13

Green Routing/Switching and Transport

Luca Chiaraviglio, Antonio Cianfrani, Angelo Coiro, Marco Listanti
and Marco Polverini
DIET Department, University of Rome "Sapienza", Rome, Italy

13.1 Energy-Saving Strategies for Backbone Networks

In the last decade, the energy consumption of Internet has been hugely fuelled by an increase of the number of connected devices causing an explosion of traffic volume. Consequently, energy-restraint strategies have become a crucial element to prevent the Internet from being throttled by an energy bottleneck. The challenge is to make the increase of the Internet operational efficiency faster than the rate of traffic growth.

Several studies [3] tried to foresee the future Internet energy consumption over the next 15 years (2010–2025). These studies divide the network in two segments: the access segment, realized with heterogeneous technologies, and a backbone segment, composed of high-capacity IP routers and transmission devices. It is foreseen that, at the end of the considered time interval, the overall carried traffic will achieve a total increment of 2–3 orders of magnitude, while the energy per bit consumed by the networking equipment will decrease in the range of 15–20% per year. Basic results of these studies show that, though at present (2013) the access network is about two orders of magnitude higher power than the backbone segment, in the near future, due to the foreseen traffic increase, the overall energy consumption of switching and transport equipment will progressively equal that of the access segment to become the dominant component after the 2020. Moreover, studies are unanimous to indicate that hardware technologies are expected to improve energy efficiency over the 2010–2020 decade, but the rate of the overall network efficiency improvement is expected to be slower than the traffic growth rate.

Green Communications: Principles, Concepts and Practice, First Edition.
Edited by Konstantinos Samdanis, Peter Rost, Andreas Maeder, Michela Meo and Christos Verikoukis.

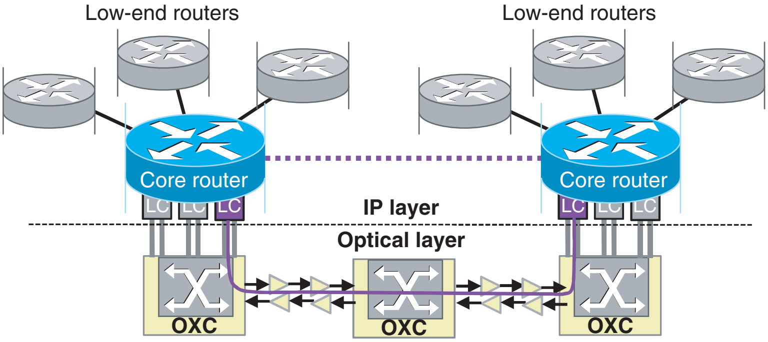

Figure 13.1 IP-over-WDM network architecture

Summarizing, the future design of IP networks oblige to reach an improvement of the energy efficiency of backbone network elements since they will be the dominant components of the overall energy consumption in the core network.

13.1.1 Backbone Networks and Energy Consumption

A typical ISP backbone network architecture is usually based on an IP-over-WDM model (Figure 13.1). The electronic IP routers make use of the huge capacity circuits, called lightpaths, provided by the optical layer. Each IP link corresponds to a dedicated WDM lightpath. A specific IP interface (NIC: Network Interface Card) is associated to each logical IP link originated or terminated by the node; a NIC interacts with Optical Layer by means of an E/O converter. More details regarding energy efficiency in optical network technologies are provided in Chapter 15.

Focusing on power consumption requirements, recent studies (see e.g. Ref. [4]) have shown that the power consumption of IP devices is significant, thus an effective energy-saving strategy must act on IP routers, trying to reduce their power consumption.

An IP router is a very complex device, having an internal structure that depends on the specific vendor. However, it is possible to define a high-level model, reported in Figure 13.2, composed of three main blocks: a router processor, a set of line cards and a switching fabric.

The Router Processor (RP) manages the whole system and implements control plane functionalities, such as routing protocol execution and forwarding table management. Line cards interconnect the router with external networks receiving and sending IP packets; in Figure 13.3 a simple scheme of a line card is shown. Basically, there is an I/O interface and an internal interface connecting the line card with the switching fabric. The core of a line card is the Layer 3 processor that elaborates the header of incoming/outgoing IP packets and executes routing table lookup to perform IP forwarding. In terms of energy consumption, the line cards consume about 45% of the overall router power consumption, and the Layer 3 processor is responsible for about the 60% of line card power consumption [5, 6]. In the rest of this chapter, when referring to the power consumption of an IP link, we specifically refer to the power consumption of the line cards at the two ends of the link. Similarly, switching off a link means here switching off the related line cards.

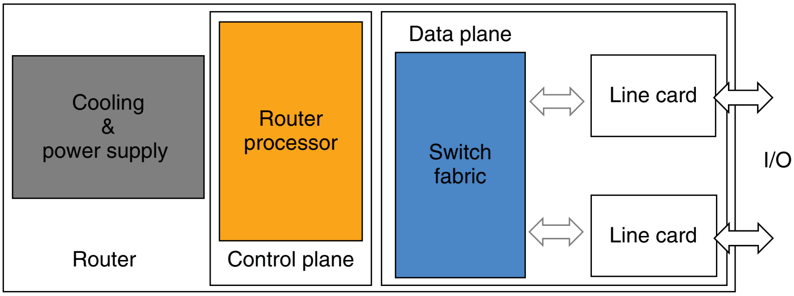

Figure 13.2 Simplified model of an IP router

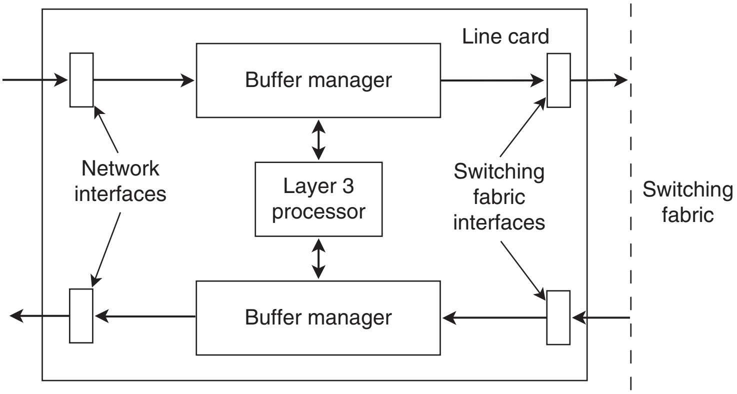

Figure 13.3 General scheme of an IP line card

13.1.2 Energy-Saving Strategies: Switch Off versus Energy Proportional

Traditional Internet design follows two main guidelines: use of equipment redundancy and over-provisioning in processing and transmission capacity normally dimensioned to support much more that the traffic peak values. In this way, network resources are usually largely under-loaded during off-peak periods. However, traffic variations are not followed by a corresponding decrease in the amount of energy consumed by the network because of the nearly independent power consumption of network devices. As a result, a key goal is to provide network energy consumption nearly proportional to the amount of carried traffic. In particular, two possible classes of strategies have been widely discussed in literature to bring the network energy consumption nearly proportional to the amount of carried traffic:

