

The Green-Game: Striking a Balance between QoS and Energy Saving

Aruna Prem Bianzino, Claude Chaudet, Dario Rossi, Jean-Louis Rougier
Telecom ParisTech,
Institut Telecom, CNRS LTCI UMR 5141, Paris, France
Email: {bianzino, chaudet, drossi, rougier}@enst.fr

Stefano Moretti
CNRS UMR 7243 Lamsade
Université Paris-Dauphine, Paris, France
Email: Stefano.MORETTI@dauphine.fr

Abstract—The energy consumed by communication networks can be reduced in several ways. A promising technique consists in concentrating the workload of an infrastructure on a reduced set of devices, while switching off the others. This technique, called “resource consolidation”, is particularly appealing when routing data traffic over a lightly loaded network. However, deciding on the set of routers that can be safely switched off requires an accurate evaluation of their criticality in the network. In this work, we define a measure of criticality that does not only take into account the network topology, but also the traffic matrix. Employing a game theoretic approach, we model the scenario as a coalitional game, and show the effectiveness of the Shapley value as a criticality index. Shapley value ranking is used to drive the resource consolidation procedure, and is compared with other classical indexes on a real network scenario. Numerical results confirm that the proposed index provides a robust and relevant criticality measure, yielding a good tradeoff between energy efficiency and network robustness.

I. INTRODUCTION

The energy consumed by communication networks has received increased attention lately, due to its environmental, economical and marketing implications. Various research efforts, grouped under the “Green networking” denomination, advocate pushing energy awareness into communication networks design, devices and protocols.

Communication networks are usually designed with *over-provisioning* and *redundancy* in mind. However, a network does not have to face a permanent high load. Indeed, traffic requests exhibit a periodical behavior with diurnal peaks and nightly valleys [1]. When the load is low, it is theoretically possible to concentrate all the data traffic over a small subset of the links and devices, allowing the others to enter a power saving mode. Besides energy savings, this strategy, called “resource consolidation” [2], needs to preserve the network connectivity and quality of service (e.g., ensuring a minimum path diversity, limiting the maximum link utilization, etc.).

[2] evokes resource consolidation as a hypothetical working direction. [3] then formulates the problem as an Integer Linear Program (ILP). They detail and evaluate greedy heuristics that consist in progressively switching off nodes and links, ordered according to the amount of routed traffic. [4] uses a similar approach, but considering only the links. In [5] we present numerical solutions to the ILP considering both nodes and links, considering real topologies and traffic matrices. [6] models the tradeoff between network robustness and energy consumption as an optimization problem whose objective

function is a weighted sum of the total energy consumed by the network and of a function of the link utilization level, to account for robustness. [7] overviews additional related work concerning the resource consolidation paradigm.

However there is no satisfactory definition of the criticality of nodes in a network. Classical indexes rank routers based either on the sole *topological* aspects, such as betweenness centrality, degree, closeness, eigenvectors, or on the sole *traffic* load. In this work, we use Game theory to define a criticality index that takes into account both aspects and to define heuristics for the resource consolidation problem.

In this article, we model the resource consolidation problem as a cooperative Transferable Utility Game, the *Green-Game* (or G-Game for short). This 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. We compute, for every node, its *Shapley value* that 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). The Shapley value on the G-Game defines a joint topology-aware and traffic-aware ranking of the network devices.

Our numerical results, obtained over a real network topology and traffic matrix, show that the proposed index provides a reliable critically measure in order to determine which nodes can be safely switched off, and in which order. Comparison with other indexes show that, for a similar level of energy savings, the network quality of service improves (e.g., maximum link load is lower) when Shapley value based ranking drives the resource consolidation process.

Sec. II introduces the necessary notation and describes the game theoretic model. To give the reader the intuition behind the G-Game model, we refer to simple but illustrative toy-case examples. Sec. II discusses also the computational complexity of the Shapley value. We propose heuristics to compute these values, from a pure theoretical point of view (e.g., decomposition in unanimity games) and relying on practical aspects (e.g., bounding the maximum path length). Sec. III details the scenarios used in the investigation, presents and then analyzes the performance of our proposal, along with a thorough comparison with other criticality indexes. Finally, Sec. IV gathers conclusive remarks and future working directions.

II. G-GAME DEFINITION AND SHAPLEY VALUE

A. Network Model and Associated Shapley Value

A communication network can be represented as a graph $G = \langle N, E \rangle$. 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., $n = |N|$). 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.

