Strategic Behavior and the Environment, 2014, 4: 131–153

The Green-Game: Accounting for Device Criticality in Resource Consolidation for Backbone IP Networks

Aruna Prem Bianzino[1], Jean-Louis Rougier[2], Claude Chaudet[2], and Dario Rossi[2]

[1]*Department of Computer and Information Science, Linköping University, Linköping, Sweden; aruna.bianzino@liu.se*
[2]*Institut Télécom/Télécom ParisTech, Paris, France*

ABSTRACT

The reduction of energy expenditure of communication networks represents a key issue for the research community. A promising technique acting in this direction is known as "resource consolidation". It consists in concentrating the workload of an infrastructure on a reduced set of devices, while switching off the others. Deciding on the set of devices that can be safely switched off requires an accurate evaluation of their criticality in the network. We define here a measure of criticality that takes into account not only the network topology, but also the traffic, and different possible network configurations. We model the scenario as a coalitional game. Shapley value ranking is efficiently used to drive the resource consolidation procedure. Numerical results, on real network scenarios, confirm the robustness and relevance of the proposed index in measuring criticality, yielding a good tradeoff between energy efficiency and network robustness, and outperforming other classical indexes.

Keywords: Coalitional games; Shapley value; backbone IP networks.

ISSN 1944-012X; DOI 10.1561/102.00000038

Introduction

The carbon footprint of Information and Communication Technologies (ICT) represents today up to 10% of the global CO_2 emissions, according do different estimations, and this contribution is constantly increasing (Webb, 2008; Global Action Plan, 2007). Among the main ICT sectors, 37% of the total emissions are due to telecommunication infrastructures and their devices, while data centers and user devices are responsible for the remaining part (Webb, 2008). In such scenario, it is not surprising that the research community, the manufacturers and the network providers are spending significant efforts to reduce the power consumption of ICT systems from different angles.

To this extent, networking devices waste a considerable amount of power. In particular, energy consumption has always been increased in the last years, coupled with the increase of the offered performance (Bolla *et al.*, 2010). Actually, power consumption of networking devices scales with the installed capacity, rather than the current load (Adelin *et al.*, 2010). Thus, for an Internet Service Provider (ISP) the network power consumption is practically constant, irrespective of traffic fluctuations, since all devices consume always the same amount of power. In turn, devices are underutilized, especially during off-peak hours when traffic is low. This represents a clear opportunity for saving energy, since many resources (i.e., routers and links) are powered on without being fully utilized, while a carefully selected subset of them can be switched off or put into sleep mode without affecting the level of Quality of Service (QoS) offered by the network.

Different approaches have been proposed in the past to reduce the level of energy *consumption*, by pushing it as proportional as possible to the capacity *offered* by the network [see Bolla *et al.* (2010) and Bianzino *et al.* (2012a) for an overview]. Among the proposed approaches, the ones referred to as ‘‘Energy-Aware Routing’’ are of particular interest, as they exploit a reasonable amount of collaboration and of knowledge sharing among network devices, so that higher energy savings can be achieved.

In this paper, we face the problem of reducing power consumption in backbone networks adopting the energy-aware routing approach. In particular, we discuss an optimal formulation of the problem, and an enhanced solution, keeping into account the criticality of devices in the considered network scenario. The proposed criticality metric allows accounting for (i) the network topology, and the importance of the devices in keeping the network connectivity, (ii) the amount of traffic that the devices are routing, and (iii) different network configurations, corresponding to devices being turned off to save energy, or, equivalently, failing. Game theory represents a powerful tool to define such criticality index, in particular, we model the energy-aware routing problem as a cooperative Transferable Utility Game (TU-Game), and use the *Shapley value* of each node (player) to rate its criticality in the specific network scenario. The defined game is referred to as Green-Game, or G-Game for short, from now on. Finally, we evaluate the proposed solution on a real network scenario, contrasting it with classical criticality metrics, showing as our solution outperforms them, being able to achieve a better tradeoff between energy saving and QoS.

This approach has been previously presented in Bianzino *et al.* (2011). With respect to Bianzino *et al.* (2011), where our main interest was on the practical exploitation of the results from a networking standpoint, this extended version provides (i) a more detailed formulation of the G-Game problem from a theoretical point of view, as well as (ii) a discussion of other possible extensions of the model, and applications to other problems.

The remaining of the paper is organized as follows. The following section discusses the formulation of the energy-aware routing problem as an optimization problem, followed by the discussion of the definition of the device criticality, and the Green-Game formulation. Then, the solution complexity, and techniques to efficiently solve the Green-Game, in case of complex network scenarios are analyzed. In the next section, we present results for the Green-Game, when applied to a real network scenario. Finally, we discuss possible extensions and other applications of the G-Game, and conclude the paper.

The Optimization Problem Formulation

Once a network has been designed (i.e., the resources that will compose it have been deployed), a periodical off-line process is applied to optimize the utilization of resources, which is usually referred to as ‘‘routing optimization’’. This classical process consists, in particular, in determining the paths used for each origin–destination pair or, equivalently, to ingress–egress routers in a transit network. A classical optimization objective for Internet service providers is to avoid congestion by, e.g., balancing the traffic as evenly as possible on the network links, or by ensuring that maximum link utilization always remains below a pre-defined threshold. In pure IP networks, the path used by each flow is determined by the Internal Gateway Protocol (IGP), based on link administrative weights. Network dimensioning is thus handled by careful weight assignments, for instance using IGP Weight Optimization (IGP-WO) algorithms (Fortz and Thorup, 2000).

Routing optimization may be effectively used to reduce the energy consumption due to underutilized devices in networks at a given time. This solution shall of course preserve connectivity and QoS, e.g., by limiting the maximum utilization over any link. In other words, the required level of performance will still be guaranteed, but using an amount of resources that is dimensioned over the actual traffic demand, rather than for the peak demand. Flows may be aggregated on a reduced set of links, for example, through a proper configuration of the routing weights in an IP network.

The problem of resource consolidation may be formalized as an optimization problem, where the objective is the minimization of the total network power consumption, and constraints include the classical connectivity constraints, and QoS constraints. In more detail, we represent the network as a directed graph, $G = (N, E)$. N is the set of vertices, whose elements, $i \in N$, represent the interconnection nodes (routers, switches, etc.). E is the set of edges, whose elements $e = \{i, j\} \in E$, represent the communication links existing between pairs of nodes $i, j \in N$. We denote by n the cardinality of N (i.e., $n = |N|$). For any network element a (node or link), we will denote by l_a its load and by c_a its capacity, i.e., the maximum load it can support.