- *Energy proportional strategies*: They work on the individual network devices and try to achieve energy consumption proportionality by adapting the speed and capacity of the devices to the actual load, over relatively short timescales [7].
- *Switch-off strategies*: They involve the network as a whole and approximate load proportionality by carefully distributing the traffic in the network so that some devices are highly

utilized while others become idle and are put in sleep modes, that is, in a state of hibernation or completely switched off [8]. As this class of strategies need an enhancement of the network control procedures, namely routing, to aggregate traffic flows so as to reduce the set of active network resources, these strategies are also often named *Energy-Saving Traffic Engineering* (ES-TE) strategies.

Obviously, solutions belonging to these two classes can be merged, so that energy proportional devices are present, that is, powered-on, while sleep mode can be leveraged to possibly save additional energy. In Ref. [9], an evaluation study is performed considering the two different classes of solutions. As ES-TE enables selected devices to be switched off by steering traffic towards powered-on devices, increasing their power consumption due to higher traffic volumes sustained by them, this study brings light on identifying the scenarios where the switch-off approach is still beneficial.

The basic achievements of the mentioned study are summarized in Figure 13.4, derived from [9]. The curves shows the energy-saving ratio E that represents the ratio among energy consumption in the case of link switch off and energy consumption in case of load proportionality, as a function of p, the fraction of links (nodes) switched off. Results are averaged for a set of network topologies. The two curves are obtained for two different values of the parameter ν, called *equivalent load*; it measures the dependency of the energy consumption of a link (or a node) on the load: if $\nu = 10$, the energy consumption of a link is practically independent of the load; on the contrary, if $\nu = 0.01$, the energy consumption rapidly increases and the load becomes high.

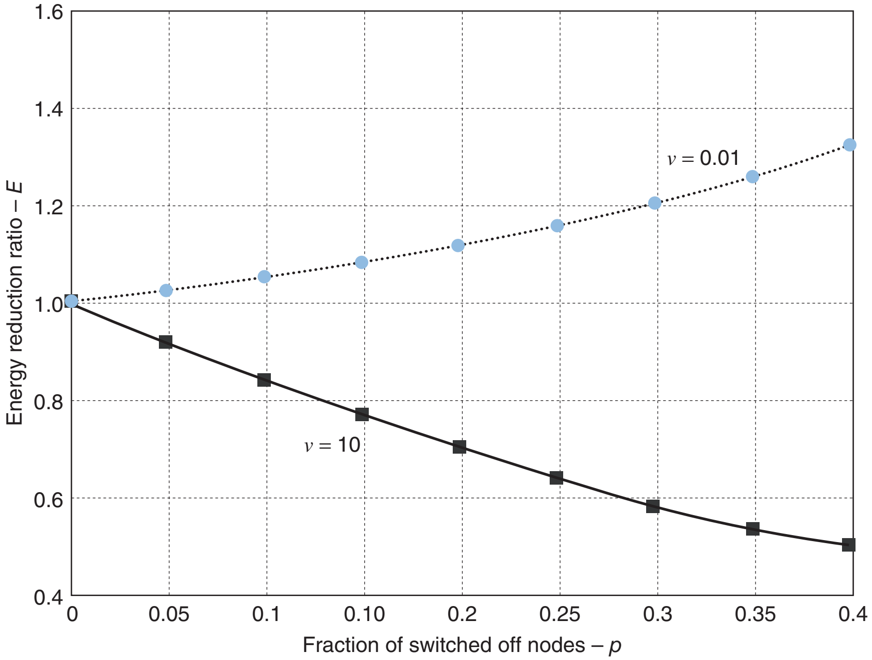

Figure 13.4 E versus p

When $\nu = 10$, the introduction of sleep modes saves energy in the network, that is, $E < 1$. Instead, if $\nu = 0.01$, because $E > 1$, the network consumes a higher amount of energy when devices are switched off. Anyway, in Ref. [9] it is also proved that link switch off is always convenient if $\nu > 0.2$, that is, even if a significant fraction of proportionality is assumed in link operation.

Summarizing, (i) switch-off methods are more convenient when the network is large with a high number of devices, whereas switch-off strategies are not convenient in small networks; (ii) if the device consumption is independent of the load, (or nearly independent, i.e. lightly variable), the use of sleep modes is very effective; on the other hand, if efficient power management techniques are implemented in the network devices, that is, devices are more energy proportional, the effectiveness of sleep mode approaches decreases.

Switch-off strategies aim at keeping the network-wide energy consumption nearly proportional to the amount of carried traffic. Traffic fluctuations are monitored in order to provide traffic aggregation over a subset of the network elements, allowing other equipment to enter power-saving mode and to wake them up only when needed.

In principle, network elements that can be switched off are both links (specifically, the network line cards associated to each IP link originated or terminated by a node) and entire nodes. Of course, an entire node can enter a sleep mode only if all its line cards are switched off. Nodes in sleep mode are essentially removed from the network topology, not able to receive, process and forward packets. Consequently, the node switch off is possible only if the node itself is neither a source nor a destination of traffic. This is the reason why most of studies exclusively concentrate on pure link switch-off strategies.

A particular Switch-Off strategy is possible when Link Aggregation functionality is enabled: A network link, referred to as a bundle, is composed of several physical links, referred to as "member" links. During low utilization times, traffic could be aggregated onto a subset of member links, providing the opportunity for the unused member links to be powered down, thus reducing the energy consumption of the overall bundle.

