reached, the corresponding solution (i.e., the subset of nodes equipped with power equalizers) is returned to the main GRASP–PEP function.

Even though this solution achieves the desired *target_maxLAR*, not all the assigned power equalizers are necessarily needed. This effect was illustrated in the example in Section 3.2. Consequently, the solution is delivered to the local search phase which attempts to decrease the number of power equalizers used, but without increasing the *maxLAR*, i.e., removing those power equalizers which do not affect the overall *maxLAR* value. A neighboring solution with respect to the current one is defined as a power equalization placement scheme in which a power equalizer is removed from one and only one node in the current solution. The local search phase iteratively updates the current solution with a better neighboring solution, i.e., a solution which uses one less power equalizer while maintaining the *target_maxLAR*, for as long as a better solution is found. If there are multiple improving neighboring solutions of the same quality, one is chosen at random to become the new current solution in next iteration of the local search.

The incumbent solution of GRASP–PEP is updated after each GRASP iteration with a better solution, if found. The pseudo-code of the GRASP–PEP algorithm follows.

5. Numerical results

The performance of the proposed algorithms for power equalization placement, i.e., Greedy–PEP and GRASP–PEP, was evaluated through simulations using the 30-node European COST 266 reference network [15] shown in Fig. 4. Both algorithms were implemented in Matlab, and run on a PC powered by Intel core i5 at 3.2 GHz and 3.5 GB of

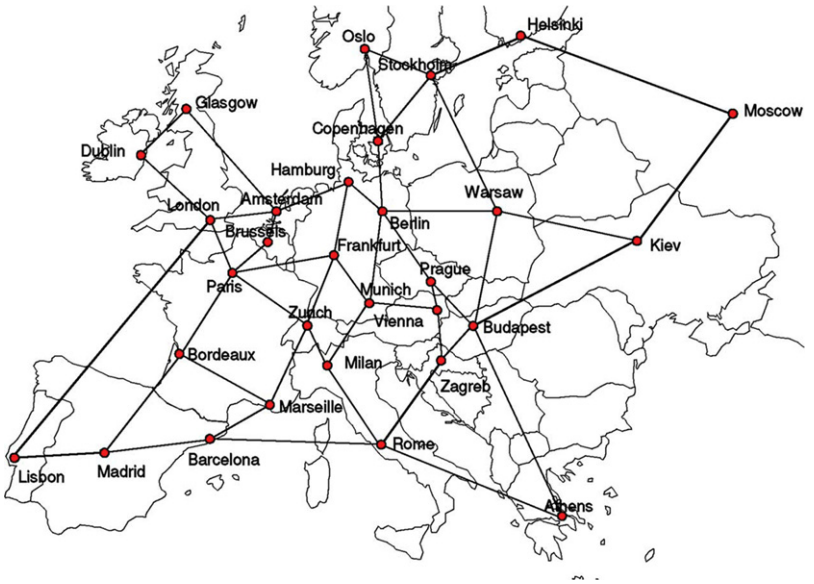

Fig. 4. The 30-node European COST 266 reference network.

Begin GRASP–PEP
//Initialization
Input $G = (V, E)$, *set of physical paths*, *maxLAR*, *target_maxLAR*;
Input *GraspIteration*; *//The maximum allowed number of GRASP–PEP iterations*
Input *Iter_No_Improve*; *//The maximum allowed number of iterations without improvement*
Input *RCL_size*; *//Size of RCL*
$S = \emptyset$; *//Solution of the current iteration*
$BestS = \emptyset$; *//Best solution found*
best_number_of_PEs $= |V|$; *//Number of power equalizers needed; initially all nodes*
$i = 1$;
$j = 0$;
while $i \leq$ *GraspIteration* **do**
  $S =$ *GreedyRandomizedConstruction* (*RCL_size,maxLAR,target_maxLAR*);
  $S =$ *LocalSearch*(*S,target_maxLAR*);
  if $|S| <$ *best_number_of_PEs* **then**
    $BestS = S$;
    best_number_of_PEs $= |S|$;
    $j = 0$;
  else
    $j = j + 1$;
  end if
  if $j =$ *Iter_No_Improve* **then**
    break;
  end if
  $i = i + 1$;
end while
End GRASP–PEP