The objective is to find the network configuration (i.e., the working point of the nodes and links of the network) that minimizes the total network power consumption, expressed as the sum of the consumptions of all nodes and links. The power consumption of a network element a, node or link, can be modeled as the sum of (i) a constant part $P_{0,a}$ (corresponding to the idle power consumption), and (ii) a variable part $P_{v,a}(l_a/c_a)$, that depends on the traffic load of the device. The constant part, which is often the predominant term, is always consumed, except when the element is switched off. To include this on/off status in the model, a binary variable x_a models the status of the element a ($x_a = 1$, whenever a is *on* and $x_a = 0$, otherwise). For an arbitrary node i, the variable x_i indicates whether i is *on* or *off* and similarly the *on*/*off* status of a link (i, j) is modeled by the variable $x_{i,j}$. Finally, links are full duplex and they are considered entirely powered on as soon as one direction conveys traffic. Since in the above graph formulation the two directions are separately modeled, the link load is the sum of both directions loads. With this model, the network total power consumption may be represented by the following expression (where the first sum is divided by a factor 2 in order to avoid counting links twice):

$$\frac{1}{2}\sum_{(i,j)\in E}(P_{v,ij}((l_{ij}+l_{ji})/(c_{ij}+c_{ji}))+x_{ij}P_{0,ij})+\sum_{n\in N}(P_{v,n}(l_n/c_n)+x_nP_{0,n}). \tag{1}$$

The load imposed to the network, seen as a whole, is defined by a *traffic matrix*, $T=(t_{s,d})_{s,d\in N}$, in which an element $t_{s,d}$ represents the volume of traffic entering the network through node s and exiting through node d. Such a traffic matrix only reports traffic loads between sources and destinations, and does not report what happens inside the network. Each traffic request from s to d is routed across the network, generating a traffic of f_{ij}^{sd} over any link $\{i,j\}$. The relationship between the traffic matrix and individual links loads defines the following set of constraints on each link (ij):

$$\sum_{(i,s,d)\in N^3} f_{ij}^{sd}-\sum_{(i,s,d)\in N^3} f_{ji}^{sd}=\begin{cases}t_{s,d}, & \forall(s,d)\in N^2, j=s,\\ -t_{s,d}, & \forall(s,d)\in N^2, j=d,\\ 0, & \forall(s,d)\in N^2, j\neq s,d.\end{cases} \tag{2}$$

As mentioned above, to preserve QoS, no links should reach a 100% utilization. In the general case, the network manager may define a maximum load percentage, that we will denote by α, which he/she considers safe enough. Expressing these constraints for every link (ij) in the network defines the following set of constraints:

$$\sum_{(s,d)\in N^2} f_{ij}^{sd}=l_{ij}\leq \alpha c_{ij}, \quad \forall(i,j)\in E. \tag{3}$$

We further assume node load to be directly proportional to the amount of traffic entering and leaving the node. For this work, we consider that they are equal, even though a more complex relationship may be defined. This translates into the following constraints to our problem for every node n:

$$l_n=\sum_{(i,n)\in L} l_{in}+\sum_{(n,i)\in L} l_{ni}, \quad \forall n\in N. \tag{4}$$

Finally, we consider that a node or a link is switched off if its load is equal to zero. This allows to relate variables x_a and l_a for any element of the network through the following sets of constraints:

$$Zx_{ij}\geq l_{ij}+l_{ji}, \quad \forall i,j\in E, \tag{5}$$

$$Zx_n\geq l_n, \quad \forall n\in N, \tag{6}$$

where Z is a ''big'' number (i.e., greater than twice the maximum between the nodes and the links capacities), used to force the variable x_a to take the value 1 when a has a load greater than 0, while allowing the device a to be switched off (i.e., $x_a = 0$) only if $l_a = 0$.

Minimizing the total energy consumption (1) while satisfying all the constraints mentioned in this section is a mixed integer program, with binary variables (x_a) and continuous variables (l_a). Results for the formulation presented above have been discussed in Bianzino *et al.* (2010) and Chiaraviglio *et al.* (2011), considering a range of different power consumption models and network scenarios.

The Green-Game Definition

The main limitation of the optimization of the energy-aware routing problem, as presented in ''The Optimization Problem Formulation'' section, is the fact that the set of devices to be switched off to save energy is chosen on the basis of the sole energy costs, and does not take into account the ''criticality'' of devices in the network scenario. As discussed in Sansò and Mellah (2009), this may lead to network configurations lacking in robustness and reliability.

We believe that the resource consolidation process should be driven by an accurate evaluation of the *criticality* of devices in the specific network scenario. Some criticality metrics have been defined in the literature for network devices, in particular, the criticality of nodes in a network can be evaluated relatively easily based on the sole topology, or on the sole volume of traffic routed by each node. For what concerns topology based rankings, the most widely used ones are based either on the connectivity of each node [degree centrality (Shaw, 1954)], on the number of shortest paths passing through each node [betweenness centrality (Bavelas, 1948)], on the average distance between each couple of nodes Closeness centrality (Beauchamp, 1965), or on the importance of nodes neighbors [Eigenvector centrality (Bonacich, 1972)]. For what concerns the amount of routed traffic, the Least Flow (LF) ranking has been proposed by Chiaraviglio *et al.* (2009), which ranks devices on the basis of the amount of traffic they *would* route in an energy-agnostic configuration. However, there is no satisfactory definition of the criticality of nodes in a network, keeping into account at the same time (i) the network topology, (ii) the network traffic conditions, and (iii) alternative network configurations.

Game theory represents a powerful tool to define a criticality index that accounts for these three aspects at the same time: if we model the resource consolidation problem as a cooperative Transferable Utility Game (TU-Game), the *Shapley value* of each node indicates how much the node contributes in the traffic delivery process, and how its absence would affect the network on "average" (i.e., over all possible network configurations). This game, namely the G-Game, takes as its only inputs the *network topology*, i.e., the set of links and devices, and the *traffic matrix*, i.e., the amount of traffic routed by the network between each pair of devices. The Shapley value on the G-Game defines a joint topology-aware and traffic-aware ranking of the network devices, that can profitably be used to drive the resource consolidation process. In the remaining of this section, we will detail the definition of the G-Game.

Considering the network notation introduced in the previous section, between any two nodes $i, j \in N$, data may be transported along one or several *paths*. A path is an ordered sequence of vertices $p_{i,j} = (i = i_0, i_1, i_2, \dots, i_{k-1}, i_k = j)$. A path that does not contain twice the same node, $\forall i_a, i_b \in p_{i,j}, i_a \neq i_b$, is called an *acyclic* (or *loop-less*) path. Moreover, we denote by v the total traffic load that the network has to route, with respect to a given traffic matrix T: $v(N) = \sum_{i,j \in N} t_{i,j}$.