Communication networks are dimensioned based on measurements and estimates of the volume of data they have to support under realistic conditions. Various scenarios such as daytime traffic, nighttime traffic, etc. may be considered. Each scenario is characterized by a *traffic matrix*, $T = (t_{i,j})_{i,j \in N}$, in which an element $t_{i,j}$ represents the volume of traffic entering the network through node i and exiting through node j . We denote by v the total traffic load that the network has to route, with respect to a given traffic matrix T : $v(N) = \sum_{i,j \in N} t_{i,j}$.

As networks are usually dimensioned with peak hours in mind, energy may be saved during some low utilization periods by switching off some nodes. The goal is to find which subset of the network nodes may be switched off, without affecting the ability of the network to support the relevant traffic matrix. Let us consider an arbitrary subgraph of G , $G_S = \langle S, E_S \rangle$ formed by the nodes $S \subseteq N$ and by the corresponding edges subset $E_S = \{\{i, j\} : i, j \in S\} \subseteq E$.

The amount of traffic that G_S can effectively transport, with respect to T , is denoted by $v(S) = \sum_{i,j \in S} t_{i,j} \mathbf{1}_{\{G_S\}}(i, j)$, where $\mathbf{1}_{\{G_S\}}(i, j) = 1$ 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*, which we call the *G-Game* from now on. A group of nodes (players), $C \subseteq N$ is called a *coalition* and the value $v(C)$ is called the *worth* 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 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 : $\Sigma_N = \{\sigma : N \rightarrow N : \sigma \text{ is a bijection}\}$. 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 :

$$\phi_i = \sum_{\sigma \in \Sigma_N} \frac{m_i^\sigma}{n!}. \quad (1)$$

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. Sec. III provides 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 [8].

B. Efficient Computation of Shapley Value

The computation of Shapley value according to (1) 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* [8]. This decomposition provides a less expensive method to calculate the Shapley value. For a set of players N , an 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 $\mathcal{P}(N)$:

$$v = \sum_{C \in \mathcal{P}(N)} \lambda_C \cdot u_C, \text{ with } \lambda_C(v) \in \mathbb{R}, \forall C \in \mathcal{P}(N), \quad (2)$$

where λ_C are called the Harsanyi dividends [9], that are defined recursively by:

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

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:

$$\phi_i = \sum_{C \in \mathcal{P}(N); i \in C} \frac{\lambda_C}{|C|}. \quad (3)$$

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 [10]: 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 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 Fig. 1a. 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, $\pi((A, B)) = \pi((B, A)) = \{A, B\}$. Let us also denote by $\mathcal{P}_k(\mathcal{M}_E(\{i, j\}))$ 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 $\pi(p) = \pi(p_1) \cup \pi(p_2) \cup \dots \cup \pi(p_k)$ for a path $p = \{p_1, p_2, \dots, p_k\}$. The following expression defines the graph-restricted game [11], by introducing a characteristic function for a unanimity game that removes the influence of augmented paths:

$$u_{i,j} = \sum_{k=1}^{K_{ij}} \left(\sum_{p \in \mathcal{P}_k(\mathcal{M}_E(\{i,j\}))} (-1)^{k+1} \cdot u_{\pi(p)} \right). \quad (4)$$

To better understand the rationale behind (4), let us consider the toy-case example of Fig. 1a. 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 path and restrict our computations on the set $\mathcal{M}_E^*(\{i, j\}) =$



(a) Two acyclic paths exist between i and j , one augmenting the other. (b) Alternate paths exist between i and j . Dark arrows represent the traffic from node i to node j .

Fig. 1: Toy examples illustrating Shapley value computation

$\{P \in \mathcal{M}_E(\{i, j\}) : \nexists Q \in \mathcal{M}_E(\{i, j\}) Q \subset P\}$. 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 $\mathbf{1}_{\{h\}}(p) = 1$ if node h belongs to p , and $\mathbf{1}_{\{h\}}(p) = 0$ 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

$$\phi_h(i, j) = t_{i,j} \sum_{k=1}^{K_{ij}^*} \left(\sum_{p \in \mathcal{P}_k(\mathcal{M}_E^*(\{i,j\}))} \frac{(-1)^{k+1}}{|\pi(p)|} \cdot \mathbf{1}_{\{h\}}(p) \right). \quad (5)$$

For the sake of illustration, let us consider the example depicted on Fig. 1b. 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:

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

This value shows 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.

C. Practical Considerations on Shapley Value Computation

Computing the Shapley value using (5) 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 [12]. 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 $(i, i_1, i_2, \dots, i_n)$ is explored, the already visited nodes i, i_0, \dots, i_{n-1} cannot be visited again due to the loop-less path constraint. Then, the branch should also avoid




Fig. 2: The reference topology.

neighbors of preceding nodes, as this would otherwise lead to augmented paths.

