

How Much Can The Internet Be Greened?

Original

How Much Can The Internet Be Greened? / Chiaraviglio, Luca; Ciullo, Delia; Leonardi, Emilio; Mellia, Marco. - (2009), pp. 1-6. (Intervento presentato al convegno 2nd International Workshop on Green Communications (GreenComm 2) tenutosi a Honolulu, Hawaii nel Dicembre 2009) [10.1109/GLOCOMW.2009.5360677].

Availability:

This version is available at: 11583/2281020 since:

Publisher:

IEEE

Published

DOI:10.1109/GLOCOMW.2009.5360677

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

How Much Can The Internet Be Greened?

Luca Chiaraviglio, Delia Ciullo, Emilio Leonardi, Marco Mellia
Politecnico di Torino, Italy, Email: {lastname}@tlc.polito.it

Abstract—The power consumption of the Internet is becoming more and more a key issue, and several projects are studying how to reduce its energy consumption. In this paper, we provide a first evaluation of the amount of redundant resources (nodes and links) that can be powered off from a network topology to reduce power consumption. We first formulate a theoretical evaluation that exploits random graph theory to estimate the fraction of devices that can be removed from the topology still guaranteeing connectivity. Then we compare theoretical results with simulation results using realistic Internet topologies. Results, although preliminary, show that large energy savings can be achieved by accurately turning off nodes and links, e.g., during off-peak time. We show also that the non-cooperative design of the current Internet severely impacts the possible energy saving, suggesting that a cooperative approach can be investigated further.

I. INTRODUCTION

The energy consumption is becoming a sensible topic to which both people and the research community are devoting increasing attention. The ICT makes no exception, and more and more activities and projects are studying how to reduce the energy waste. Current estimates indicate that ICT is responsible for a significant fraction of the world power consumption, ranging between 2% and 10% (the latter figure including also the manufacturing and cooling costs of ICT devices), as reported in [1]. To reduce energy consumption (and costs), large data centers and telecommunication networks, as well as the Internet, are identified as possible targets for optimization.

To this extent, the study of power-saving network devices has been introduced over these years, starting from the pioneering work of [2]. In [3] we faced the problem of defining which is the minimum set of routers and links that have to be used in order to support a given traffic demand under QoS constraints. Unfortunately, the complexity of the problem does not allow to study cases with large networks.

In this paper, we focus on the Internet-wide backbone network, and in particular we aim at estimating the worldwide amount of resources that potentially are redundant in the current Internet topology. We do not directly tackle the energy consumption figure, since the actual energy footprint of devices is hard to know. We rather simply count the number of resources (nodes and links) that can be possibly powered down still guaranteeing the service, e.g., during the off-peak periods when network load is sufficiently light (at night, during weekends, holidays etc). Our goal is to have a first estimate of the possible savings that the adoption of smart energy saving policies may entail. We base our considerations on purely topological properties of the Internet graph, ignoring the effect of traffic flowing in the system. We recognize that

this work is somehow preliminary, since a careful evaluation of the network devices to be switched off cannot ignore the traffic flowing in the network. Nevertheless our study permits in a simpler way to estimate the possible gain margins, providing a first answer to the question of whether it could be worth to include in the future Internet design the capability of selectively turning on/off nodes and links still matching the traffic demand and QoS.

We tackle this problem using both analysis and simulation. Given the graph that models the Internet router level topology, nodes belong to two classes: actual source and destination nodes (called terminals in the remaining of the paper), and pure transit nodes, i.e., nodes that are neither source nor destination of information. Clearly, nodes in the first class cannot be turned off, while transit nodes (and corresponding links) can be switched off, still guaranteeing that the residual graph is connected. We consider both *One-connected graphs*, i.e., graphs in which a single path is guaranteed for each source destination node pair, and *Two-connected graphs*, i.e., graphs in which at least two distinct paths exist to guarantee eventual failure recovery. Given the terminal set, the minimum set of links and nodes that are part of a one-connected graph forms a Steiner Tree [4]. Since the minimum Steiner Tree is known to be a NP-hard problem in general graphs, we exploit random graph theory to estimate the number of devices that can be eventually powered off. This allows us to evaluate the average figure, and to consider very large graphs to see how the saving figure evolves considering a worldwide topology. Then, in order to validate analytical results accounting for more complex (and possibly realistic) graphs, we consider synthetic Internet topologies and evaluate the number of nodes that can be switched off by using a heuristic to get the Steiner tree. Finally, we evaluate the impact of the current Internet design, in which the global topology is partitioned into subgraphs (i.e., Autonomous Systems - AS) that implement autonomous decisions, therefore limiting the possibility of turning off the devices. We compare results in a possible future scenario in which ASs cooperate with the aim of reducing energy waste.