Let us consider an arbitrary subgraph of G, $G_S = (S, E_S)$ formed by a subset of the nodes $S \subseteq N$ and by the corresponding subset of edges $E_S = \{\{i,j\}\colon i, j \in S\} \subseteq E$. The amount of traffic that G_S can effectively transport, with respect to T, is denoted by $v(S) = \sum_{i,j \in S} t_{i,j} \mathbb{1} G_S(i,j)$, where $\mathbb{1} G_S(i,j) = 1$ whenever i and j are connected in G_S (i.e., there exist a path in G_S from i to j) and zero otherwise. By convention $v(\emptyset) = 0$.

Let us denote by $\mathcal{P}(N)$ the set of parts (i.e., subsets) of N. As N is a finite set of elements and as v is a function of $\mathcal{P}(N)$ into $\mathbb{R}$, the couple (N, v) defines a *coalitional game*, precisely the *G-Game*. A group of nodes (players), $C \subseteq N$ is called a *coalition* and the value $v(C)$ is called the *worth* (utility) of the coalition C, while v is called the *characteristic function* of the game. The problem of determining which network elements can be safely switched off without disrupting the network can be modeled as the search for a coalition with the same worth as the full network, but with a reduced size.

In other words, given any traffic matrix, we need to identify the most important nodes in the network: in case the problem is modeled as a coalitional game, the solution is represented by the Shapley value. The Shapley value averages the marginal contribution of each node over many possible scenarios, which makes it perfectly suited to find a good tradeoff between saving energy and preserving QoS.

Let us denote by Σ_N the set of permutations over N: $\Sigma_N = \{\sigma\colon N \to N : \sigma \text{ is a bijection}\}$. We also denote by $B[i,\sigma]$ the set of nodes that appear *before* node i with respect to permutation σ, including i itself: $B[i,\sigma] = \{j \in N \text{ s.t. } \sigma^{-1}(j) \leq \sigma^{-1}(i)\}$. $B(i,\sigma)$ is similarly defined as the set of nodes that appear *before* node i with respect to permutation σ, excluding i: $B(i,\sigma) = \{j \in N \text{ s.t. } \sigma^{-1}(j) < \sigma^{-1}(i)\}$. The *marginal value* of node $i \in N$, with respect to the order σ is defined as:

$$m_i^\sigma = v(B[i,\sigma]) - v(B(i,\sigma)).$$

Intuitively, the marginal value of a node according to an order represents its importance in maintaining the network performance when nodes are switched off (or fail), one by one, following the order σ. The *Shapley value* ϕ_i of node $i \in N$ is defined as the average of the marginal values associated to i for all possible permutations of N:

$$\phi_i = \sum_{\sigma \in \Sigma_N} \frac{m_i^\sigma}{n!}, \tag{7}$$

where ϕ_i defines a ranking on the nodes, which appears particularly relevant for our problem. For each node i, ϕ_i increases with the number of coalitions that i participates to and with the importance of i in each coalition. The Shapley value takes indeed into account the number of primary and backup paths each node lays on, reflecting the position of the node in the topology in a similar way to centrality measures. "Results on a Real Network Scenario" section provides also a comparison with other classical centrality indexes. Exploring every path, the Shapley value grants higher values to nodes whose removal would disconnect the graph, or to nodes belonging to small sets whose presence in the network is essential for traffic delivery. Another important advantage of this approach is that this ranking takes into account the characteristic function v, defined as the volume of traffic transported by a coalition. In other words, the higher the value ϕ_i for a device i is, the higher its contribution to traffic routing on average over all coalitions will be. For more insight on the Shapley value, we refer the interested reader to Moretti and Patrone (2008).

On the Efficient Computation of the Shapley Value

The computation of Shapley value according to Equation (7) is computationally expensive, as it requires considering all the $n!$ potential permutations of N. However, any coalitional game can be decomposed as a linear combination of *unanimity games* (Moretti and Patrone, 2008). This decomposition provides a less expensive method to calculate the Shapley value. For a set of players N, a unanimity game, (N, u_R), is defined over a subset of nodes $R \subseteq N$ by its characteristic function, u_R, which associates to any subset $C \subseteq N$ a Boolean value: $u_R(C) = 1$ if and only if $R \subseteq C$, $u_R(C) = 0$ otherwise. By convention, $u_R(\emptyset) = 0$. Any coalitional game (N, v) admits a unique decomposition in unanimity games over $\mathcal{P}(N)$:

$$v = \sum_{C \in \mathcal{P}(N)} \lambda_C \, u_C, \quad \text{with } \lambda_C(v) \in \mathbb{R}, \forall C \in \mathcal{P}(N), \tag{8}$$

where λ_C are called the Harsanyi dividends (Shapley, 1953), that are defined recursively by:

$$\lambda_C = \sum_{B \subset C} (-1)^{|C|-|B|} \, v(T). \tag{9}$$

The Shapley value of a node $i \in N$ is fully determined by these dividends, considering all the subsets of N in which i appears:

$$\phi_i = \sum_{C \in \mathcal{P}(N);\, i \in C} \frac{\lambda_C}{|C|}. \tag{10}$$

The complexity of this computation is $O(3^n)$, considering that this expression requires at most a computation of all the 2^n Harsanyi dividends. Each $\lambda_C(v)$ computation requires to enumerate all the subsets B included in C (i.e., $2^{|C|}$ sets). Ordering the sets C by increasing cardinality, we can thus see that the total complexity for computing all the dividends can be expressed as: $\sum_{k=0}^{n} \binom{n}{k} 2^k = 3^n$. Even though 3^n is asymptotically lower than $n!$, the algorithm complexity remains exponential.

Fortunately, a further simplification is introduced in Gómez *et al.* (2004): as the Shapley value reflects the importance of a node in the routing process, we do not need to consider the whole $\mathcal{P}(N)$, but only the elements that represent valid paths in which the node participates. In addition, "augmented" paths shall not be considered. Let us consider two paths, P and Q between

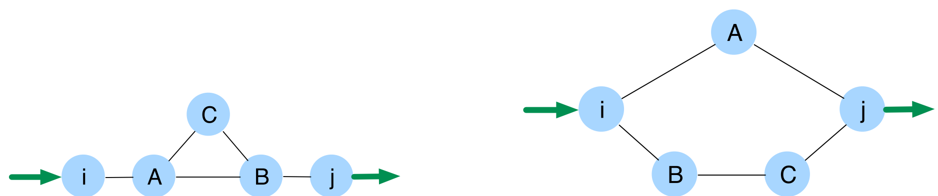

Figure 1. Toy examples illustrating the Shapley value computation. On the left graph, two acyclic paths exist between i and j, one augmenting the other. On the right graph, alternative paths exist between i and j. Dark arrows represent the traffic request from node i to node j.

i and j, such that $Q = P \cup R$. Q is an "augmented" path, since $P \subseteq Q$. For example, let us consider paths $P = (i, A, B, j)$ and $Q = (i, A, C, B, j)$ in Figure 1 (left). Nodes in $R = Q \setminus P$ (i.e., C in the example) do not provide any alternative when a node in P is switched off. Therefore, they should not increase their score for participating in path Q. Note that cyclic paths are special cases of augmented paths, meaning that only acyclic paths are of interest.