Let us consider for instance the network represented in Fig. 2. Let us consider the path from node 1 to node T and focus on the branch $(1, 2, 13)$. Node 13 has two neighbors: 12 and 14, but the exploration only needs to consider node 14. The branch $(1, 2, 13, 12)$ would actually build augmented paths, such as $(1, 2, 13, 12, T)$, with respect to the branch $(1, 2, 12)$ that would reach T with a shorter path $(1, 2, 12, T)$. Therefore, the exploration has to skip node 12, as it is already a neighbor of node 2 (i.e., a neighbor of a predecessor node in the branch). The taboo list is then populated by the set $\bigcup_{n=i, i_0, \dots, i_{n-1}} \mathcal{N}_n$, where \mathcal{N}_n represents the set of neighbors of a node n .

Even when only the limited set of valid paths are considered, the Shapley value computation from (5) becomes intractable as the number of paths 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 (5)). 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. Fig. 3 reports the G-Game Shapley value of all nodes, for




Fig. 3: Node ranking for different maximum path lengths.




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

growing maximum path length L , ranging from $L = 4$ hops (i.e., the diameter of the network in Fig. 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.

The path length bound L makes Shapley value computation feasible. Fig. 4 represents, for example, the number of paths in the set $\mathcal{M}_E^*({1, T})$ for different maximum length L for the flow between node 1 and node T in Fig. 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 saturates). Looking more closely however, we can notice that the path of length greater or equal to 6 are already very "long paths" that are unlikely to be used in an operational networks (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 reduce 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 Fig. 3).

III. NUMERICAL RESULTS

Once the Shapley value is computed for every node in the network, it can be used as ranking to determine in which order nodes should be switched off. In this section, we evaluate the tradeoff between QoS and energy savings, comparing the proposed method with other classical node ranking schemes.

To provide a relevant evaluation, we take special care in 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 Fig. 2, represents a portion of the ISP access/metropolitan network segment. The light-shaded nodes (1 to 8) are access nodes, source and destination of traffic requests, and can not be switched off. The dark nodes (9 to 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 [13], widely accepted in the literature. The power consumption P_i (in Watts) of a node, is related to its switching capability C_i (in Mb/s) according to $P_i = C_i^{2/3}$. We consider that a node is able to switch twice the capacity of its entire set of connected links – allowing the ISP to add a reasonable number of links to the same device, avoiding the need for an upgrade. As link power consumption is negligible with respect to node consumption [5], 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).

A. The G-Game vs. Other Possible Criticality Rankings

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 [14]), on the number of shortest paths passing through each node (Betweenness centrality [15]), on the average distance between each couple of nodes (Closeness centrality [16]), or on the importance of nodes neighbors (Eigenvector centrality [17]). A completely different criticality criterion, proposed by [3] and denoted by “Load” hereafter, merely consists in sorting nodes depending on the amount of load they effectively carry in a standard network configuration.

The above indexes either consider the topology or the traffic, but not both. The Shapley value used in the G-Game instead takes into account (i) the traffic expressed by the traffic matrix




Fig. 5: Node criticality when considering only the network topology in the G-Game, i.e., G-Game U-TM (a), and when considering also the real traffic matrix, i.e., the full G-Game (b).

and (ii) the importance of the node in the routing process. The 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.

All the above listed criticality indexes 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. Fig. 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 the Pearson correlation coefficients between every pair of rankings. Results are summarized in Tab. I, 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: Load 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

TABLE I: Correlation coefficients between the rankings defined by different criteria

	G-Game (U-TM)	Betweenness	Degree	Closeness	Eigenvector	G-Game	Load
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	
Load	0.43	0.49	0.48	0.60	0.19	0.56	1.00

G-Game U-TM hereafter. Degree and Eigenvalues also form a distinct family which is omitted below for space reasons and as they are less pertinent.

B. Energy Savings

Let us first evaluate the ability of different criticality indexes to reduce energy consumption. 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 the (iii) Load ranking, based only on the traffic matrix. The resulting orders of nodes are reported in Tab. II.

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 Tab. II. The process can stop as soon as one check fails (e.g., switching off nodes 9 and 15 for the G-Game ranking), or it can continue until the whole list is scanned (i.e., switching off all the underlined nodes in Tab. II). Alternatively, the search procedure may be stopped not at the first check failure, but at the second, leaving a gap of powered on-nodes among the switched off nodes, tolerating a few consecutive ranking “errors” (e.g., switching off nodes 9, 15, 16 and 21 for the G-Game).