Begin GreedyRandomizedConstruction
  (*RCL_size,maxLAR,target_maxLAR*)
$N = \emptyset$; *//Temporary set of nodes with power equalizers*
current_maxLAR $=$ *maxLAR*;
for all $i \in V$ **do**
  load(i) = number of outgoing lightpaths from i + number of transit lightpaths at i;
end for
while *current_maxLAR* $>$ *target_maxLAR* **do**
  Sort nodes $i \in V \backslash N$ in the descending order of their respective *load*(i);
  most_loaded_RCLnodes $= \{|RCL\_size - 1|$ of the most loaded nodes $i\}$;
  random_RCLnode = {random($i \in V \backslash (N \cup most\_loaded\_RCLnodes)$)};
  RCL $=$ *most_loaded_RCLnodes* $\cup$ *random_RCLnode*;
  Node j = random(*RCL*);
  $N = N \cup \{j\}$;
  current_maxLAR $=$ *maxLAR* if nodes in N equipped with power equalizers;
end while
$S = N$;
return S;
End GreedyRandomizedConstruction

RAM. Twelve different sets of lightpath demands, i.e., virtual topologies, were created as follows. First, 12 traffic matrices were generated using the method from [16], where a fraction F of the traffic is uniformly distributed over the range $[0, (C/a)]$, and the remaining traffic is uniformly distributed over range $[0, (C^*\gamma/a)]$. The values were set to $C = 1250$, $a = 20$, $\gamma = 10$ and $F = 0.7$, as in [16]. To obtain a virtual topology from each traffic matrix, lightpath requests were assigned to node pairs in decreasing order of their corresponding traffic, with at most one lightpath between each pair of nodes and 10 transmitters and receivers available per node. This gave on average 298 lightpaths per virtual topology. To find physical routes for each set of lightpath demands, three types of routing schemes were considered: (i) shortest path routing, (ii) an attack-aware tabu-search heuristic routing

Begin LocalSearch(*S,target_maxLAR*)
$l = 1$;
$S_{temp} = S$;
while $l = 1$ **do**
  $S_{it} = \emptyset$; *//Set of solutions*
  for $k = 1, \ldots, |S_{temp}|$ **do**
    $S_{neighbor} = S_{temp} \backslash S_{temp}(k)$;
    Find *maxLAR* if nodes in $S_{neighbor}$ equipped with power equalizers;
    if $maxLAR_{Sneighbor} \leq target\_maxLAR$ **then**
      $S_{it} = S_{it} \cup S_{neighbor}$;
    end if
  end for
  if $S_{it} = \emptyset$ **then**
    $l = 0$;
  else
    Choose randomly among solutions in S_{it} and set one of them as S_{temp};
  end if
end while
$S = S_{temp}$;
return S;
End LocalSearch.

algorithm, called *TS_LAR*, aimed at minimizing the *maxLAR* from [2], and (iii) *k*-shortest path routing [17] for a total of 36 different routing sets. In the *k*-shortest path routing approach, *k* was set to 3, i.e., the physical path of each lightpath request was randomly selected amongst its 3 shortest paths. The characteristics of the obtained test sets are shown in Table 1, i.e., their initial *maxLAR* and congestion values. Parameter tuning for GRASP–PEP was done experimentally giving the best results for the following parameters: *GraspIteration* = 1000, *Iter_No_Improve* = 150, and *RCL_size* = 10.

Fig. 5 shows the number of power equalizers in the solutions obtained by the proposed algorithms needed to achieve a desired *target_maxLAR* equal to its lower bound, i.e., the value of the network congestion. As stated before, placing power equalizers at all nodes in order to reduce the *maxLAR* to the congestion can be expensive. Our approaches achieve the same minimal *maxLAR* value (i.e., congestion) by using a significantly smaller percentage of power equalizing nodes. Namely, the Greedy–PEP algorithm obtained such solutions by placing power equalizers on average at only 10.5 (35%), 17.08 (56.9%), and 18 (60%) of the network nodes for the shortest path, *TS_LAR*, and *k*-shortest path routing schemes, respectively. The GRASP–PEP algorithm performed even better, yielding solutions with on average 7.16 (23.87%), 13.41 (44.7%) and 14.5 (48.34%) power equalizing nodes for the shortest path, *TS_LAR*, and *k*-shortest path routing schemes, respectively. In other words, GRASP–PEP obtained a *maxLAR* equal to congestion with on average less than 50% of power equalizing network nodes over all routing schemes.