More formally, let us now denote by $\mathcal{M}_E(\{i, j\})$ the set of all acyclic paths between i and j in G, and let K_{ij} denote the cardinality of this set. For each path p, we denote by $\pi(p)$ the unordered set of nodes composing p. For instance, $\pi\left((A, B)\right) = \pi\left((B, A)\right) = \{A, B\}$. Let us also denote by $\mathcal{P}_k\left(\mathcal{M}_E(\{i, j\})\right)$ the set composed by all the combinations of the union of k paths in $\mathcal{M}_E(\{i, j\})$. Let us extend the π notation to a set of paths, by posing $\pi(p) = \pi(p_1) \cup \pi(p_2) \cup \cdots \cup \pi(p_k)$ for a path $p = \{p_1, p_2, \ldots, p_k\}$. The following expression defines the graph-restricted game (Myerson, 1977), by introducing a characteristic function for an unanimity game that removes the influence of augmented paths:

$$u_{i,j} = \sum_{k=1}^{K_{ij}} \left(\sum_{p \in \mathcal{P}_k(\mathcal{M}_E(\{i,j\}))} (-1)^{k+1} \, u_{\pi(p)} \right). \tag{11}$$

To better understand the rationale behind (11), let us consider the toy-case example of Figure 1 (left). First, the set of acyclic paths is composed of $K_{i,j} = 2$ elements: $\mathcal{M}_E(\{i, j\}) = \{p_1 = (i, A, B, j), p_2 = (i, A, C, B, j)\}$.

Applying the previous formula, we may express $u_{i,j}$ as:

$$\begin{aligned} u_{i,j} &= u_{\pi(p_1)} + u_{\pi(p_2)} - u_{\pi(\{p_1,p_2\})} \\ &= u_{\{i,A,B,j\}} + u_{\{i,A,B,C,j\}} - u_{\{i,A,B,C,j\}} \\ &= u_{\{i,A,B,j\}}. \end{aligned}$$

We may thus neglect the augmented paths and restrict our computations on the set $\mathcal{M}_E^*(\{i,j\}) = \{P \in \mathcal{M}_E(\{i,j\}) : \nexists Q \in \mathcal{M}_E(\{i,j\})\ Q \subset P\}$. For a given path $p \in \mathcal{M}_E^*(\{i,j\})$, the value of $u_{\pi(p)}$ is equal to 1 for every subset of nodes part of this path, leading to a Shapley value increase proportionally to $t_{i,j}$ and inversely proportionally to the path length. For a path p and a node h, let us define $\mathbb{1}h(p) = 1$ if node h belongs to p, and $\mathbb{1}h(p) = 0$ otherwise. Denoting by K_{ij}^* the cardinality of $\mathcal{M}_E^*\{i,j\}$, and by $\phi(i,j)$ the Shapley value of the unanimity game $u_{i,j}$, the Shapley value granted to a node h is thus $\phi_h = \sum_{i,j} \phi_h(i,j)$, with:

$$\phi_h(i,j) = t_{i,j} \sum_{k=1}^{K_{ij}^*} \left(\sum_{p \in \mathcal{P}_k(\mathcal{M}_E^*\{i,j\})} \frac{(-1)^{k+1}}{|\pi(p)|}\, \mathbb{1}h(p) \right). \tag{12}$$

For the sake of illustration, let us consider the example depicted in Figure 1 (right). If we consider that the traffic matrix only has one non-null element, say $t_{i,j} = 1$, the resulting Shapley value $\phi = (\phi_i, \phi_A, \phi_B, \phi_C, \phi_j)$ is:

$$\begin{aligned} \phi &= \left(\left(\frac{1}{3}, \frac{1}{3}, 0, 0, \frac{1}{3}\right) + \left(\frac{1}{4}, 0, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right) - \left(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}\right) \right) \\ &= \left(\frac{23}{60}, \frac{8}{60}, \frac{3}{60}, \frac{3}{60}, \frac{23}{60}\right). \end{aligned}$$

These values show that the traffic source and destination, i and j, are the most critical nodes, as their Shapley value is maximal. Then comes A, which lies on the shortest path from i to j, and finally B and C are granted the smallest values, as they represent a longer, backup path.

Practical Considerations on Shapley Value Computation

Computing the Shapley value using (12) is still computationally intensive when considering a realistic, and hence complex, network scenario, so specific heuristics are needed. First, for every non-null entry in the traffic matrix,

$t_{i,j}$, we need to find all valid (i.e., non augmented) paths from i to j. We can determine these paths using a taboo search procedure (Glover and Laguna, 1997). Taboo search explores the network similarly to a Breadth First Search (i.e., by neighbors), but with different stopping conditions. To produce valid paths, the search ignores some branches, avoiding (i) loops and (ii) augmented paths. First, when a branch $(i, i_1, i_2, \ldots, i_n)$ is explored, the already visited nodes $i, i_0, \ldots, i_{n-1}$ cannot be visited again due to the loop-less path constraint. Then, the branch should also avoid neighbors of preceding nodes, as this would otherwise lead to augmented paths.

Let us consider for instance the network represented in Figure 2. Let us consider the path from node 1 to node T and focus on the branch $(1, 8, 9)$. Node 9 has two neighbors: 21 and 10, but the exploration only needs to consider node 21. The branch $(1, 8, 9, 10)$ would actually build augmented paths, such as $(1, 8, 9, 10, 11, T)$, with respect to the branch $(1, 10)$ that would reach T with a shorter path $(1, 10, 11, T)$. Therefore, the exploration has to skip node 10, as it is already a neighbor of node 1 (i.e., a neighbor of the predecessor nodes in the branch). The taboo list is then populated by the set $\bigcup_{n=i,i_0,\ldots,i_{n-1}} \mathcal{N}_n$, where $\mathcal{N}_n$ represent the set of neighbors of a node n.