The paper is organized as follows. Sec. II describes the main graph models used to represent Internet. Sec. III presents our theoretical models. Simulation results are shown in Sec. IV. Finally conclusions are drawn in Sec. V.

II. INTERNET TOPOLOGY MODELS

The Internet is a complex, distributed and evolving system: understanding, measuring and modeling it is a complex task. In this paper, we are interested in evaluating the amount of

resources (in terms of nodes and links) that are redundant and thus can be switched off to reduce the energy consumption ¹

The Internet topology is typically modeled as an undirected graph $G(v, e)$, being v the set of vertices (i.e., nodes) and e the set of edges (i.e., bidirectional links). Which graph better models the current Internet graph is still a matter of discussion. Indeed, knowing the actual topology is almost impossible, since, on the one hand, ISPs are not willing to share their actual topology, and, on the other hand, the size of the Internet is so large that it is impossible to experimentally infer it. Nonetheless, in the literature several graph models have been proposed, all based on the idea of random graphs, i.e., a random process that generates a graph with known properties. Initially, simple random graphs have been proposed, such as the Erdős and Rényi [5] model, in which nodes are connected by links according to a given probability. Unfortunately, such simple graphs do not match properties that have been actually observed in the Internet, such as the “small-world” property, according to which even if degree (i.e. the number of edges per node) of nodes is rather limited, the diameter of the graph (i.e. the maximum distance in terms of hops between nodes) is very small. Moreover, the degree distribution of nodes $P(k)$ is known to follow a power-law distribution, i.e., $P(k) \sim k^{-\gamma}$. Therefore, more complex random graph models have been introduced, among which Barábasi and Albert (BA) [6] is generally accepted as a good (and simple) model. Indeed, the BA model both matches the small-world property, and the power-law distribution of edges experienced in the actual Internet Topology. The BA model builds a graph by iteratively adding a new node to the already existing graph. Each new node has L edges, that are randomly placed to connect to already existing nodes. The probability of selecting a node can follow a “preferential attachment paradigm”, so that nodes with larger degree (large number of edges) are selected with higher probability than nodes with smaller degree. The intuition suggests that nodes that are added earlier to the topology will have a higher probability to become “hubs”, while nodes that are added later will have fewer edges, being selected with smaller probability and by fewer nodes.

In this paper, we consider the BA model to study the probability that a node can be removed from a graph without producing a disconnected graph, i.e., the probability that a node is redundant.

Besides random graph models which are suitable for analytical evaluation of Internet topology properties, several random topology generators have been proposed in the literature. They generate a synthetic topology starting from a (possibly more complex) random graph model. For example, hierarchical models that better mimic the actual Internet routing policies can be modeled, as in Brite tool [7]; in particular, a top-down approach is adopted: first links between different Autonomous Systems (AS) are placed according to a simple Erdős and Rényi model, then routers in the same AS are interconnected

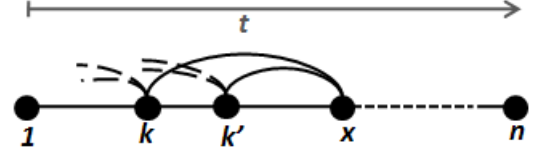


Fig. 1. Example of essential node: node k is essential for node x , while node k' can be removed without eventually disconnecting node x .

using a BA model. Other approaches are based on actual Internet topology, that are used to generate different size graphs which show the same properties (edge distribution, path length, etc.) of the sample topology given as reference. For example, topologies can be generated from the well known Hot [9] and Skitter [10] topologies, scaled with the tool Orbis [11], [12].