Note that the congestion, i.e., the used *target_maxLAR*, depends on the routing scheme. The congestion of shortest path routing is the highest of the three cases (Table 1) requiring the fewest power equalizers to achieve it, but consequently limiting the *maxLAR* to the highest value. The routing obtained by *TS_LAR*, which *a priori* reduces

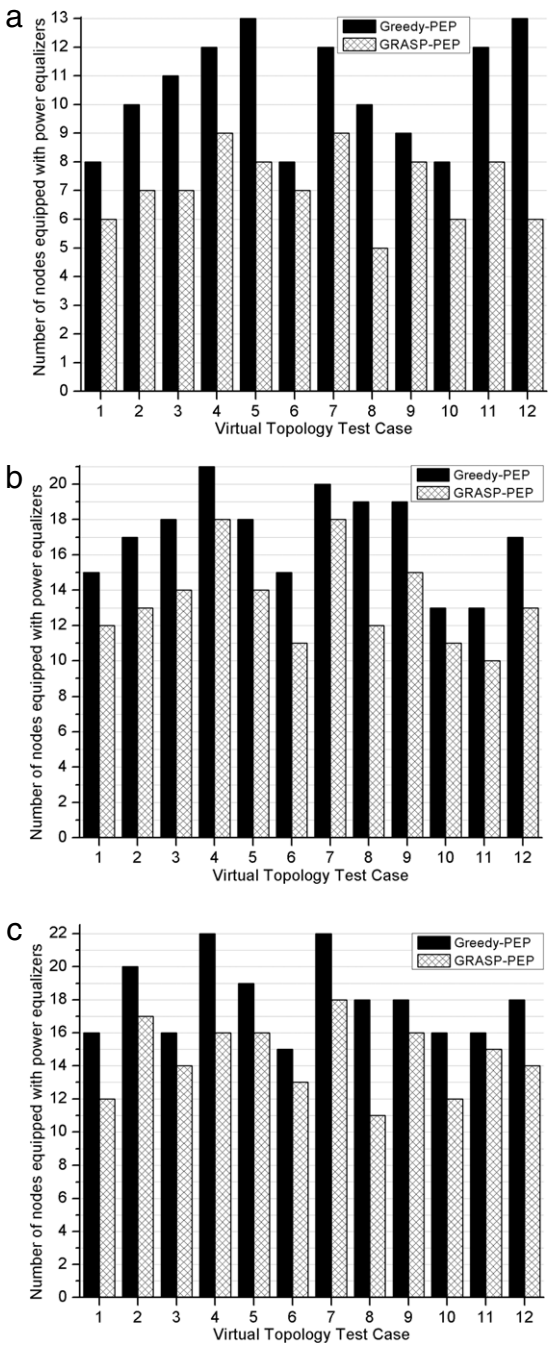

Fig. 5. The number of power equalizers in the solutions obtained by Greedy–PEP and GRASP–PEP with a *target_maxLAR* = network congestion for the (a) shortest path routing, (b) TS_LAR [2], and (c) *k*-shortest path routing schemes corresponding to the tested virtual topologies.

the *maxLAR*, has significantly lower congestion values than the *k*-shortest path schemes, making the power equalization problem more challenging. Nonetheless, both Greedy–PEP and GRASP–PEP reach the lower *TS_LAR* congestion value using fewer power equalizers for the *TS_LAR* routing scheme than for the *k*-shortest paths. This indicates that using attack-aware routing (e.g., *TS_LAR*)