Even when only the limited set of valid paths are considered, the Shapley value computation from (12) becomes intractable as the number of paths

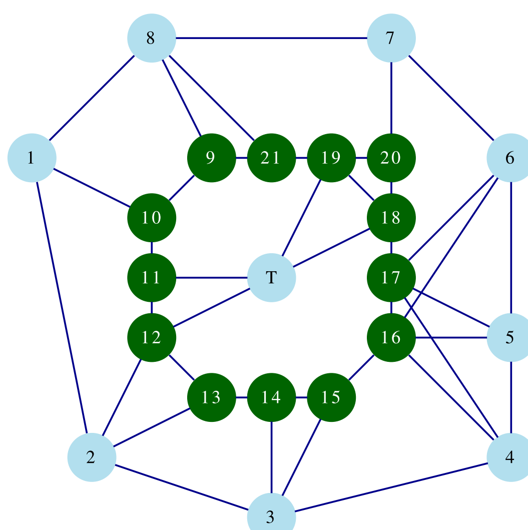

Figure 2. The reference topology.

grows: the formula requires indeed, for any non-null flow (i, j), to consider all the possible combinations of the $K_{i,j}$ paths that have been found, hence $2^{K_{i,j}}$ iterations per flow. At the same time, it is usually possible to limit the value of $K_{i,j}$ while still obtaining accurate results, by simply bounding the maximum path length L. Actually, every path brings a contribution inversely proportional to its length to the Shapley value of each traversed node (as shown, e.g., in (12)). In addition, the use of very long paths (i.e., greater than the network diameter) is rare in real networks, as they would only be used in extreme cases when multiple link/node failures occur simultaneously. Network design and routing optimization processes seldom consider such situations, as they are extremely rare. Hence, bounding the maximum path length to a value L greater than the diameter would not affect the practical relevance of the solution from a networking standpoint.

Finally, and most important, the Shapley value for the bounded and unbounded maximum path length become very close provided that the maximum path length is large enough. Figure 3 reports the G-Game Shapley value of all nodes, for growing maximum path length L, ranging from $L = 4$ hops (i.e., the diameter of the network in Figure 2), to $L = 7$ hops (nearly twice the diameter). As we can see, the difference between absolute Shapley values for $L \geq 6$ gets negligible, and, most important, the order of nodes following the Shapley rank remains the same, confirming that longer paths, contributions are marginal.

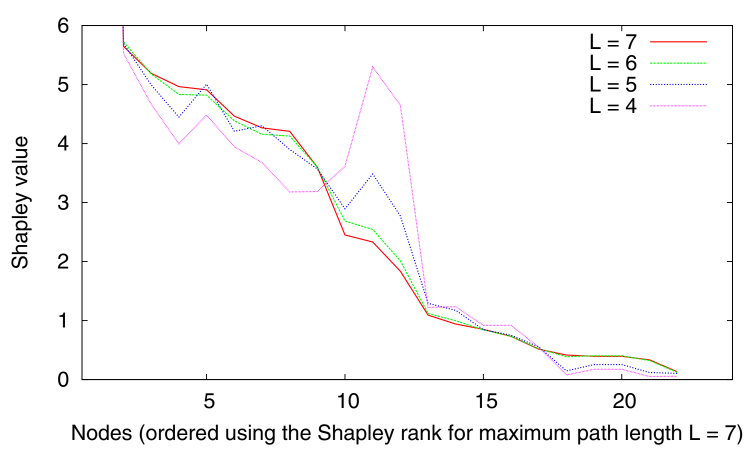

Figure 3. Node ranking for different maximum path lengths.

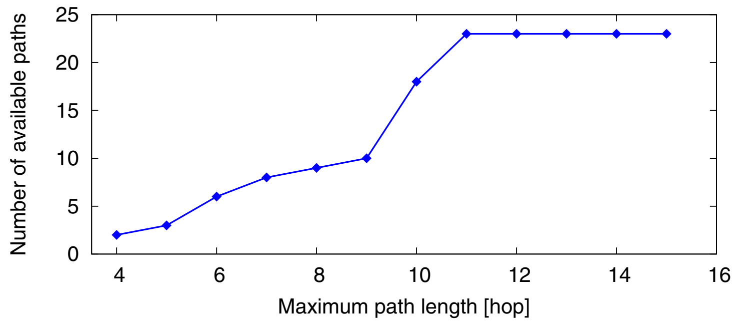

Figure 4. Number of available paths between 1 and T as a function of the maximum path length.

The path length bound L makes Shapley value computation feasible. Figure 4 represents, for example, the number of paths in the set $\mathcal{M}_E^*(\{1,T\})$ for different maximum length L for the the flow between node 1 and node T in Figure 2. As expected, the number of discovered paths increases rapidly as the maximum path length grows, until the longest acyclic path is found (after which the number of available valid paths saturate). Looking more closely however, we can notice that the path lengths greater than or equal to 6 are already very "long paths" that are unlikely to be used in an operational network (at least two nodes should be simultaneously down, e.g., 2 and 10, or more). Thus, a limit of, e.g., $L = 6$ hops would allow a reasonable number of backup paths from the network viewpoint, while at the same time limiting the number of iterations to just $2^6 = 64$ for the flow $(1, T)$.

As a conclusion, the use of taboo search and maximum path length limitation considerably reduces the Shapley value computational complexity, focusing only on paths relevant for the network operation. Moreover, Shapley value ranking for nodes is still accurate if the bound L is appropriately selected, as a function of the network diameter. In the following, we impose a maximum path length of $L = 7$ hops (approximately the double of the network diameter), which corresponds to the smallest length at which the Shapley ranking does not evolve anymore (i.e., ranking for $L = 6$ hops is identical to $L = 7$, as shown in Figure 3).

Results on a Real Network Scenario

Once a criticality ranking has been defined, it can be utilized in the energy-aware routing process to decide the order in which network devices are attempted to be switched off. As described in Chiaraviglio *et al.* (2011), in fact, a periodical off-line process may be applied to networks. Each process occurrence acts in a specific traffic condition [i.e., time period (What Europeans do at Night, 2009)], and attempts at switching off nodes, one by one, following the considered criticality ranking (i.e., from the less to the most critical). A switch off attempt consists in verifying if the network would be able to work in normal conditions without the considered node, i.e., deliver all the traffic requests, respecting the QoS constraints imposed by the network administrator.

In this section, we evaluate the tradeoff between QoS and energy saving, comparing the proposed ranking with other classical ranking schemes. To provide a relevant evaluation, we take care of building a realistic scenario. As far as the network is concerned, we consider the reference topology of an ISP participating in the TIGER2 project, and the corresponding traffic matrix. This network, depicted in Figure 2, represents a portion of the ISP access/metropolitan network segment. The light-shaded nodes (1–8) are access nodes, source and destination of traffic requests, and cannot be switched off. The dark nodes (9–21) are transit nodes, performing only traffic transport, and can be switched off. Node T is the traffic collection point, providing access to the core network and the big Internet, with whom nodes typically exchange the majority of the traffic.