These three stopping strategies are labeled (according to the number of gaps allowed in scanning the list) “none”, “any” and “one” in Tab. III, that compares their respective achievable energy savings. When scanning the whole list, both the G-Game and the Load orderings are able to obtain the maximum achievable energy saving, reaching the optimal configuration obtained as in [5]. Notice that nodes that can be switched off are less critical in the G-Game ranking with respect to the Load ranking: hence, they are found earlier during the list scan, which explains why G-Game achieves better results in the “one” strategy.

TABLE III: Maximum achievable energy saving, in percentage, when considering different criticality rankings and different stopping criteria.

Ranking	Number of gaps (stop criterion)		
	none	one	any
G-Game	10.85	22.47	29.40
G-Game (U-TM)	0.0	13.43	18.88
Load	10.85	16.27	29.40
Baseline	0.0	0.0	0.0




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

C. Energy Savings vs QoS

So far, we have relatively compared different ranking strategies, and have evaluated their efficiency in reducing energy consumption. However, 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 $N_{off} = 3$ 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 Tab. II. 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 [18] to perform an optimization of the routing weights (using the IGP-WO algorithm [19]) and route the traffic enabling Equal Cost Multi Path (ECMP). It follows that we are able to not only 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 Tab. IV, together with the average path length \bar{l} ; we also report a weighted

TABLE II: List of the network nodes, ordered (left to right) from the least to the most critical one, according to different criticality rankings. Underlined values identify nodes that can be switched off such that the topology remains able to carry the traffic matrix. Bold values identify nodes considered in Sec. III-C and Sec. III-D.

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

average path length \bar{l}_{TM} , where paths are weighted by the amount of traffic they transport over the traffic matrix. 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 consider again the toy case of Fig. 1b, 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 $\bar{l} = \bar{l}_{TM} = 2.5$ 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, $\bar{l} = \bar{l}_{TM} = 3$ hops), or to a *decrease* (i.e., only the (i,A,j) path is available, $\bar{l} = \bar{l}_{TM} = 2$ hops) in the average path length, depending on which nodes are switched off.

Finally, Fig. 6 reports link utilization distributions for the different rankings when $N_{off} = 3$ and the baseline configuration $N_{off} = 0$. Notice that G-Game yield to excellent performances, as the link distribution is roughly equivalent to the baseline configuration, where no node is switched off. Especially, maximum link utilization does not increase under G-Game with respect to the full network configuration: this means that energy saving is obtained without compromising the expected QoS. Conversely, some links reach an utilization higher than 90% for the U-TM and Load strategies. The Load 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).

D. Sensitivity Analysis to Traffic Matrix Variation

To gather consistent results, we consider further variations of the original scenario. The original traffic matrix presents high centralization, in the sense that most of the traffic has node T as source or destination. We therefore “smooth” the traffic matrix in a controlled fashion, keeping constant the overall traffic volume and number r of traffic demands. In more details, let $x_{s,d}$ denote a traffic request in the original traffic matrix, for a given source node s , and destination node d . In the case of a smoothed traffic matrix, every traffic request

TABLE IV: 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}_{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
Load	9, 15, 21	2.64	3.25	16.27
Baseline	None	2.64	3.13	0.0




Fig. 7: Variation of the link utilization distribution, for scenarios ranging from the original traffic matrix ($p = 0$) to the smoothed traffic matrix ($p = 1$).

between access nodes pairs is then equal to:

$$\bar{x} = \frac{1}{r} \sum_{(s,d) \in N} x_{s,d}$$

We can now define a smoothing parameter $p \in [0, 1]$ to tune the traffic between the original traffic matrix ($p = 0$) and the smoothed traffic matrix ($p = 1$). In any intermediate scenario individuated by a given value of p , the elements in the traffic matrix(p) are set to:

$$x_{s,d}(p) = x_{s,d} - p(x_{s,d} - \bar{x})$$

We consider again the case $N_{off} = 3$, and evaluate values of p ranging from 0 to 1 by steps of 0.1. Fig. 7 reports, for every value of p the average and maximum link load, for the baseline and G-Game configurations. As we can see, sensitivity to traffic matrix variations is limited, with similar




Fig. 8: Variation of the energy saving achieved by the G-Game, for scenarios ranging from the original traffic matrix ($p = 0$) to the smoothed traffic matrix ($p = 1$).

maximum link utilization in both cases (i.e., between 60 and 80%). Also the energy savings achieved by the G-Game only minimally vary over the same traffic matrix range (i.e., varying between 16.3% for $p = 0$, to 17.1% for $p = 1$). Results are reported in Fig. 8).