III. THEORETICAL MODELS

We consider a BA graph $G(v, e)$, $|v|$ is the cardinality of v , i.e., the number of vertices. According to the BA model, vertices are sequentially added to the graph, so that at step x , vertex x is added, and L edges connect x to L randomly selected vertices from the set $\{1, 2, \dots, x-1\}$ of vertices already in the graph. The average vertex degree is therefore $2L$. Fig. 1 shows vertex x that randomly selects $L = 2$ vertices k' and k at time x .

First, we focus on the one-connectivity problem. Observe that by construction at every step x , the graph comprising the first x vertices is connected; in particular, for $L > 1$ it is L connected (i.e., there exist at least L disjoint paths between any two vertices). This implies that when a vertex k is removed from $G(v, e)$, only the vertices $x > k$ can be potentially disconnected from the *principal component* of the remaining graph, which in turn comprises all vertices $y < k$. Furthermore, a vertex $x > k$ having an edge pointing to k can be disconnected from the principal component for effect of k removal only if none of the edges of x is pointing to any vertex $y < k$. In this latter case, k is declared *essential* for x . Conversely, a vertex k , which is non essential for every vertices $x > k$, is declared *non-essential*. As a consequence, the following theorem holds:

Theorem 3.1: Given a BA graph $G(v, e)$, all the non-essential vertices can be removed from $G(v, e)$ without disconnecting the graph.

Proof: We sequentially scan all the non essential vertices in graph $G(v, e)$ and remove them, following a reverse order. Let $G_k^+(v, e)$ be the graph obtained after the removal of all the non essential nodes $x > k$. Let $G_k^-(v, e)$ be the graph after the removal of non essential node k . We claim that $G_k^-(v, e)$ is connected, provided if $G_k^+(v, e)$ is connected too.

Indeed, observe that the non-essential vertex k can be removed from G_k^+ without disconnecting the remaining graph, since every other vertex $x > k$ is directly (through an edge) or indirectly (through a path) connected to a vertex $y < k$. This can be easily seen by contradiction. Assume there is a vertex

¹We assume that the rerouting of traffic on the devices that are left on induces a negligible increment of power consumption.

$x > k$ that is disconnected from the principal component for effect of the removal of k , i.e., it has no path connecting it to vertices $y < k$. Since the node x was connected before k removal, necessarily, there was a path from x to a vertex $y < k$ passing through k . Let x_1 be the vertex that precedes k along the path $x \rightarrow y$. By construction, x_1 is a neighbor of k . If $x_1 < k$, then x_1 is part of the principal component and we are in contradiction. If $x_1 > k$, since k is not essential for x_1 , x_1 must have an edge leading to a vertex $y_1 < k$, and then we are in contradiction.

At last, by induction over the nodes that are removed, the assertion immediately follows. ■

Note that Theorem 3.1 expresses a necessary condition only, i.e., a vertex k that is essential for x can be removed without necessarily resulting into the disconnection of x .

Proposition 3.2: Given a BA graph $G(v, e)$, some essential vertex in v can be removed without necessarily causing G to be disconnected.

We denote with $Pr(k, x)$ the probability that k is essential to x , i.e., vertex x has an edge pointing to k and no other edges pointing to vertices $y < k$.

At last, observe that the events i) k is essential for $x_1 > k$ and ii) k is essential for $x_2 > k$ are independent. Thus we can easily compute the probability $Pr_{off}(k)$ that node k is non-essential as function of $Pr(k, x)$ according to:

$$Pr_{off}(k) = \prod_{x=k+1}^n [1 - Pr(k, x)] \quad (1)$$

Due to Proposition 3.2, $Pr_{off}(k)$ provides a *lower bound* to the probability that k can be removed without disconnecting the remaining graph.

In conclusion, recalling that nodes can be removed only if they are non-terminal vertices, the average fraction of nodes of the graph that are jointly non-terminal and non-essential F_{off} provides a lower bound to the fraction of nodes that can be potentially switched-off without disconnecting the graph. It follows that:

$$F_{off} = \sum_{k=1}^n (1 - P_t(k)) Pr_{off}(k) \quad (2)$$

being $P_t(k)$ the probability that node k is a terminal node.