We adopt the node power consumption model proposed by Tucker *et al.* (2008), widely accepted in the literature, and already used in Bianzino *et al.* (2010). The power consumption P_i (in Watts) of a node, is related to its switching capability C_i (in Mb/s) according to $P_i = C_i^{2/3}$. Since the node switching capability has not been disclosed for the considered scenario, we consider that a node is able to switch twice the capacity of its entire set of connected links. As, in this technological scenario, nodes are responsible for the majority of the network power consumption (Bianzino *et al.*, 2010), we focus on methods to switch-off *nodes* and neglect the energy that might be further saved by switching off links (i.e., network interfaces).

Considering the different criticality metrics, i.e., the classical ones discussed at the beginning of ‘‘The Green-Game Definition’’ section, and the one defined by the G-Game, we can see how node importance in the topology is evaluated in the G-Game by taking into account the number of paths a node lies on, similarly to the betweenness centrality. However, unlike betweenness centrality, the Shapley value takes into account failure scenarios by considering not only the shortest paths, but also longer paths that can provide alternate paths in degraded scenarios. Note that a scenario in which a node has been switched off to save energy, or the same node fails, are equivalent from the point of view of the routing, and of the G-Game.

All the above listed criticality indices have been evaluated on the reference network scenario. We also compare two different versions of the Shapley value: (i) a simplified index that reflects only the network topology, considering the G-Game with a uniform traffic matrix, referred to as G-Game U-TM hereafter; (ii) the full G-Game earlier defined, that considers the actual traffic matrix. Figure 5 offers a graphical representation of the difference between the G-Game and the G-Game U-TM: in this representation, the size and color of a node represent its criticality in the considered game (the bigger and the darker the node, the higher its criticality). As expected, the collection point T has the largest worth in the G-Game due the amount of traffic transiting to/from the big Internet, whereas transit nodes $i \in [9, 21]$ have a lower worth as they are interchangeable. As long as the traffic matrix is satisfied, there is no preference among transit nodes.

Recall that, to switch off nodes, we are only interested in the order of criticality among nodes, rather than in the evaluation of the precise values of node criticality. Therefore, to compare the different rankings we compute

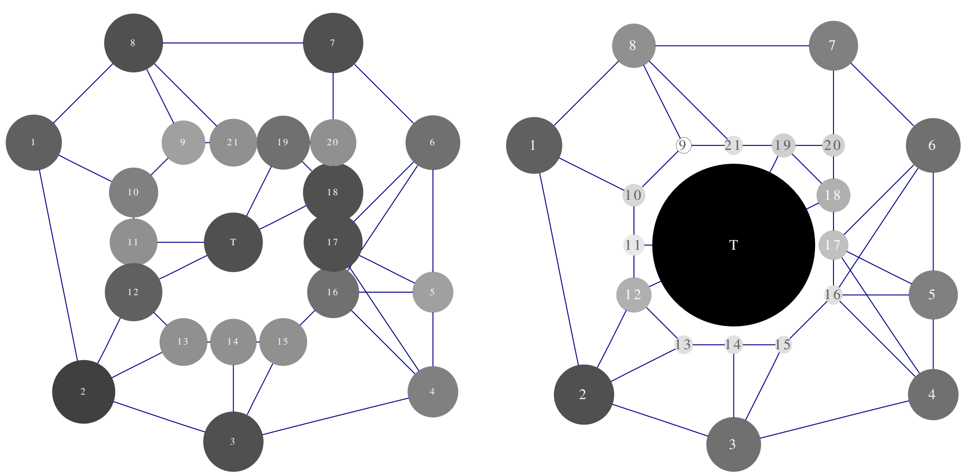

Figure 5. Node criticality when considering only the network topology in the G-Game, i.e., G-Game U-TM (left), and when considering also the real traffic matrix, i.e., the full G-Game (right).

Table 1. Correlation coefficients between the rankings defined by different criteria.

	G-Game (U-TM)	Betweenness	Degree	Closeness	Eigenvector	G-Game	LF
G-Game (U-TM)	1.00						
Betweenness	0.97	1.00					
Degree	0.46	0.53	1.00				
Closeness	0.87	0.91	0.62	1.00			
Eigenvector	-0.01	0.08	0.73	0.18	1.00		
G-Game	0.41	0.43	0.25	0.51	−0.02	1.00	
LF	0.43	0.49	0.48	0.60	0.19	0.56	1.00

the Paerson correlation coefficients between every pair of rankings. Results are summarized in Table 1, where coefficients range from −1 to 1: a value close to 1 reflects a direct correlation (i.e., same order); a value close to −1 reflects an inverse correlation (i.e., inverse orders), and a value close to 0 reflects the absence of correlation. From these results, four families of rankings appear: LF and Shapley value produce singular rankings (i.e., that are not correlated with any other). Most topology-related rankings (Betweenness, Closeness, G-Game U-TM) are similar (correlation ∼0.9) and are evaluated only through the G-Game U-TM hereafter. Degree and Eigenvalues also form a distinct family which is omitted below as resulting less pertinent, and performing poorly.

Energy Savings vs. QoS

Energy saving capability is evaluated with respect to the energy-agnostic configuration, in which all nodes are powered on (referred to as "Baseline" configuration). We focus on three different node rankings: (i) the one obtained by the full G-Game, (ii) the one obtained by the G-Game U-TM, based only on topology, and representative of the "topology-related" ranking

Table 2. List of the network nodes, ordered (left to right) from the least to the most critical one, according to different criticality rankings.

Ranking	Node ID																					
G-Game	**9**	**15**	13	14	**16**	*21*	11	10	20	*19*	17	18	12	8	5	7	4	3	6	1	2	T
G-Game (U-TM)	5	**9**	**14**	**20**	21	15	*13*	11	10	4	16	19	6	1	12	17	8	T	7	3	18	2
LF	**9**	**15**	8	7	5	4	**21**	20	2	3	1	6	14	11	10	*19*	*16*	13	12	17	18	T

Note: Underlined values nodes that can be switched off such that the network remains able to carry the traffic matrix. Bold values indicate nodes considered in the resource consolidation process.

family, and the (iii) LF ranking, based only on the traffic matrix. The resulting orders of nodes are reported in Table 2.