A classification of the *Switch-Off* strategies can be based on the constraints considered for path computation of packet traffic, that is, the constraints that are applied to determine the new routes that packets have to follow in order to put in sleep mode the highest number of network elements. Two main classes of *Switch-Off* strategies can be defined: (i) *flow-based* solutions and (ii) *destination-based* solutions. In the flow-based solutions, each packet is classified as belonging to a flow, implicitly identified by the source-destination addresses or explicitly denoted by a label or a set of rules, while the routing decision depends on the specific flow a packet belongs to. These approaches may require a connection-oriented layer 2 IP routing, typically implemented by Multi Protocol Label Switching (MPLS) or Open Flow. Hence, energy saving is obtained in a different protocol layer than IP; for this reason such strategies can be either called *MPLS* or Open Flow *solutions*. In the destination-based solutions, the forwarding decisions at each router are taken considering uniquely the destination IP address carried by the incoming packets. Specifically, the destination-based solutions maintain the traditional principles of IP routing.

Alternative strategies that aim to define low-power states avoiding switching-off nodes are based on deactivating a subset of router functionalities, concentrating on those regarding IP packets processing and routing. The selection of these functions is made since they are the most power hungry [4] and a considerable energy saving can be obtained through their deactivation. These strategies are classified as *Forwarding Freezing* approaches and can achieve significant energy saving without switching off entire hardware components. In this way, they are able to limit the impact on network performance, stability, and reliability levels. In Section 13.4, we discuss a specific *Forwarding Freezing* solution, referred to as Table Lookup Bypass (TLB).

13.1.3 Energy-Saving Strategies: Deployment Issues

One of the main limitations for fast deployment of energy-saving techniques in real networks is their integration in actual network protocols. Both Switch-Off and TLB strategies require the implementation of two main features: (i) an extension of the control protocol for the computation of new energy-aware paths and (ii) a device-level functionality to enable the specific power state (switch off or TLB) on the device (i.e. link or node).

The control protocol extension consists in integrating the new energy-aware routing algorithm to be executed during low traffic levels. The routing algorithms can be classified in two main categories: centralized algorithms and distributed algorithms. A centralized algorithm is executed on a single network element and the results (links to power down or node in TLB state) are then applied in the network by specific control-level messages. A distributed algorithm is executed on each network node, so avoiding the generation of control-level messages and then allowing each node to independently manage forwarding functionality. The implementation choice strongly depends on the network scenario considered: a destination-based or a flow-based scenario. In the case of a destination-based scenario, a distributed energy-saving routing algorithm is needed to maintain the principles of IP routing [10]. In the case of flow-based scenario, both centralized and distributed approaches can be used; in the first case, each router executes a distributed algorithm [11]; in the second case a centralized path computation element, such as Path Computation Element (PCE) for GMPLS networks, executes the energy-saving algorithm and manages the routing configurations of network routers [12].

The second step in the deployment of energy-aware solutions is the enabling of low-power state on network device. In the case of actual devices, where no advanced power-saving capabilities are available, the only possible solution is to power off the specific device (i.e. node or link); in this case, the TLB solution cannot be deployed. This solution has a drawback: the generation of a considerable amount of routing protocol-related messages and the network convergence. For instance, in the case of a link-state routing protocol such as OSPF, the removal of a link leads to the generation of a Link State Advertisement message that reaches all network nodes; each node performs path re-computation and updates its routing table. During the convergence phase, the network can experience performance degradation due to temporary unfeasible paths, that is, loops. New energy-aware devices, available in the near future [13], implement advanced power management techniques: (i) the standby of a line card with periodical wake-up allows to maintain physical layer synchronization and to generate routing protocol periodical messages (such as Hello messages for OSPF), making possible a more effective implementation of Switch Off; (ii) the management of line-card functionalities for the implementation of TLB techniques. The coordination among control-level algorithms and advanced hardware capabilities reduces the impact of energy-saving solutions on existing protocol and network performance.

In the case of energy-proportional devices, the implementation issues are limited: The modification of network paths is not needed and so each device acts independently, modifying its processing capability as a function of traffic load. The only aspect to carefully consider is the ability for the device to follow in real time the traffic behaviour and properly react to sudden traffic peaks.

13.2 Switch-Off ILP Formulations

The reduction of energy consumption in backbone networks can be obtained modifying network paths so as to use a subset of network resources; the computation of new paths and the detection of network devices to power off will be the result of energy-saving strategies. In this section, we define this problem by means of an Integer Linear Programming (ILP) formulation.

A network can be modelled as a graph $G(V,E)$, where V is the set of nodes and E is the set of directed edges (or links). The nodes represent the network routers and $N = |V|$ is the total number of router in the networks; the edges represent the network links and $L = |E|$ is the total number of directed edges in the network. The generic node i has a power consumption P_i. The generic edge from node i to node j, indicated with notation (i,j), has a power consumption P_{ij} and a capacity C_{ij}. We consider directed links to have a more general scenario, since the modification to an undirected links scenario can be easily derived. To represents switch-off operation in the formulation, two variables must be introduced:

- x_{ij}, a binary variable indicating if link (i,j) is active ($x_{ij} = 1$) or it is switched off ($x_{ij} = 0$);
- y_i, a binary variable indicating if node i is active ($y_i = 1$) or it is switched off ($y_i = 0$).

The energy minimization problem can be formulated by means of the following objective function:

$$\min \left[\sum_i \sum_j x_{ij} P_{ij} + \sum_i y_i P_i \right] \tag{13.1}$$

The scope is to minimize the overall consumption of nodes and links.

The traffic demand among network nodes can be represented by means of an $N \times N$ traffic matrix $T = [t_{sd}]$, where the element t_{sd} represents the amount of traffic originated from node s and directed to node d.

The complete formulation of the energy minimization problem for an IP network requires the introduction of several constraints: In particular, traffic-related constraints are needed to determine the network routing able to minimize the overall energy consumption, while still satisfying traffic requirements. These constraints are thus related to the network scenario considered, that is the flow-based or the destination-based one. For the sake of simplicity in the problem formulation, we introduce an assumption regarding paths computation: For each source-destination pair, a single path is computed. The extension to the multi-path case can be easily done but since it will lead to a more complex formulation we limit our analysis to the single path case.

13.2.1 Flow-Based Routing Formulation

In a flow-based network scenario, each traffic relationship t_{sd} is routed independently by means of a dedicated flow, that is, a dedicated path from the source to the destination. To take this into account, the binary variable f_{ij}^{sd} is introduced: f_{ij}^{sd} is equal to 1 only if the flow from s to d is routed through edge (i,j).