Since we are interested in evaluating the asymptotic probability to remove vertices, we consider an infinite graph by computing the limit for $n \rightarrow \infty$, so that vertices k can be mapped to a unitary segment space by defining $\alpha = k/n \in [0, 1]$. Then, we can approximate Pr_{off} as:

$$F_{off} \simeq \int_0^1 (1 - P_t(\alpha)) Pr_{off}(\alpha) d\alpha \quad (3)$$

In the following we compute Pr_{off} considering different cases. For the sake of simplicity, we assume that the probability of being a terminal node is the same, so that $P_t(k) = P_t$, $\forall k$. Then

$$F_{off} \simeq (1 - P_t) \int_0^1 Pr_{off}(\alpha) d\alpha. \quad (4)$$

In the remaining of this Section, we specify how $Pr(k, x)$ can be evaluated both for the Uniform Attachment and Preferential Attachment models.

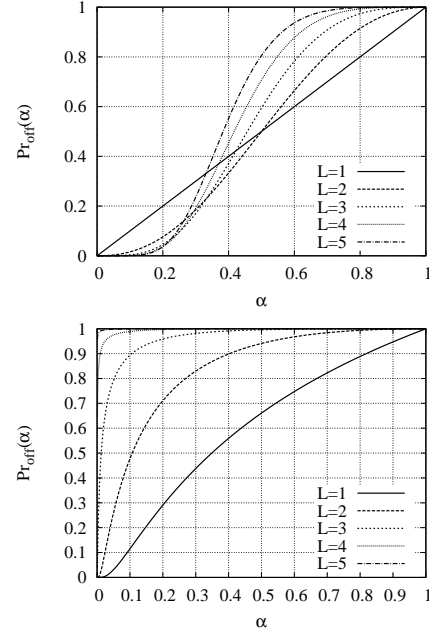


Fig. 2. One-connected graph: probability to remove a vertex in the Uniform Attachment model (top plot) and Preferential Attachment model (bottom plot)

A. Uniform Attachment model

Let us consider first a BA graph obtained using Uniform Attachment (UA) paradigm, according to which each new vertex is connected with equal probability to the vertices already present in the graph. A new vertex x is connected to a vertex $k < x$ with probability: $Pr_{att}(x, k) = L/(x-1) \quad \forall k \in [1, x-1]$. The probability $Pr(x, k)$ that k is essential for x is then given by:

$$Pr(k, x) = L \frac{1}{x-1} \prod_{i=1}^{L-1} \left(\frac{x-k-i}{x-i} \right) \simeq L \frac{1}{x} \left(\frac{x-k}{x} \right)^{L-1} \quad (5)$$

for $x \geq k + L$, while it is null for $x < k + L$.

From Eq.(1) and some approximations reported in [14], we derive the final expression of the probability:

$$Pr_{off}(k) \geq \alpha^\beta \quad (6)$$

with $\beta = L(1 - \alpha)^{L-1}$, and $\alpha = k/n$.

B. Preferential Attachment model

In the Preferential Attachment (PA) model, new vertices connect preferentially to highly connected vertices so that the probability for vertex x to select vertex k is proportional to vertex k degree. Being $K_k(x)$ the degree distribution of vertex k at time vertex x , then it holds (see [8] for details):

$$K_k(x) \simeq L \sqrt{\frac{x}{k}}$$

we can derive the probability that vertex x selects vertex k as:

$$Pr_{att}(x, k) = L \frac{K_k(x)}{\sum_{i=1}^{(x-1)} K_i(x)}$$

for $x \geq k + L$, while it is null for $x < k + L$.

Therefore, the probability that the vertex k is essential for a vertex x can be approximated by:

$$Pr(k, x) \simeq Pr_{att}(x, k) \left[\frac{\sum_{i=k+1}^{x-1} K_i(x)}{\sum_{i=1}^{x-1} K_i(x)} \right]^{L-1} \quad (7)$$

After some approximations, from Eq.(1) and Eq.(7) we derive the final expression of the average probability of removing vertex k :

$$Pr_{off}(k) \geq \exp \left[\left(\frac{1-\sqrt{\alpha}}{L} \right)^{L-1} \left(1 - \frac{1}{\sqrt{\alpha}} \right) \right]$$

C. Two-connected graph

The previous arguments can be generalized to evaluate the average fraction of vertices that can be removed still guaranteeing that the remaining graph is two-connected. In this case a vertex k is declared essential for a vertex x if the x has less than two edges pointing to vertices $y < k$, and all the non-essential vertices can be removed maintaining the two-connected properties.