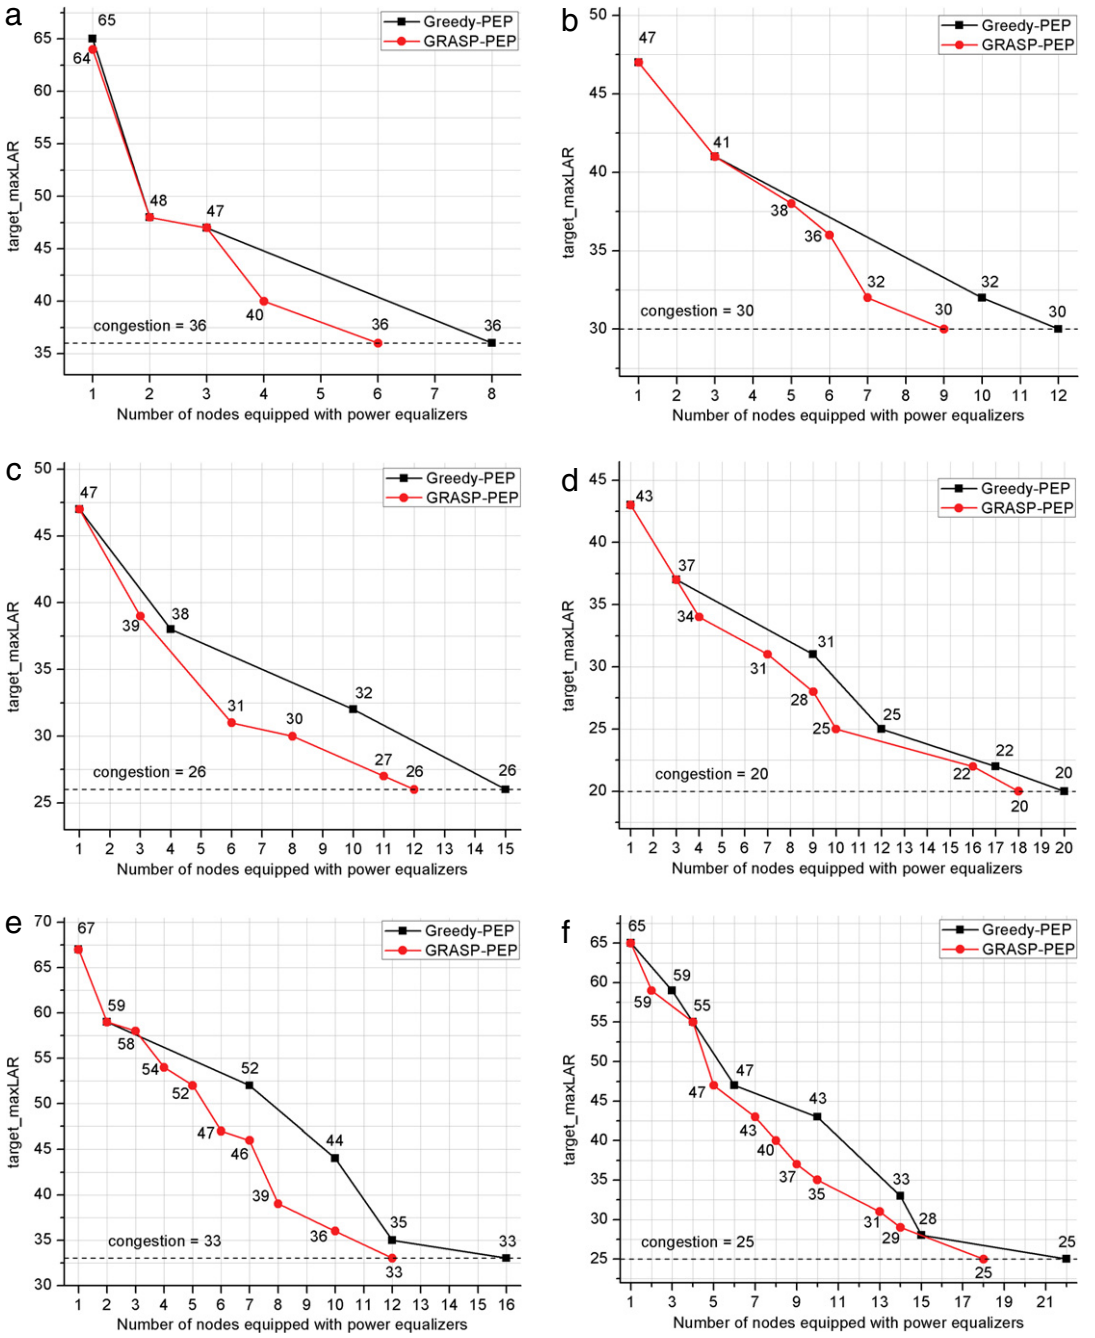

Fig. 6. The number of power equalizers needed by Greedy–PEP and GRASP–PEP to achieve various values of the *target_maxLAR* for [(a), (b)] shortest path routing, [(c), (d)] *TS_LAR* routing, and [(e), (f)] *k*-shortest path routing for virtual topology test cases [(1), (7)], respectively.

in conjunction with a PEP algorithm may be a favorable planning approach.

When comparing the performance of the proposed algorithms, we can see from Fig. 5 that GRASP–PEP outperforms Greedy–PEP in all cases, as expected. On average, GRASP–PEP uses 30.5%, 21.6%, and 19.2% less power equalizers than Greedy–PEP for shortest path, *TS_LAR* and *k*-shortest path routing, respectively. However, the average execution time of Greedy–PEP (1.6 s) is significantly lower than that of GRASP–PEP (1004.8 s). Considering that the PEP problem is an offline planning problem, the increase in complexity of the GRASP in comparison to the Greedy approach is not critical and can be used to solve realistic problem instances.

For more insight into the behavior of the algorithms, we tested both GRASP–PEP and Greedy–PEP for various levels of the desired *maxLAR* (i.e., *target_maxLAR*) for all the considered routing schemes. Fig. 6 presents the results for virtual topology test cases 1 and 7 for all three routing schemes. The results for the remaining test

Table 1
Initial *maxLAR* and network congestion of the considered routing schemes.

Virtual topology test case	Shortest path		TS_LAR		*k*-shortest path	
	maxLAR	Congestion	*maxLAR*	Congestion	*maxLAR*	Congestion
1	68	36	51	26	73	33
2	75	35	52	24	83	28
3	75	33	55	23	76	30
4	75	32	55	23	86	27
5	63	32	45	23	76	29
6	64	34	54	27	85	33