The first constraint to be considered is the classical flow conservation constraint:

$$\sum_{j=1}^{N} f_{ij}^{sd} t^{sd} - \sum_{j=1}^{N} f_{ji}^{sd} t^{sd} = \begin{cases} t^{sd} & \text{if } i = s \\ -t^{sd} & \text{if } i = d \\ 0 & \text{if } i \neq s, d \end{cases} \qquad \forall s, d \in V \tag{13.2}$$

The second one is the bandwidth constraint: allowing a maximum utilization value α_{ij} for each edge (i,j), the amount of traffic routed on each link is limited by the following equation:

$$\sum_{s=1}^{N} \sum_{d=1}^{N} f_{ij}^{sd} t^{sd} \leq x_{ij} \alpha_{ij} C_{ij} \quad \forall (i,j) \in E \tag{13.3}$$

Another effect of equation (13.3) is that only unused links can be switched off ($x_{ij} = 0$).

The last constraint to be considered regards the node switch-off condition, and use the big-*M* notation to make possible the switch off of a node only if all its links are switched off:

$$\sum_{j=1}^{N} x_{ij} + \sum_{j=1}^{N} x_{ji} \leq M y_i \quad \forall i \in V \tag{13.4}$$

The *M* value is an integer constant assuming a "big" value so that even when all x_{ij} variables are equal to 1 the first term of Eq. (13.4) is smaller than *M*; on the other side, a node can be switched off ($y_i = 0$) only when all x_{ij} variables are equal to zero.

The objective function (13.1) and the constraints (13.2), (13.3) and (13.4) represent the general energy minimization problem in a flow-based network scenario.

The general formulation can be adapted to specific cases introducing further constraints. One of the most common implements a physical requirement of router interfaces: It is possible to switch off a directed link only if also the link on the other direction is switched off:

$$x_{ij} = x_{ji} \quad \forall (i,j) \in E \tag{13.5}$$

13.2.2 Destination-Based Routing Formulation

The destination-based formulation is introduced to model the conventional data plane of IP networks and in particular the forwarding action performed by IP routers. A router selects the outgoing interface for an incoming IP packet exclusively based on the destination IP address; so, all the packets entering a router and directed to the same destination will be forwarded on the same output link, regardless of the source that generated them or the communication flow they belong to. Therefore, packets emitted by the network nodes directed to a given destination node are constrained to follow a *Reverse Path Tree* (RPT) having as root the destination node itself.

To formulate the energy minimization problem in such a scenario, it is necessary to modify the traffic assumption provided for the flow-based case. Traffic demands cannot any longer be split in source-destination flows, which are treated independently, since different flows having the same destination node must "share" respective parts of their route towards such common destination. To take this into account, a new binary variable n_{ij}^{d} must be introduced in the MILP formulation: n_{ij}^{d} is equal to 1 only if node *i* uses edge *(i,j)* to route traffic directed to destination *d*; in other words, using the IP terminology, *j* is the next-hop router of *i* to reach destination *d*. Otherwise, it remains equal to 0.

Starting from the flow-based formulation, two further constraints must be introduced to define the destination-based problem. The first one is needed to relate the new n_{ij}^{d} variables with the flow variables:

$$\sum_{s=1}^{N} f_{ij}^{sd} \leq M n_{ij}^{d} \quad \forall i, d \in V \tag{13.6}$$

The second one is needed to maintain the single path assumption, that is, the use of a single next hop for each destination:

$$\sum_{j=1}^{N} n_{ij}^{d} = 1 \quad \forall i, d \in V \tag{13.7}$$

13.2.3 Comparison of Flow-Based and Destination-Based Formulations

The MILP formulations proposed in the previous section highlight that the solutions space of the energy minimization problem in the destination-based scenario is included in the solution space of the same problem in the flow-based scenario. A more rigorous analysis on the optimal solution in both cases is not possible since the problems are both NP-hard and can only be solved for small network scenarios. In Figure 13.5, we provide the solution of both formulations for a 12-node full mesh network considering a reference peak traffic matrix scaled to obtain nine different traffic scenarios.

Figure 13.5 shows that when the traffic is lower than 60% of peak values, the two problems have the same optimal solution, while in other cases, the flow-based optimal solution is slightly better than the equivalent destination-based one.

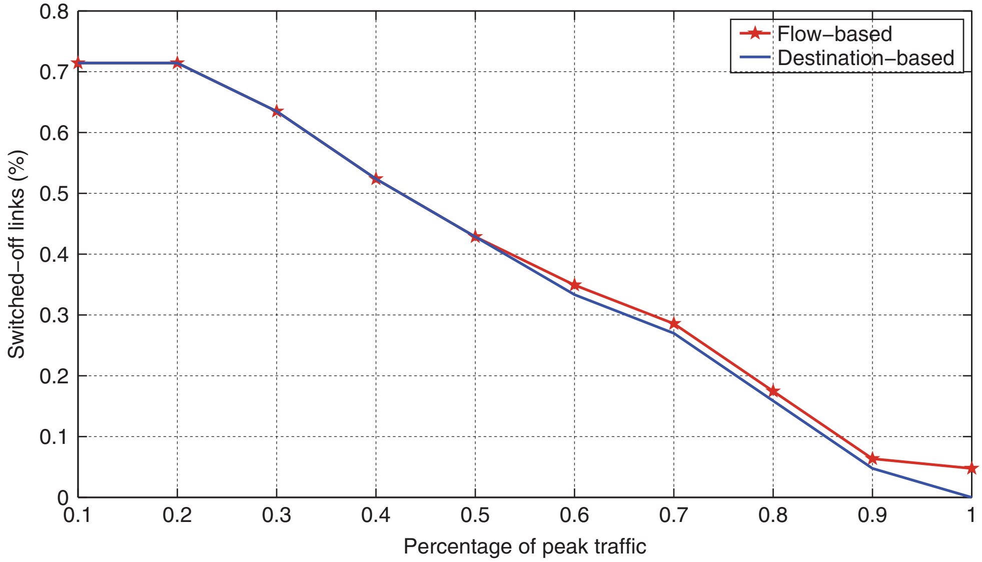

Figure 13.5 Optimal solution of flow-based (FLOW) and destination-based formulations (DBF) in a 12-node network scenario

As mentioned before, the problem formulations analyzed can be solved only for small networks. In real network scenarios, a low-complexity energy-saving algorithm must be considered, in order to compute the set of devices to switch off and the new routing configuration in a reasonable time. The next section provides a brief overview of the most common flow-based and destination-based energy-saving algorithms.

13.3 Switch-Off Algorithms

In this section, we provide a brief description of energy-saving algorithms for backbone networks proposed in literature classifying the solutions into flow and destination-based.