Fig. 9 reports maximum and weighted average path length for the same set of scenarios. We see that, as traffic matrix smoothness increases, average path length decreases, as nodes tend to exchange more traffic with neighbors. On the other hand, the maximum path length remains constant, and equal to the network diameter (which further confirms the soundness of maximum path length L bound in Shapley value computation).

IV. CONCLUSIONS AND PERSPECTIVES

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 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 ranking used to drive a resource consolidation process, i.e., to switch off nodes. Numerical results on a realistic network scenario show that Shapley value ranking yield 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 the full network scenario.

In future work, we aim at broadening our experimental studies over a wider set of network topologies and traffic matrices. In particular, we aim at digging further the correlation between the criticality of nodes and different traffic matrices for any given topology. Another open point that we aim at exploring is to refine the evaluation of robustness, by considering the impact of unexpected faults, or changes in the traffic conditions, to an already consolidated network.

ACKNOWLEDGMENT

This work was funded by the Eureka Celtic TIGER2 project and by the NimbleNet project.




Fig. 9: Variation of the maximum and weighted average path length, for scenarios ranging from the original traffic matrix ($p = 0$) to the smoothed traffic matrix ($p = 1$).

REFERENCES

- [1] What Europeans do at Night, “<http://asert.arbornetworks.com/2009/08/what-europeans-do-at-night/>.”
- [2] M. Gupta and S. Singh, “Greening of the Internet,” in *ACM SIGCOMM’03*, (Karlsruhe, Germany), ACM, Aug. 2003.
- [3] L. Chiaraviglio, M. Mellia, and F. Neri, “Reducing Power Consumption in Backbone Networks,” in *IEEE ICC’09*, (Dresden), June 2009.
- [4] W. Fisher, M. Suchara, and J. Rexford, “Greening Backbone Networks: Reducing Energy Consumption by Shutting Off Cables in Bundled Links,” in *Proceedings of 1st ACM SIGCOMM Workshop on Green Networking*, (New Delhi, India), August 2010.
- [5] A. Bianzino, C. Chaudet, F. Larroca, D. Rossi, and J.-L. Rougier, “Energy-Aware Routing: a Reality Check,” in *3rd International Workshop on Green Communications (GreenComm3)*, (Miami, USA), December 2010.
- [6] B. Sansò and H. Mellah, “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), Oct. 2009.
- [7] A. P. Bianzino, C. Chaudet, J. Rougier, and D. Rossi, “A survey of green networking research,” *IEEE Communications Surveys & Tutorials (to appear)*, 2011.
- [8] S. Moretti and F. Patrone, “Transversality of the Shapley value,” *Top*, vol. 16, no. 1, pp. 1–41, 2008.
- [9] L. Shapley, “A Value for n-Person Games,” *Contributions to the Theory of Games II*, pp. 307–317, 1953.
- [10] D. Gómez, E. González-Arangüena, C. Manuel, G. Owen, M. del Pozo, and J. Tejada, “Splitting graphs when calculating Myerson value for pure overhead games,” *Mathematical Methods of Operations Research*, vol. 59, no. 3, pp. 479–489, 2004.
- [11] R. Myerson, “Graphs and cooperation in games,” *Mathematics of Operations Research*, vol. 2, no. 3, pp. 225–229, 1977.
- [12] F. Glover and M. Laguna, *Tabu Search*. Kluwer Academic Publishers, 1997.
- [13] R. Tucker, J. Baliga, R. Ayre, K. Hinton, and W. Sorin, “Energy consumption in IP networks,” in *34th European Conference on Optical Communication (ECOC’08)*, 2008.
- [14] M. Shaw, “Group structure and the behaviour of individuals in small groups,” *Journal of psychology*, no. 38, pp. 139–149, 1954.
- [15] A. Bavelas, “A mathematical model for small group structures,” *Human Organization*, no. 7, pp. 16–30, 1948.
- [16] M. Beauchamp, “An improved index of centrality,” *Behavioral Science*, vol. 10, no. 2, pp. 161–163, 1965.
- [17] P. Bonacich, “Factoring and Weighting Approaches to Status Scores and Clique Identification,” *Journal of Mathematical Sociology*, no. 2, pp. 113–120, 1972.
- [18] S. Balon, J. Lepropre, O. Delcourt, F. Skivée, and G. Leduc, “Traffic Engineering an Operational Network with the TOTEM Toolbox,” *IEEE Transactions on Network and Service Management*, vol. 4, pp. 51–61, June 2007.
- [19] The Interior Gateway Protocol Weight Optimizer (IGP-WO) algorithm, “<http://totem.run.montefiore.ulg.ac.be/algo/igpwo.html>.”