1) *Uniform Attachment model:* The probability $Pr(x, k)$ that k is essential for x is similar to (5), with a further term due to the presence of the second path:

$$Pr(k, x) = L \frac{1}{x} \left(\frac{x-k}{x} \right)^{L-1} - 2 \binom{L}{2} \frac{1}{x} \frac{k}{x} \left(\frac{x-k}{x} \right)^{L-2} \quad (8)$$

The probability to switch-off vertex k becomes:

$$Pr_{off}(k) \geq \alpha^\beta \quad (9)$$

with $\beta = (1 - \alpha)^{L-1} \left[L + \frac{L! \alpha}{(L-2)!(1-\alpha)} \right]$.

2) *Preferential Attachment model:* We follow the same procedure of Sec.III-B. The probability that vertex k is essential for vertex x is:

$$Pr(k, x) \simeq L \frac{K_k(x)}{\sum_{i=1}^x K_i(x)} \left[\frac{\sum_{i=k+1}^x K_i(x)}{\sum_{i=1}^x K_i(x)} \right]^{L-1} - \frac{L!}{(L-2)!} \frac{K_k(x)}{\sum_{i=1}^x K_i(x)} \frac{\sum_{i=1}^k K_i(x)}{\sum_{i=1}^x K_i(x)} \left[\frac{\sum_{i=k+1}^x K_i(x)}{\sum_{i=1}^x K_i(x)} \right]^{L-2} \quad (10)$$

From Eq.(1) and Eq.(10), we obtain the final expression of the average probability of removing a vertex:

$$Pr_{off}(k) \geq \exp \left[\left(\frac{1-\sqrt{\alpha}}{L} \right)^{L-1} \left(1 - \frac{1}{\sqrt{\alpha}} \right) \left(1 + \frac{\alpha L!}{L(1-\sqrt{\alpha})(L-2)!} \right) \right] \quad (11)$$

D. Results

Top plot of Fig. 2 reports, for different values of L , the probability of removing vertex $k = n\alpha$ in a BA graph generated according to the Uniform Attachment model. Considering the case $L = 1$, the figure shows that a vertex $k = n\alpha$ can be removed with a probability proportional to the time at which it joined the graph. Indeed, since only one edge ($L = 1$) is used, a vertex k that is selected by any vertex $x > k$ is clearly an essential vertex. Therefore, the probability of

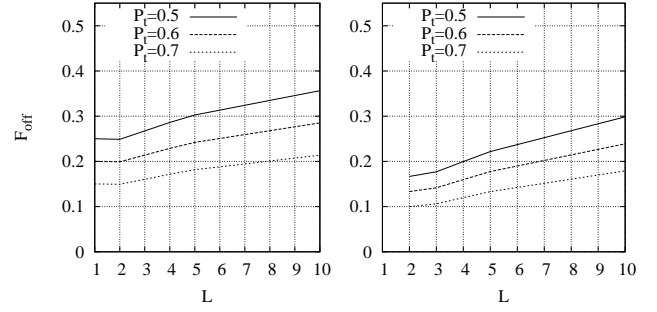


Fig. 3. Uniform Attachment model: F_{off} in the One-connected graph (left) and Two-connected graph (right).

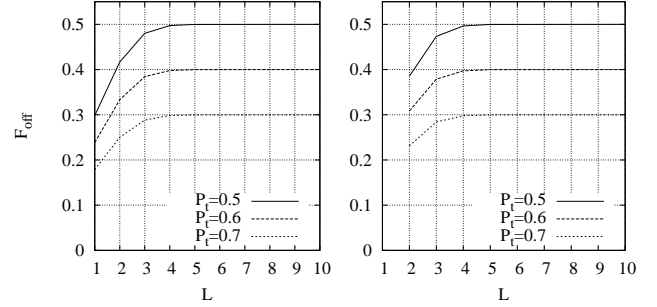


Fig. 4. Preferential Attachment model: F_{off} in the One-connected graph (left) and Two-connected graph (right).

removing a vertex is proportional to the number of times other vertices select it, i.e., to the time the vertex is added to the graph. Eq. (6) becomes then simply $Pr_{off}(k) \geq k/n$. When L increases, the probability of removing vertex k becomes smaller for vertices that join the graph early on (small α), while it increases for vertices that are added later. This is due to the fact that early vertices become “hubs” that guarantee connectivity for larger number of vertices (therefore allowing to remove late vertices with higher probability). Considering the Preferential Attachment case (bottom plot of Fig. 2), the bias induced by the preferential selection of hubs is even more evident, so that also in the case $L = 1$ early vertices have higher probability of being essential, while late vertices can be removed with higher probability.