To evaluate the pertinence of the different rankings, we select a set of nodes that can be switched off by scanning the list sorted by increasing criticality (i.e., safest first). The algorithm examines each node in turn, by checking whether its removal, in addition to nodes previously turned off, would prevent the network from routing the whole traffic matrix (by means of a linear program). Nodes that can be switched off, for the different considered orders, are underlined in Table 2. Notice that nodes that can be switched off are less critical in the G-Game ranking with respect to the LF ranking: hence, they are found earlier during the list scan.

Indeed, the energy saving objective shall affect neither the offered QoS, nor the network robustness. Yet, the greedy switch-off approaches considered so far tend to leave little space to redundancy, and even less means to *control* the redundancy level. An alternative option to control redundancy is to stop the process when reaching a preconfigured target maximum number of switched off nodes, selected by scanning the whole list if necessary.

To evaluate the impact of this strategy on the reference network, let us fix a limit of $N_{\text{off}} = 3$ off nodes, so that at most 25% of the transit nodes can be switched off at the same time; the nodes selected by each strategy are those highlighted in boldface in Table 2. After a network configuration is selected (i.e., after N_{off} nodes are switched off), we compute the link load by routing the traffic matrix on the resulting topology: in more detail, we use TOTEM (Balon *et al.*, 2007) to perform an optimization of the routing weights (using the IGP-WO algorithm) and route the traffic enabling Equal

Table 3. Variation of the average path length (in number of hops) and achievable energy saving, considering different criticality rankings.

Ranking	Off nodes	Avg. path length $\bar{l}$	$\bar{l}_{\mathrm{TM}}$	Energy saving (%)
G-Game	9, 15, 16	2.45	2.99	17.05
G-Game (U-TM)	9, 14, 20	2.92	3.40	13.43
LF	9, 15, 21	2.64	3.25	16.27
Baseline	None	2.64	3.13	0.0

Cost Multi Path (ECMP). It follows that we are able not only to evaluate topological properties, but also to precisely measure the load on individual links. This is an important point, as the distribution of the link utilization is a very relevant Traffic Engineering (TE) indicator for carriers.

The resulting energy saving is reported in Table 3, together with the average path length $\bar{l}$; we also report a weighted average path length $\bar{l}_{\mathrm{TM}}$, where paths are weighted by the amount of traffic they transport over the traffic matrix (TM). The increase of the average path length is a logical consequence of switching off some nodes. Notice that the average path length is minimal for the G-Game, and reduces with respect to the baseline configuration. To get an intuition on how the average path length may decrease by switching off nodes, let us consider again the toy case of Figure 1 (right), and suppose for the sake of illustration that traffic shall be shared evenly on paths (i, A, j), and (i, B, C, j), resulting in $\bar{l} = \bar{l}_{\mathrm{TM}} = 2.5$ hops. The resource consolidation process may disable one of the two paths, bringing either to an *increase* (i.e., only the (i, B, C, j) path is available, $\bar{l} = \bar{l}_{\mathrm{TM}} = 3$ hops), or to a *decrease* (i.e., only the (i, A, j) path is available, $\bar{l} = \bar{l}_{\mathrm{TM}} = 2$ hops) in the average path length, depending on which nodes are switched off.

Finally, Figure 6 reports link utilization distributions for the different rankings when $N_{\mathrm{off}} = 3$ and the baseline configuration (i.e., $N_{\mathrm{off}} = 0$). Notice that the G-Game yields to excellent performances, as the link distribution is roughly equivalent to the one of baseline configuration, where no node is switched off. Especially, maximum link utilization does not increase under G-Game, with respect to the "all-on" network configuration: this means that energy saving is obtained without compromising the expected QoS.

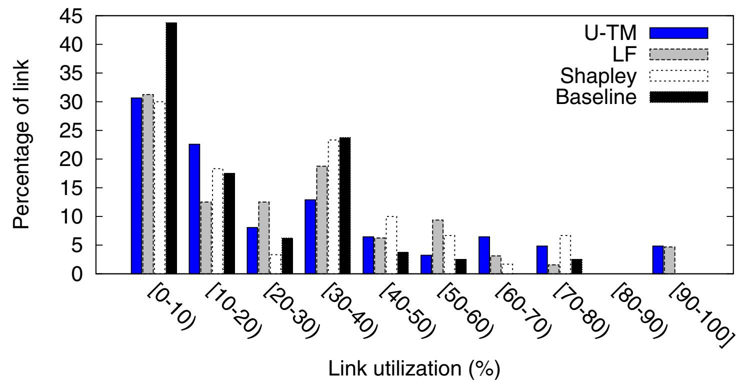

Figure 6. Distribution of the link utilization, considering different ranks and in the Baseline configuration.

Conversely, some links reach an utilization higher than 90% for the U-TM and LF strategies. The LF strategy results in worse link distribution since it passes through longer alternate paths (i.e., considers only routing paths as in the baseline configuration and ignores fault cases), while the worse QoS results of the U-TM strategy are due to its traffic unawareness (i.e., it takes into account only the topology). In contrast, G-Game explicitly considers existing paths for different node combinations, which means that it explores configurations where some nodes are excluded (i.e., which is precisely what happens when nodes are switched off in the resource consolidation process).

Possible Extensions and Other Applications

G-Game Extensions to Other Network Scenarios

Different network scenarios, involving different transmission technologies, may require different resource consolidation approaches. In particular, nowadays network routers do not allow sleeping states for the whole device, but at most for single components (e.g., ports, or linecards). Moreover, the boot-time of such device is typically non-negligible, so that it may be difficult to let *network nodes* sleep in practice.

At the same time, we point out that different transmission technologies may require power hungry devices (e.g., opto/electronics signal regenerators/amplifiers) along communication links, which may move a considerable share of the energy consumption from nodes to links. For the above reasons, we acknowledge that it may be interesting to apply a resource consolidation procedure considering *network links* rather than nodes.

Fortunately, the G-Game formulation of ‘‘The Green-Game Definition’’ section may be properly modified to account for the criticality of links rather than nodes. In this case, players would represent links, and coalitions represent set of powered-on links, which are carrying traffic requests. The rest of the definition remains valid, and the Shapley value computed solving the G-Game still represents a valuable criticality ranking for links, accounting for the network topology, the traffic conditions, and the network robustness. On the other hand, the definition of ‘‘augmented paths’’ is substantially different: a path which is not offering alternatives in case of node failures (or switching off), may instead offer alternatives in case of link failures. For the sake of example, consider again Figure 1 (left). As previously discussed, with respect to the path $\{i, A, B, j\}$, the path $\{i, A, C, B, j\}$ represents an augmented path for the node-level G-Game, as it is not offering alternatives in case of node failures. When considering links as players, instead, the path $\{i, A, C, B, j\}$ represents a valid alternative to $\{i, A, B, j\}$ in the case of link $\{A, B\}$ switch-off.