7	63	30	49	20	80	25
8	71	38	56	28	83	36
9	63	28	51	22	79	27
10	64	36	54	25	86	36
11	61	32	56	27	75	30
12	70	34	53	24	79	30

cases are analogous and omitted for lack of space. As before, GRASP–PEP obtained the same or better solutions in all cases, its superiority increasing as the *target_maxLAR* decreases and the problem becomes more challenging.

6. Conclusion

In this paper, we propose the use of wavelength-selective attenuators as power equalizers at optical network nodes in order to limit the propagation of high-power jamming attacks in transparent WDM networks. Due to the high cost of such equipment, it is desirable to place them at only a subset of nodes in the network. A series of fast greedy approaches and a Greedy Randomized Adaptive Search Procedure (GRASP) are proposed to minimize the number of power equalizers required to thwart jamming attack propagation for any given routing scheme to a desired level. The simulation results indicate that the proposed approaches can obtain placements which, in most cases, equip less than 50% of the nodes with power equalizers, while achieving the same attack propagation characteristics of the case where all nodes are equipped with such functionality. The GRASP heuristic outperforms the greedy approach in all cases, at a small trade-off with increased execution time. As future work, we plan to investigate the benefits of power equalization placement in conjunction with the reconfigurable capabilities of network nodes equipped with power equalizers in order to limit attack propagation for partially dynamic lightpath provisioning. We also plan to consider additional physical-layer attack scenarios in the context of attack-aware optical networks planning.