Fig. 3 shows the average probability of removing a vertex in the Uniform Attachment model. Left plot shows UA model considering a one-connected graph, while two-connected graph results are reported in the right plot. Different values of P_t are reported, specifically $P_t \in \{0.5, 0.6, 0.7\}$. Fig. 4 shows the average probability of removing a vertex in the Preferential Attachment model: the one-connected (left plot) and two-connected cases (right plot) are reported as well. Results show that the PA model allows to easily remove a lot of vertices, so that for $L \geq 4$ practically all non-terminal nodes can be removed still guaranteeing one- and two-connected properties. Since estimates of average node degree in the actual Internet show that $L \in [2, 3]$ [8], results show that the number of redundant nodes that can be removed can be quite large. For

example, when $L = 3$ and $P_t = 0.6$, about 38% of nodes can be removed, i.e., 95% of non-terminal nodes are unnecessary. Even considering the Uniform Attachment model, a large portion of the nodes can be removed still guaranteeing the connectivity constraints.

IV. SIMULATION RESULTS

In this section, we consider more complex and realistic graphs and evaluate the actual minimum number of resources that guarantee any terminal node to connect to any other terminal node. We consider both flat (one-level) and hierarchical (two-levels) Internet topologies.

Finding the minimum set of nodes and links in a graph that is strictly necessary to guarantee the connectivity constraint among the terminals is equivalent to compute the Steiner Tree. This problem is known to be NP-hard, and in the last years different algorithms have been proposed that give approximated solutions. We choose the Selective Closest Terminal First algorithm (SCTF) [13], because the accuracy in the solution and the computational cost can be tuned as input parameter. In particular, the SCTF algorithm builds the Steiner Tree by selecting the minimum shortest path between the set of terminals T and a set of nodes in the Steiner tree set S . At step i , one terminal node $t_i \in T$ is selected, and k minimum-cost paths from t_i to $s_j \in S$, $j = 1, \dots, k$ are evaluated. Then, the minimum-cost path is selected, whose destination is s_j^* . Vertices and links from t_i to s_j^* are added to S , and t_i is removed from T . The algorithm ends when T is empty.

A. One-level Topologies

The Hot and Skitter topologies are well-known topologies often used as benchmark dataset in the Internet design field. To evaluate the impact of the number of nodes, we artificially “scale” both the Hot and Skitter samples using the Orbis tool, which allows to scale a graph, while keeping the same macroscopic features. Finally, given a topology, terminal nodes must be selected. For the sake of simplicity, we select terminal nodes according to a constant probability, i.e., P_t is constant.

Fig. 5 shows the percentage η_N of nodes that can be removed considering different network sizes and different values of P_t . The Hot and Skitter topologies are considered in the top and bottom plots respectively². Values are averaged over 5 different topologies, and maximum and minimum values are reported as well. The upper bound represents the percentage of non-terminal nodes. Notice that in all cases, the percentage of non-terminal nodes that are part of the Steiner tree is very limited, so that more than 80% of non-terminal nodes can be easily removed. In the Hot case, it is interesting to observe that the results presented are similar to the ones obtained by the BA graph with Preferential Attachment model (Fig. 4) considering

²Being $|e| = L|v|$ the number of edges in the initial graph $G(v, e)$, the percentage of links η_L that are not part of the final solution in any Steiner tree comprising $|S|$ nodes, are

$$\eta_L = \frac{|e| - (|S| - 1)}{|e|} \approx 1 - \frac{|S|}{L|v|} = 1 - \frac{1 - \eta_N}{L} \quad (12)$$