13.3.1 Flow-Based Algorithms

13.3.1.1 Least Flow Algorithm (LFA), Most Power Algorithm (MPA) and L-Game

The Least Flow Algorithm (LFA) [14], the Most Power Algorithm (MPA)[14] and L-Game [15] are different algorithms based on the same operative modality: the set of device (either a node or a link) is first sorted according to a given rule, and then the devices are selectively powered off. For each link powered off, the heuristics re-compute the routing, and check the connectivity and the maximum link utilization constraints. If such constraints are satisfied, the current device is left powered off; otherwise it is powered on again. The heuristics then try to power off another device, until all the devices are considered. The variations among the different heuristics rely on the rules adopted to sort list the network devices. In particular, LFA sorts the devices according to the load switching off first devices with the least amount of traffic flowing through them, since the least used devices are more likely to be safely powered off. On the contrary, MPA sorts the devices with increasing power, resulting in a more aggressive behaviour. Finally, L-Game adopts a game theory approach to sort the network links using the Shapley value. This metric takes into account both the network topology and the current traffic conditions. These values can then be used to identify the less “critical” links, which will be the first to be switched off. The types of algorithms have the potential to provide significant energy savings ranging between 30% and 80% according to respective evaluation studies.

13.3.1.2 Energy Profile Aware Routing (EPAR)

The Energy Profile Aware Routing (EPAR) algorithm [16] targets the reduction of power of links in a backbone network. The main idea is to define an energy profile for each link. The profile represents the dependency of power with respect to the load flowing on the link. Different assumptions are taken into account, ranging from a classical on–off model (i.e. the link can be either powered off or powered on) to more complex profiles in which power scales logarithmically or exponentially with the load. While the on–off model is realistic for current devices, the other profiles can be representative of future devices scaling the power proportionally with the load.

The problem is then formulated as an optimization model in which the objective function is the minimization of the total link power, still guaranteeing maximum link load and connectivity among nodes. The problem is then optimally solved over a realistic case study, showing that

EPAR can save more than 35% of energy compared to a classical formulation adopting shortest path routing and not targeting explicitly the power of links.

13.3.1.3 Green Distributed Algorithm (GRiDA)

The Green Distributed Algorithm (GRiDA) [17] aims at reducing the power consumption by switching off links while guaranteeing Service Level Agreements (SLAs). Different from LFA/MPA/L-Game approaches, switch-off decisions are completely decentralized. In particular, each node continuously monitors the utilization and the power status of the incident links. Based on this information, the GRiDA algorithm (i) switches off links when they are underutilized and (ii) switches on certain links back when traffic increases. The switch-on and switch-off decisions are taken at random intervals. Such decisions are based on the minimization of a utility function that integrates the power consumption of links with history information regarding past decisions. Moreover, the GRiDA algorithm assumes periodic Link State Advertisements (LSAs) exchanged among nodes reporting the status of the network considering the following two values: normal operation and congestion. When link congestion is detected, this information is flooded in the network towards other nodes, which power on the appropriate incident links. The GRiDA algorithm is able to achieve up to 50% energy savings with a limited impact on the network SLAs.

13.3.1.4 Distributed and Adaptive Interface Switch Off for Internet Energy (DAISIES)

Distributed and Adaptive Interface Switch off for Internet Energy Saving (DAISIES) [11] is an algorithm that exploits the control mechanisms provided by MPLS to switch-off network links. The actual amount of traffic carried by each Label Switched Path (LSP) is monitored by the ingress node on a fine granular observation period (e.g. tens of seconds). Then, whenever traffic goes beyond a prefixed threshold, the ingress node re-computes the path of the LSP and reroutes it updating both the path and the reserved bandwidth (make-before-break). The information about available (unreserved) bandwidth advertised by the Traffic Engineering (TE) routing protocol is used by DAISIES to properly compute link weights, and in turn aggregate the traffic on a reduced set of links. With this solution, energy savings between 35% and 50% are possible.

13.3.1.5 Green Traffic Engineering (GreenTE)

The Green Traffic Engineering (GreenTE) [18] is an intra-domain traffic engineering mechanism, which focuses on switching off links, considering the maximum link utilization and the network delay constraints. Rather than exploring the full space of paths among each source and destination, the authors consider a heuristic solution in which only a subset of paths is considered as input to the problem. Each candidate path has to meet two properties: (i) the candidate path cannot be larger than the diameter of the network, and (ii) the candidate path length cannot be larger than twice the shortest path for a given source destination node pair.

GreenTE is run by a centralized controller which continuously collects information from network routers, solves the traffic engineering problem and then notifies the corresponding routers about the links that should be powered-off and/or powered-on. Such a process is suggested to be run every 5–15 minutes to prevent frequent routing disruptions. The network topology, the link load and the state of the links (power on and power off) are distributed via LSAs adopting the Traffic Engineering Metric of the Open Shortest Path First protocol (TE-OSPF). Results show potential power saving between 20% and 40% can be obtained, while keeping both link delay and queuing delay close to the conventional case without GreenTE.

13.3.1.6 Energy-Aware Traffic Engineering (EAT)

Energy-aware traffic engineering (EAT) [19] considers the energy consumption of links while maintaining the same traffic rates between source and destination pairs as the traditional non-energy-aware approaches. The main assumption of this approach is the capability of links to operate at different transmission rates, while the basic idea is to spread the load across different paths with the objective to increase energy saving. In particular, the proposed algorithm shifts traffic towards links that can forward it without increasing their operation rate, provided that this would reduce the operation rate of other links. Every source node makes an independent decision based on the information collected from paths used towards certain destinations using a drop margin parameter. If the link utilization is below the drop margin, traffic is shifted from the link to other links. Hence, the drop margin plays a crucial role to select the number of links that will be considered to increase energy saving. In particular, the higher is the drop margin, the higher will be the number of links operating in low-power mode.

The EAT algorithm collects information regarding the candidate links considered to shift away their traffic and uses a metric to specify the distance from the lower operating rate. This metric is normalized by the utilization and represents the fraction of traffic that needs to be removed in order to enable the link to operate in a lower energy state. Since decisions are taken at the source nodes, the destination simply reports the amount of traffic that needs to be shifted and then the source checks the feasibility of such suggestion before performing the appropriate traffic alternations. EAT is able to switch off between 15% and 31% of the links and between 10% and 24% of the nodes.

