



HAL
open science

Resilience in Computer Network Management

Marcelo F. Vasconcelos, Ronaldo M. Salles

► **To cite this version:**

Marcelo F. Vasconcelos, Ronaldo M. Salles. Resilience in Computer Network Management. 11th International Networking Conference (NETWORKING), May 2012, Prague, Czech Republic. pp.109-120, 10.1007/978-3-642-30045-5_9. hal-01531124

HAL Id: hal-01531124

<https://inria.hal.science/hal-01531124v1>

Submitted on 1 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Resilience in Computer Network Management

Marcelo F. Vasconcelos and Ronaldo M. Salles

Instituto Militar de Engenharia,
Praça General Tiburcio, 80. Rio de Janeiro, Brasil
marcelo.vasconcelos@globocom.com, salles@ieee.org

Abstract. The protection of network infrastructure and services is one of the main issues network managers face today. This paper investigates this matter through the study of network resilience. A resilience factor (RF) is proposed to take into account topological aspects of the network and also the amount of traffic losses under stress conditions. The proposed factor is evaluated using a real network backbone (RNP). Results obtained showed that the resilience factor is consistent and can be employed to support network managers decisions about network expansions, link additions and removals. Such decisions improve network robustness making it less vulnerable to attacks and failures.

Keywords: Resilience, Network Management, Network Link Failures, Traffic Loss.

1 Introduction

Among the various tasks a network manager has to perform, the study on how to improve robustness to better support faults and attacks to the network infrastructure and services is a critical issue. Such study usually involves the analysis of redundancies, alterations and expansion of the network resources.

In this sense, decisions about where to efficiently invest at a reasonable cost often depend on the experience of each network manager. Despite taking into account information concerning the operation of the network, the process is very subjective and may lead to poor or not effective choices.

The main goal of this work is to propose a measure of resilience in order to better quantify this important property, and thus provide a methodology for the use of resilience in the management of computer networks. The metric can be employed to assist network design and also to support decision-making regarding expansions or changes in the topology of an operational network. It is thus possible to decide between one or another change in the topology based on the impact of each one to the resilience of the network.

The first step toward this goal is to propose a resilience factor (RF) in order to quantify the network level of resilience. The proposed resilience factor takes into account both the topological aspects as well as the traffic demands posed to the network. After being able to measure network resilience through RF, one can search for alterations in topology in order to obtain a better outcome in terms of RF.

Finally, the proposed methodology is evaluated using a real network scenario where results obtained conforms to expectations of network managers.

This paper is organized as follows. Section II reviews the related work on the topic. Section III presents our proposal: the resilient factor, methodology and algorithm. Section IV discusses the results obtained using a real network scenario, and finally Section V presents the conclusions and suggestions for future work.

2 Related Works

The quest for increasing robustness of the network led to the study of resilience. Network resilience is a broad area, this work considers the following definition: *“Resilience in networks is the ability of an entity to tolerate, endure and automatically recover from challenges under the conditions of the network, coordinated attacks and traffic anomalies”* [1].

From the above definition, several works in the literature studied resilience using different techniques and methods. They can be roughly divided into two main categories: reactive and proactive approaches. Reactive strategies are mostly devoted to routing and other network configuration mechanisms that are triggered to react to a given fault or anomaly condition observed in the network, whereas proactive approaches try to condition the network beforehand to better resist to faults, anomalies and attacks. The focus of this work is on proactive approaches.

A quantitative and statistical analysis of the frequency and duration of faults are carried out in [2], [3] and [4], the results are equivalent and were obtained in real topologies. It was shown that failures are part of network operation and most of them last for less than 10 minutes. Regarding interruptions of network services, scheduled interruptions correspond to 20 % and unplanned interruptions correspond to 80 % of the total. Within the 80 % of unscheduled interruptions, 70 % correspond to a single link failures and 30 % are attributed to multiple failures of equipment and fiber optic networks. This information is important and could be taken into account in the development of resilience metrics.