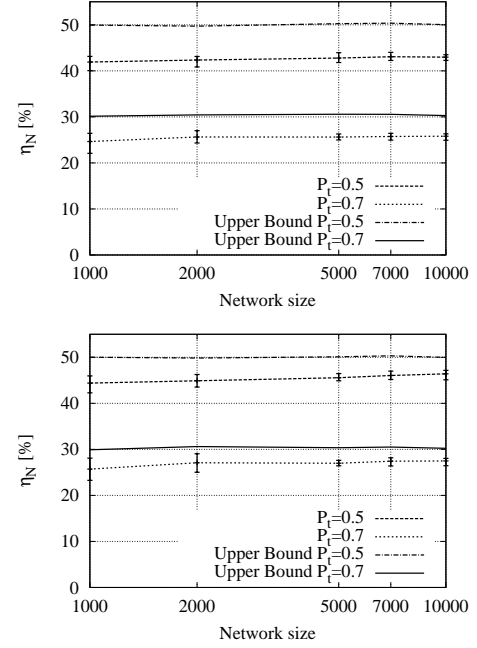


Fig. 5. Percentage of nodes that can be removed versus the network size. Hot (top) and Skitter (bottom) topologies.

the corresponding degree ($L = 2.2$) for this topology. Also, η_N increases slightly with the number of nodes, hinting that the border effects marginally impact the solution. Considering the Skitter topology (bottom plot), the higher average node degree ($L = 5.8$) guarantees to remove more non-terminal nodes, as predicted by the analytical results. This confirms that the analytical results can accurately predict the number of useless nodes in actual Internet topologies.

B. Two-level Topologies

In this Section we consider the two-level topologies in which Autonomous Systems are interconnected by a Tier-1 topology. The Brite tool is used, with the node degree parameter equal to 2 for interconnecting both intra-AS and inter-AS nodes. Moreover, the ASs are randomly interconnected using a Erdős and Rényi model, while the BA model with Preferential Attachment is used for intra-AS topology generation.

Beside considering more complex graphs, our aim is to compare the possible savings between i) the *current* Internet, where the decision in turning off devices is operated by each AS independently from the others, and ii) a *cooperative* Internet, where the ASs cooperate to minimize the global power consumption. In the first case, inter-AS devices (peering routers) cannot be removed (even if those are non-terminal nodes by definition). In the latter case this is possible, e.g., in any given ASs, one out of two peering routers can be potentially removed.

In order to test the effectiveness of cooperation, we use networks with 8192 nodes in total, while the number of Autonomous Systems N_{AS} varies between 4096 and 16, e.g., a scenario with many small ASs, to a scenario with few large

ASs. Terminal nodes are then uniformly selected from intra-AS nodes only, and the SCTF algorithm is then run to obtain the Steiner tree in each AS, so that the final topology results as the union of N_{AS} Steiner trees. Inter-AS nodes and links are always present in the final graph.

Fig. 6 (top plot) shows the percentage of nodes $\eta_N(N_{AS})$ that are removed versus the number of AS in the network, N_{AS} . The plot clearly shows that the non-cooperative approach imposed by the hard partitioning of the graph into independent ASs results in a large inefficiency of the final solution. The possible saving, indeed, decreases rapidly for large number of ASs, since a large number of devices has to be powered on to guarantee the inter-AS connectivity. On the contrary, for topologies with small number of ASs, the impact of the partitioning is marginal, suggesting that it is worth investigating a cooperative and global approach.

To dig further into the impact of cooperativeness, we compare results obtained considering the previous cases with the global Steiner tree considering the topology not partitioned into several ASs, i.e., a scenario in which a global algorithm is envisioned to achieve resource saving. Let

$$\delta_N(N_{AS}) = \eta_N(N_{AS}) - \eta_N(1)$$

be the efficiency loss in terms of nodes respectively when considering a network that is partitioned into k non-cooperative ASs versus a single ($k = 1$) cooperative network. Fig. 6 (bottom plot) confirms the previous intuition, showing that a large waste is achieved if the number of ASs increases (and therefore the number of nodes in each AS decreases). Indeed, in a cooperative Internet large savings are possible for N_{AS} up to 256 ASs, with the major loss due to devices that cannot be removed from the graph, i.e., inter-AS devices.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we faced the study of the amount of resources that can be eventually removed from the Internet topology still guaranteeing connectivity among terminals. The aim is to study the eventual amount of energy saving obtained once nodes and links in the Internet can be selectively turned-off. Results, obtained with both analytical and simulation methodologies, show that there is potential room to investigate further whether in the current and future Internet it is possible to reduce the power consumption by turning off devices that are not necessary, e.g., during off-peak periods. Both the small-world and power-law distribution of nodes degree that are observed in the actual Internet help to keep the topology connected, allowing to remove up to 80% of transport nodes.