13.3.1.7 Greening Backbone Networks with Bundled Links (GBNB)

Green backbone networks with bundled links [20] aim to switch off links that compose a link aggregation bundle at off-peak times when the traffic load is low, in order to save as much energy as possible. In particular, links are first sorted considering different rules and then are selectively powered off until no further link can be removed without causing service disruption. The heuristic initially computes the amount of flow on each link. In particular, an ILP model is formalized. The objective function is the minimization of the amount of flow on each link. Then, as a second step, the algorithm checks the amount of spare capacity on each link. Based on this information, the algorithm then turns off some of the links composing the bundle. Energy savings between 35% and 75% are achievable.

13.3.1.8 Green MPLS Traffic Engineering (GMTE)

Green MPLS Traffic Engineering [21] employs a traffic engineering means to switch off routers in a backbone network. The proposed heuristic takes as input the traffic demands and

Table 13.1 Main features of the flow-based algorithms

Algorithm	Computation	Targeted devices	Constraints	Complexity
LFA/MPA	Centralized	Routers/links	Connectivity, maximum link utilization	Central controller computing the set of devices to be powered off
L-Game	Centralized	Links	Connectivity, maximum link utilization	Central controller computing the set of links to be powered off
EPAR	Centralized	Links	Connectivity, maximum link utilization	Central controller computing the set of links to be powered off
Grida	Distributed	Links	Connectivity, maximum link utilization	Single node maximizing a utility function
DAISIES	Distributed	Links	Connectivity, maximum link utilization	Ingress nodes monitoring a traffic threshold
GreenTE	Centralized	Links	Connectivity, maximum link utilization, link delay	Central controller computing the set of links to be powered off
EAT	Distributed	Routers/links	Connectivity, maximum link utilization	Source node that shifts traffic the paths to the destinations
GBNB	Centralized	Bundled Links	Connectivity, maximum link utilization	Central controller computing the part of bundle to be powered off
GMTE	Centralized	Routers	Connectivity, maximum link utilization	Central controller computing the set of routers to be powered off

the routing information, and produces as output the set of powered off routers and the set of alternate label-switched paths with energy saving for re-routing traffic. The heuristic works as follows: Routers are first sorted according to the number of LSPs that they carry. Then, the routers are selectively considered starting from the ones with the least number of LSPs. For each router, traffic is shifted on new LSPs that are created. If the maximum link utilization does not exceed a threshold, the router is powered off. Otherwise, the router is kept powered on and the initial LSPs are restored. Results show that energy savings between 18% and 30% are achievable with the proposed solution.

Table 13.1 summarizes the main features of each algorithm.

13.3.2 Destination-Based Algorithms

13.3.2.1 Energy Saving IP Routing Strategy (ESIR)

The Energy-Saving IP Routing Strategy (ESIR) [10] is designed to be integrated into the OSPF routing protocol. The path computation strategy, realized by means of a modified Dijkstra algorithm, is distributed and fully compatible with OSPF mechanisms. At the same time ESIR is able to satisfy QoS requirements by maintaining traffic load on all the network links under fixed configurable values. The modified version of the Dijkstra algorithm is able to select a subset of paths to route the traffic, leaving the unused interfaces to enter low-power mode. The main advantage of this approach is that the IP topology does not change, and consequently no exchange of LSAs packets is needed. To achieve this goal, the set of routers is divided in importers and exporters. The main idea is the exportation mechanism, which allows to share a Shortest Path Tree between neighbours routers, so that the overall set of active links is minimized. Only the importer routers modify their shortest path tree starting from the tree of the exporter routers. This allows to avoid triggering an entire path re-computation inside the network. In order to work properly, the exporter has to be the neighbour of the importer. Moreover, an importer can be object of a single exportation, that is, it can receive the shortest path tree from a single exporter. Finally, once a node has become an exporter, its shortest path tree cannot be modified any more. The obtained results show that it is possible to reduce a percentage equal to about 40% the number of active links.

13.3.2.2 Energy Saving Based on Algebraic Connectivity (ESACON)

Energy Saving based on Algebraic Connectivity (ESACON) [22] adopts the algebraic connectivity to identify network links to be powered off. The algebraic connectivity is a metric developed in the context of complex networks, while in this work it is used as a parameter to keep the network connectivity above a suitable level. In particular, the algorithm works in two steps: It first creates a set of ordered links and then a subset of links is powered off. The links are ordered based on the value of the algebraic connectivity, since the aim is to switch off those links that have a low impact on the network connectivity. In the second step, a set of links is powered off in order to maintain the network connectivity above a given threshold. Despite not considering directly traffic flowing in the network, the ESACON algorithm reveals comparable performance with respect to traffic-aware solutions. Result show that the connectivity threshold plays a crucial role in determining the fraction of links powered off. Therefore, this parameter needs to be carefully set. Results show that up to 73% of links can be switched off in a realistic scenario, while maintaining the same link utilization as the LFA algorithm.

13.3.2.3 Ant Colony-Based Self-Adaptive Energy Saving Routing for Energy-Efficient Internet

Ant colony based energy-aware routing [23] relies on the concept of traffic centrality to switch off links, maximizing the energy saving. In particular, the traffic centrality is a measure of traffic volume on a link connecting a node to a neighbour. The authors proved that when this metric is maximized the energy consumption is minimized. Therefore, they propose a heuristic solution to maximize the traffic centrality. In particular, the heuristic, called A-ESR, is based on the ant colony technique, in which a number of artificial individual ants in the employed colony explore the network in real time. There are two types of artificial ants: the forward ant and the backward ant. At regular intervals, forward ants are launched from the nodes to other randomly chosen nodes. While moving to the destination, forward ants gather network information, including the arrival time and the identifier of each node. Once the forward ant

Table 13.2 Main features of the destination-based algorithms

Algorithm	Computation	Targeted devices	Constraints	Complexity
ESIR	Centralized	Links	Connectivity, maximum link utilization	Central controller computing the list of exporters and importers
ESACON	Centralized	Links	Connectivity	Central controller computing the link to be powered off
A-ESR	Distributed	Links	Connectivity, maximum link utilization	Single nodes computing the links to be powered off considering the information injected by "ants" in the network

has received the destination, a backward ant is generated. While moving to the original source, the backward ant applies the decision of the energy-aware algorithm for all nodes along the path. Results show an energy efficiency up to 71% in a realistic scenario.

Table 13.2 summarizes the main features of each algorithm.

13.4 Table Lookup Bypass

The "table lookup bypass" (TLB) is an alternative strategy to improve the energy efficiency of IP backbone networks [22] by avoiding using entirely or partly the forwarding capacity of a router, allowing a low-power mode for the hardware specifically dedicated to this function, that is, the L3 Processor. Since a relevant part of the power consumed by an IP router is wasted in performing layer 3 processing and, specifically, in forwarding packets, such an approach is very promising and could be a valid alternative to link switch off.