Acknowledgments

The work presented in this paper was carried out with the support of the project "A Security Planning Framework for Optical Networks (SAFE)", funded by the Unity Through Knowledge Fund (UKF) in Croatia, 036-0362027-1641, funded by the Ministry of Science, Education and Sports, Croatia, the project "Security in All-Optical Networks" funded by VINNOVA (The Swedish Governmental Agency for Innovation Systems), and by BONE ("Building the Future Optical Network in Europe"), a Network of Excellence funded by the European Commission through the 7th ICT-Framework Programme.

References

[1] J. Berthold, A.A.M. Saleh, L. Blair, J.M. Simmons, Optical networking: past, present, and future, IEEE/OSA Journal of Lightwave Technology 26 (9) (2008) 1104–1118.
[2] N. Skorin-Kapov, J. Chen, L. Wosinska, A new approach to optical networks security: attack-aware routing and wavelength assignment, IEEE/ACM Transactions on Networking 18 (3) (2010) 750–760.
[3] C. Mas, I. Tomkos, O.K. Tonguz, Failure location algorithm for transparent optical networks, IEEE Journal on Selected Areas in Communications 23 (8) (2005) 1508–1519.
[4] M. Furdek, N. Skorin-Kapov, M. Grbac, Attack-aware wavelength assignment for localization of in-band crosstalk attack propagation, Journal of Optical Communications and Networking 2 (11) (2010) 1000–1009.
[5] I. Chlamtac, A. Ganz, G. Karmi, Lightpath communications: an approach to high-bandwidth optical WANs, IEEE Transactions on Communications 40 (7) (1992) 1171–1182.
[6] A. Jirattigalachote, N. Skorin-Kapov, M. Furdek, J. Chen, P. Monti, L. Wosinska, Limiting physical-layer attack propagation with power equalization placement in transparent WDM networks, in: Proc. of Asia Communications and Photonics Conference and Exhibition, ACP, December 2010.
[7] A. Mocki, T.M. do Amaral, A.A.P. Pohl, 8-channel dynamic optical power equalizer for WDM networks, in: Proc. IEEE International Telecommunications Symposium, September 2006.
[8] I.W. Way, D. Chen, M.A. Saifi, M.J. Andrejco, A. Yi-Yan, A. von Lehman, C. Lin, High gain limiting erbium-doped fibre amplifier with over 30 dB dynamic range, IEEE Electronics Letters 27 (3) (1991) 211–213.
[9] O.C. Graydon, M.N. Zervas, R.I. Laming, Erbium-doped-fiber optical limiting amplifiers, Journal of Lightwave Technology 13 (5) (1995) 732–739.
[10] Cisco, Cisco ONS 15454 Optical Filter Cards. http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/product_data_sheet09186a00801a5572.pdf (retrieved on: January 2011).
[11] Thorlabs, Optical Attenuators. http://www.thorlabs.de/NewGroupPage9.cfm?ObjectGroup_ID=1385 (retrieved on: January 2011).
[12] Neophotonics, Variable Optical Attenuator: FVOA2000. http://www.neophotonics.com/product/detail.php?id=70 (retrieved on: January 2011).
[13] S. Zsigmond, External report on physical-layer attacks in optical networks, Technical Report, Project SAFE. (http://www.fer.hr/tel/en/research/safe) Supported by the Unity through Knowledge Fund (UKF), Ministry of Science, Education and Sports, Croatia, 2011.
[14] M.G.C. Resende, C.C. Ribeiro, Greedy randomized adaptive search procedures, in: F. Glover, G. Kochenberger (Eds.), Handbook of Metaheuristics, Kluwer Academic Publishers, 2003, pp. 219–249.
[15] R. Inkret, A. Kuchar, B. Mikac, Advanced infrastructure for photonic networks: extended final report of COST action 266, Faculty of Elect. Eng. Comput., Univ. Zagreb, Zagreb, Croatia, 2003, pp. 19–21.
[16] D. Banerjee, B. Mukherjee, Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs, and a reconfiguration study, IEEE/ACM Transactions on Networking 8 (5) (2000) 598–607.
[17] E.Q.V. Martins, M.M.B. Pascoal, A new implementation of Yen's ranking loopless paths algorithm, 4OR: A Quarterly Journal of Operations Research 1 (2) (2003) 121–133.