Finally, considering the current rigid partition of the Internet topology into several and non-cooperative Autonomous Systems, we show that much higher savings are achievable in a possible future Internet, in which ASs form a cooperative network globally targeting energy saving.

Our work is somehow preliminary, since a careful evaluation of the network devices that can be switched off cannot ignore

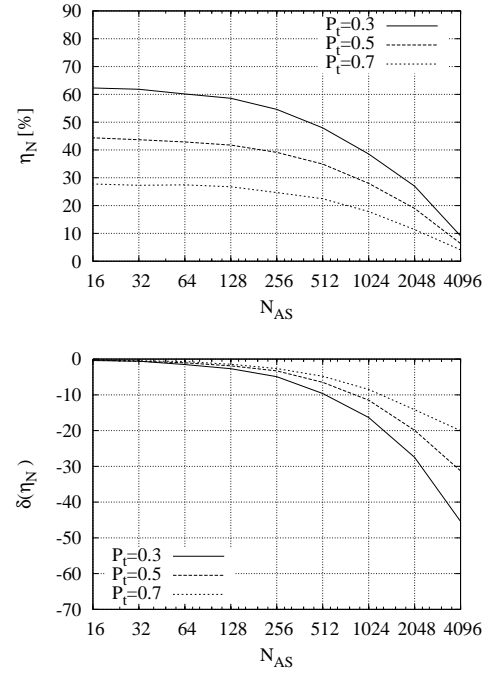


Fig. 6. Percentage of nodes (top plot) and efficiency loss for the number of nodes (bottom plot) that can be removed versus the number ASs.

the traffic flowing in the network, the protocol and device support for remote power management, etc. Nonetheless, results are encouraging for future investigation in this field.

ACKNOWLEDGMENT

The work described in this paper was performed with the support of the BONE project (Building the Future Optical Network in Europe), funded by the EU through the FP7 call.

REFERENCES

- [1] SMART 2020 Report, <http://www.theclimategroup.org>, 2008.
- [2] M. Gupta, S. Singh, *Greening of the Internet*, Proceedings of ACM SIGCOMM, Karlsruhe, Germany, August 2003.
- [3] L. Chiaraviglio, M. Mellia, F. Neri, *Reducing Power Consumption in Backbone Networks*, IEEE ICC, Dresden, DE, June 2009.
- [4] F. K. Hwang, D. S. Richards, "Steiner Tree Problems", Networks, Vol.22, pp.55-89, 1992.
- [5] P. Erdős and A. Rényi, "On random graphs" Publicationes Mathematicae,6:290-297, 1959.
- [6] A.-L. Barabási, R. Albert and H. Jeong, "Mean field theory for scale-free random networks", Physica A, 272 (1999), pp. 173-187.
- [7] <http://www.cs.bu.edu/brite/>.
- [8] R. Pastor-Satorras, A. Vespignani "Evolution and structure of the internet: A statistical physics approach", Cambridge University Press, 2004.
- [9] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the Internets router-level topology, SIGCOMM, 2004.
- [10] <http://www.caida.org/tools/measurement/skitter/>.
- [11] <http://orbis.ucsd.edu/>.
- [12] P. Mahadevan, C. Hubble, D. Krioukov, B. Huffaker, A. Vahdat, Orbis: rescaling degree correlations to generate annotated internet topologies, SIGCOMM 2007.
- [13] S. Ramanathan, "Multicast tree generation in networks with asymmetric links", IEEE/ACM Transactions on Networking, 1996.
- [14] L. Chiaraviglio, D. Ciullo, E. Leonardi, M. Mellia, "How much can the Internet be greened?", Technical Report No.290609Polito, <http://www.telematica.polito.it/chiaraviglio/papers/TRG2.pdf>