The main idea under TLB is to deterministically forward some or all packets entering a node to a specific and fixed next-hop router regardless of their destination address. By doing so, the lookup operation of the routing table, which is the main operation performed by the L3 Processor of IP routers, is bypassed and, consequently, the related hardware can be frozen.

Similarly to link switch off, this general mode of operation may take place during low traffic load periods during which the actual routing operation that should be performed by a specific router is delegated to some other routers within the network. Moreover, as TLB does not determine a reduction of the transmission capacity of the network, it may also take place during medium/high traffic load periods. In more detail, the TLB strategy may lead to changing the routing of some traffic flows within the network as elaborated in detail in the next section, which will generally experience a path length increase. Consequently, some spare capacity is generally needed within the network in order to support the increased load. In contrast, link switch off leads to both a path length increase of traffic flows and a reduction of the active transmission capacity. Therefore, link switch off generally needs higher reductions of the offered load in order to allow a significant number of links to be powered off.

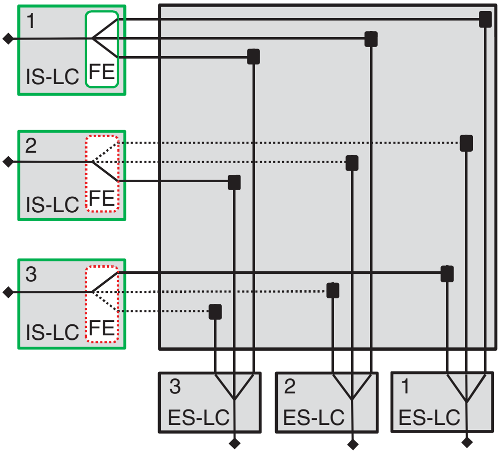

Figure 13.6 Graphical representation of a router with TLB enabled on all line cards

13.4.1 General Model and Implementation Aspects of TLB

Let us consider the router model presented in Figures 13.2 and 13.3, wherein each line card implements its own L3 Processor which is responsible for forwarding the packets received by the line card. This is performed by inspecting the forwarding table in order to find the correct next-hop router to forward packets and consequently the correct egress line card.

In such a model, it is possible to freeze the L3 Processor of a line card regardless of the state of the other line cards of the same router. Line card whose L3 Processor has been frozen cannot perform TL operation and must forward packets to a predetermined egress LC. LCs in this state are said to be in TLB state.

Figure 13.6 shows a graphical representation of a node having TLB-enabled LCs. IS-LCs are depicted on the left side of the figure whilst ES-LCs on the bottom. The node is represented as a set of cross points, which identify all possible switching operations between each IS-LC and each ES-LC. These switching operations are actually allowed only if there exists a path to reach the related cross point. According to the state of each IS-LC, only a subset of these paths can be actually used to forward packets, which are represented by continuous lines, whilst other ones, represented by dotted lines, cannot. The first IS-LC (top-left of the figure) is working normally forwarding received packets to any ES-LC. The second IS-LC (middle-left of the figure) is in TLB state towards the third ES-LC. Finally, the third IS-LC (bottom-left of the figure) is in TLB state towards the first ES-LC. Note that, if an LC is in TLB state, the related IS-LC can forward packets to only one ES-LC, whilst the related ES-LC can receive packets from any IS-LC. This mode of operation leads to change the routing of some traffic flows but allows to still exploit all the capacity deployed in the network to carry traffic demands.

Figure 13.7 shows a simple example of a TLB operation applied on the LC related to link (A,B) on node B, which is put in TLB state towards the outgoing interface on link (B,C). In the depicted scenario, the traffic demand t_{AC} follows its original path whilst the traffic demand t_{AD} is deviated from its original path A → B → D when enters node B (the dashed line in Figure 13.7) and routed on the path A → B → C → D. This operation allows the LC of B on (A,B) to enable the TLB and so to reduce the power consumption.

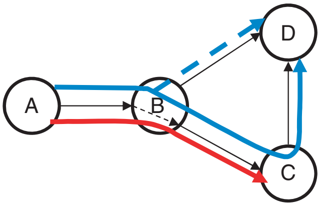

Figure 13.7 Example of paths modification in a network with a LC in TLB mode

A key point concerns how to implement the TLB functionality in current routers LCs. With reference to the LC scheme presented in Figure 13.3, we can note that, when an LC is in TLB state, packet redirection can be easily performed outside the IP packet processor and the lookup engine. For instance, this operation may be included in the hardware chip performing buffer management operations, since this module usually drives I/O operations towards other line cards and local ports. In this way, the L3 Processor represented in Figure 13.3 can be switched off (or put in a low-power state).

13.4.2 Network-Wide Solution

The application of TLB strategy foresees that each specific LC may be either in TLB state or work as usual (FULL state) regardless of the state of any other LC, including those belonging to the same router. Moreover, for each LC in TLB state, a specific egress LC (or equivalently a next-hop router) must be chosen.

A network-wide solution for TLB strategy may provide the state of each LC and the specific predetermined next-hop router for each LC in TLB state. The solution must consider routing and capacity constraints and should maximize the number of LCs put in TLB state.

Such a problem has been modelled as ILP [22], whilst sub-optimal heuristic solutions have been proposed in Ref. [23].

The ILP formulation includes typical flow conservation and capacity constraints of multi-commodity flow problems in addition to a set of specific constraints related to the unique state of each LC. The complexity of the general TLB problem turns out to be higher than the link switch off problem making the problem intractable even for small networks (i.e. even ten nodes).

In order to give an estimation of the energy savings achieved by the TLB approach in comparison to the link switch off, tone needs to consider (i) how much energy can be saved by putting in TLB state a single LC and (ii) how many LCs can be put in TLB state.

Concerning the first point, we have to specifically refer to the amount of power wasted by the L3 processor with respect to the total power consumed by an LC. As outlined in Section 13.4.1, L3 functionalities are responsible of about 60% of the power consumed by an LC. Thus, we expect that the power consumption of an LC in TLB state is reduced of about 50%.