The different definition of paths reflects on the practical computation of the Shapley values, which can no longer leverage on the taboo search procedure to determine the set of non-augmented paths, but should adopt a more generic Breadth First Search. As a consequence, other heuristics should be introduced to make the computation treatable when considering large scale networks. A preliminary discussion for a formulation of the G-Game considering communication links as players are presented in Bianzino *et al.* (2012b).

G-Game Application to Non-Networking Scenarios

We believe that the G-Game, as formulated in ‘‘The Green-Game Definition’’ section, can be profitably used in the evaluation of the criticality of other scenarios, different from communication networks.

In particular, scenarios in which fault tolerance is extremely important, like power or water distribution systems, may be modeled using the G-Game, considering the connection points as players, and defining a traffic matrix on the basis of the plant (or spring) productions, and of the use-point requests.

The Shapley value of the obtained game may help in the evaluation of the criticality of specific connection/distribution points, and drive the emergency planning procedure.

Conclusions

Classical measurements for the criticality of a device in a communication network take into account either (i) the network topology (and limitedly the shortest path between node pairs) or (ii) the traffic matrix. The G-Game, that we presented in this paper, provides a powerful way to measure such criticality, jointly taking into account the traffic conditions and the network robustness (i.e., possible failure scenarios, and multiple paths between node pairs beyond the shortest one).

We have compared different criticality rankings, that we use to drive a resource consolidation process, i.e., to switch off network nodes. Numerical results on a realistic network scenario show that Shapley value ranking yields to high energy savings, with little or no impact on the expected QoS levels on the network: in particular, the maximum link utilization does not increase much on G-Game with respect to that of the full network scenario (i.e., where all nodes are active and no consolidation is in place).

On the one hand, we have shown that G-Game can be applied to problems outside the networking domain, which is an interesting direction to explore. On the other hand, the G-Game is flexible enough to solve different resource consolidation problems within the networking domain, that is needed whenever the underlying assumption change (e.g., node versus link consumption).

Finally, we believe that the criticality index for network devices, as defined by the G-Game, may be profitably used also in the network design process. In this case, G-Game may allow, for instance, to distribute the criticality as much as possible among network devices (i.e., introduce back-up devices aside the most critical ones, and remove the less critical devices, taking into account the corresponding investment).

References

Adelin, A., P. Owezarski, and T. Gayraud. 2010. "On the Impact of Monitoring Router Energy Consumption for Greening the Internet." In *IEEE/ACM International Conference on Grid Computing (Grid 2010)*, Bruxelles, Belgique.

Balon, S., J. Lepropre, O. Delcourt, F. Skivée, and G. Leduc. 2007. "Traffic Engineering an Operational Network with the TOTEM Toolbox." *IEEE Transactions on Network and Service Management* 4(1): 51–61.

Bavelas, A. 1948. "A Mathematical Model for Small Group Structures." *Human Organization* (7): 16–30.

Beauchamp, M. 1965. "An Improved Index of Centrality." *Behavioral Science* 10(2): 161–163.
Bianzino, A., C. Chaudet, F. Larroca, D. Rossi, and J. L. Rougier. 2010. "Energy-Aware Routing: A Reality Check." In *3rd International Workshop on Green Communications (GreenComm3)* Miami, USA.
Bianzino, A. P., C. Chaudet, S. Moretti, D. Rossi, and J. L. Rougier. 2011. "The Green-Game: Striking a Balance Between QoS and Energy Saving." In *23rd International Teletraffic Congress (ITC 2011)* San Francisco, USA.
Bianzino, A. P., C. Chaudet, D. Rossi, and J. L. Rougier. 2012a. "A Survey of Green Networking Research." *IEEE Communication Surveys and Tutorials* 2(2): 1–18.
Bianzino, A. P., C. Chaudet, S. Moretti, J.-L. Rougier, L. Chiaraviglio, and E. Le Rouzic. 2012b. "Enabling Sleep Mode in Backbone IP-Networks: A Criticality-Driven Trade-off." In *IEEE ICC'12 Workshop on Green Communications and Networking* Ottawa, Canada.
Bolla, R., R. Bruschi, F. Davoli, and F. Cucchietti. 2010. "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures." *IEEE Communication Surveys and Tutorials* 13(99): 1–22.
Bonacich, P. 1972. "Factoring and Weighting Approaches to Status Scores and Clique Identification." *Journal of Mathematical Sociology* (2): 113–120.
Chiaraviglio, L., M. Mellia, and F. Neri. 2009. "Reducing Power Consumption in Backbone Networks." In *IEEE ICC'09* Dresden, Germany.
Chiaraviglio, L., M. Mellia, and F. Neri. 2011. "Minimizing ISP Network Energy Cost: Formulation and Solutions." *IEEE/ACM Transactions on Networking* (99): 1–1.
Fortz, B. and M. Thorup. 2000. "Internet Traffic Engineering by Optimizing OSPF weights." In *IEEE INFOCOM 2000* Tel-Aviv, Israel.
Global Action Plan. 2007. "An Inefficient Truth." Global Action Plan Report, http://globalactionplan.org.uk.
Glover, F. and M. Laguna. 1997. *Tabu Search*. Kluwer Academic Publishers.
Gómez, D., E. González-Arangüena, C. Manuel, G. Owen, M. del Pozo, and J. Tejada. 2004. "Splitting Graphs When Calculating Myerson Value for Pure Overhead Games." *Mathematical Methods of Operations Research* 59(3): 479–489.
Moretti, S. and F. Patrone. 2008. "Transversality of the Shapley value." *Top* 16(1): 1–41.
Myerson, R. B. 1977. "Graphs and Cooperation in Games." *Mathematics of Operations Research* 2(3): 225–229.
Sansò, B. and H. Mellah. 2009. "On Reliability, Performance and Internet Power Consumption." In *Proceedings of 7th International Workshop on Design of Reliable Communication Networks (DRCN 2009)* Washington, D.C., USA.
Shapley, L. 1953. "A Value for n-Person Games." *Contributions to the Theory of Games II*, pp. 307–317.
Shaw, M. 1954. "Group Structure and the Behaviour of Individuals in Small Groups." *Journal of Psychology* (38): 139–149.
Tucker, R., J. Baliga, R. Ayre, K. Hinton, and W. Sorin. 2008. "Energy Consumption in IP Networks." In *ECOC 2008*, Brussells, Belgium.
Webb, M. 2008. *SMART 2020: Enabling the Low Carbon Economy in the Information Age*. London: The Climate Group.
What Europeans do at Night. 2009. http://asert.arbornetworks.com/2009/08/ what-europeans-do-at-night/.