The actual energy saving that can be achieved via TLB use depends on the amount of LCs that may enter TLB state according to the offered load. In this respect, it has been shown that a significant percentage of LCs can be put in TLB state even during peak traffic load periods. Specifically, in Ref. [22] it is shown to be between 30% and 80% depending on the average nodal degree of the network, with more meshed network experiencing higher performance. Moreover, it has also been shown that, as the traffic load decreases, up to 80% of LCs can enter TLB state even in little meshed networks. These results refer to small networks for which it was possible to optimally solve the problem. However, these general results have also been confirmed for large real ISP topologies considering heuristic approaches [23], which lead to about 40% of LCs in TLB state when traffic is at the peak level and about 90% when it is scaled down to 10% of the peak.

13.5 Conclusion

In this chapter, we proposed a classification of energy-aware techniques for wired backbone networks. We focused our attention on Switch-Off strategies, since they represent the preferred choice in the near future from a technological point of view. We first provided two different formulations for the energy minimization problem: the flow-based and the destination-based formulations; then we described the most known heuristics proposed in literature, showing that a considerable energy saving is possible in actual backbone networks exploiting both network overprovisioning and daily traffic behaviour. Finally we introduced an advanced power saving techniques for backbone networks: the TLB; this technique allows to manage the forwarding capabilities of a network device so that to enable low-power state during low-level traffic periods. The TLB technique represents a first step towards mid-term energy-aware techniques that will allow energy consumption of network devices “to closely follow” the traffic levels.

References

[1] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch: “Energy consumption of telecommunication networks and related improvement options”, IEEE J. Select. Top. Quantum Electron., vol. 17, no. 2, pp. 285–295, 2011.

[2] F. Idzikowski, L. Chiaraviglio, and F. Portoso, “Optimal Design of Green Multi-layer Core Networks”, Third International Conference on Future Energy Systems (e-Energy 2012), Madrid, Spain, 2012.

[3] R. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. Ayre, and W. Sorin, “Evolution of WDM optical IP networks: A cost and energy perspective”, J. Lightwave Technol., vol. 27, no. 3, pp. 243–252, 2009.

[4] S. Aleksic, “Analysis of power consumption in future high-capacity network nodes”, IEEE/OSA J. Opt. Commun. Netw., vol. 1, no. 3, pp. 245–258, 2009.

[5] L. A. Barroso and U. Holzle, “The case for energy-proportional computing”, IEEE Comput., vol. 40, no. 12, pp. 33–37, 2007.

[6] R. Bolla, R. Bruschi, A. Cianfrani, and M. Listanti, “Enabling backbone networks to sleep”, vol. 25, no. 2, pp. 26–31, 2011.

[7] L. Chiaraviglio, D. Ciullo, M. Mellia, and M. Meo, “Modeling Sleep Mode Gains in Energy-Aware Networks”, Comput. Netw., vol. 57, no. 15, pp. 3051–3066, 2013.

[8] A. Cianfrani, V. Eramo, M. Listanti, and M. Polverini, “An OSPF-integrated routing strategy for QoS-aware energy saving in IP backbone networks”, IEEE Transactions on Network and Service Management, vol. 9, no. 3, pp. 254–267, September 2012.

[9] A. Coiro, M. Listanti, A. Valenti, and F. Matera, “Energy-aware traffic engineering: A routing-based distributed solution for connection-oriented IP networks”, Comput. Netw., vol. 57, no. 9, pp. 2004–2020, 2013.

[10] A. Tzanakaki et al., “Energy Efficiency in integrated IT and optical network infrastructures: The GEYSERS approach,” INFOCOM Workshop, 2011.

[11] R. Bolla, R. Bruschi, F. Davoli, L. Di Gregorio, P. Donadio, L. Fialho, M. Collier, A. Lombardo, D. Reforgiato, and T. Szemethy, “The green abstraction layer: a standard power management interface for next-generation network device,” IEEE Internet Comput., vol. 17, no. 2, pp. 82–86. 2013.

[12] L. Chiaraviglio, M. Mellia, and F. Neri, Minimizing ISP Network Energy Cost: Formulation and Solutions, EEE/ACM Transactions on Networking, vol. 20, no. 2, pp. 463–476, April 2012.

[13] A. P. Bianzino, C. Chaudet, S. Moretti, J.-L. Rougier, L. Chiaraviglio, and E. Le Rouzic, Enabling Sleep Mode in Backbone IP-Networks: a Criticality-Driven Tradeoff, IEEE ICC'12 Workshop on Green Communications and Networking, Ottawa, Canada, 2012.

[14] J. Restrepo, C. Gruber, and C. Mas Machuca, "Energy profile aware routing", Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on. IEEE, 2009, Dresden, Germany, 2009.

[15] A. P. Bianzino, L. Chiaraviglio, M. Mellia, and J.-L. Rougier, GRiDA: green distributed algorithm for energy-efficient IP backbone networks, Comput. Netw., vol. 56, no. 14, pp. 3219–3232, 2012.

[16] M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-Aware Traffic Engineering", IEEE ICNP 2010, Kyoto, Japan.

[17] N. Vasić and D. Kostić. "Energy-aware traffic engineering". Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking. ACM, Passau, Germany, 2010.

[18] W. Fisher, M. Suchara, and J. Rexford, "Greening backbone networks: Reducing energy consumption by shutting off cables in bundled links," in Proc. 1st ACM SIGCOMM Workshop Green Netw., New Delhi, India, pp. 29–34, 2010.

[19] H.W. Chu, C.C. Cheung, K.H. Ho, and N. Wang, "Green MPLS traffic engineering", Australasian Telecommunication Networks and Applications Conference (ATNAC), IEEE, pp. 1–4, 2011.

[20] F. Cuomo, A. Abbagnale, A. Cianfrani, and M. Polverini: "Keeping the Connectivity and Saving the Energy in the Internet", IEEE INFOCOM 2011 Workshop on Green Communications and Networking, Shanghai (China), 2011.

[21] Y.-M. Kim, E.-J. Lee, H.-S. Park, J.-K. Choi, and H.-S. Park, "Ant colony based self-adaptive energy saving routing for energy efficient internet", Comput. Netw., vol. 56 no. 10, pp. 2343–2354, 2012.

[22] A. Coiro, M. Polverini, A. Cianfrani, and M. Listanti, "Energy saving improvements in IP networks through table lookup bypass in router line cards," Computing, Networking and Communications (ICNC), 2013 International Conference on, pp. 560–566, 2013.

[23] A. Cianfrani, A. Coiro, M. Listanti, and M. Polverini, "A Heuristic Approach to Solve the Table Lookup Bypass Problem", The 24th Tyrrhenian International Workshop on Digital Communications, Genoa (Italy), 2013.