The work [6] brings a comprehensive analysis of network resilience, addressing topics such as failures in networks, mechanisms of resilience and analysis. The purpose of that paper is to present a methodology based on the calculation of autonomous systems edge connections availability using distribution functions of link occupation. Such functions depend on network failures and traffic variations. Failures can be single, double, triple and so on. The traffic considered is proportional to the size of the population in a given area. Results show that link additions to the network provide a greater network availability.

The work in [7] was one of the first to introduce a method to measure the fault-tolerance of a network. The measure was defined as the number of faults a network may suffer before being disconnected. The authors computed an analytical approximation of the probability of the network to become disconnected and validated their proposal using Monte Carlo simulation results. The simulation scenario employed three particular classes of graphs to represent network

topologies: cube connected-cycles, torus and n -binary cubes – all of them are symmetric and with same node degrees.

The authors in [8] were interested in evaluating the robustness of a network against attacks and node failures. They applied the concepts of node connectivity and topology symmetry. The evaluation of their method was carried out using different groups of networks: symmetric topologies, scale-free, and random topologies. In [9] a similar method was presented, but evaluation was performed considering only network topologies that are scale free to better adhere to real Internet characteristics. A scale free topology [10] is such that the degree of its nodes follows the power law distribution [11].

The method proposed in [12] used the k -connectivity property of a graph to construct a resilience factor for network topologies. The authors showed that the method was consistent and more efficient than previous approaches. However, the computational cost to compute the factor is quite high and the method does not take into account the capacity of links and network traffic.

On the other hand, both network traffic and link capacity were considered by the authors in [13] and [14]. The parameters were related by the concept of *delivered value: the ratio of traffic routed through a network after a capacity decrease or loss and the value sent before the reduced capacity*. According to the previous definition:

$$DV = \frac{(DeV - LOSS)}{DeV} \quad (1)$$

where DV is the delivered value, DeV is the amount of traffic that should pass through the link and $LOSS$ is the amount of traffic that was not delivered due to network changes (capacity decrease or loss). In their work a complete disconnection was not considered but the capacity of a given link could be reduced due to failures, which might impact DV .

It can be observed from previous works that in order to provide a more complete view about network resilience it is important to consider not only connectivity aspects of the network topology as in [12], but also to compute the impact on traffic caused by network changes as in [13] and [14].

3 Proposed Method

This paper proposes a method to quantify resilience of a given network through the resilience factor, RF . As mentioned before, the proposed method does not focus on a single network aspect but can be viewed as a composite metrics of delivered traffic and topology characteristics as follows:

$$RF = (ADV, R) \quad (2)$$

where ADV is the average delivered value and R accounts for the number of redundancies the network topology enjoys. In this work, R is also defined as the *network order*. Details about the computation of RF will be presented next.

3.1 Resilience Factor

The first aspect to analyze is how network topology is considered in the term R that composes RF .

From the definition and all previous works discussed in Section 2, it is possible to realize that resilience is associated to redundancy in the following way: a network enjoys a high level of resilience if it has a large number of redundant resources to cover possible losses during attacks and/or failures. We try to quantify this notion using the concept of network order.

In this work, *network order* (or simply R) is computed by the number of redundant links the network has to connect its nodes. In other words, R is defined as the maximum number of links that can be removed and still preserves network connectivity – the resulting topology becomes a spanning tree.

For instance, suppose a network where nodes are connected only by the spanning tree (ST) itself, in this case any failure causes a disconnection and traffic losses since there is no spare resources to keep nodes connected. Hence, a ST network is of zero order: $R = 0$.

On the other hand, a full mesh network has $\frac{n(n-1)}{2}$ links, where n is the number of nodes in the topology. For this network the spanning tree (ST) is composed by $n - 1$ links, and so there are $\frac{n(n-1)}{2} - (n - 1)$ extra links connecting nodes, yielding $R = \frac{(n-2)(n-1)}{2}$.

Therefore, the parameter *network order* is limited by these two practical situations that can be found in real network topologies:

$$0 \leq R \leq \frac{(n-2)(n-1)}{2} \quad (3)$$

We would like to study resilience taking into account redundant links only and how their disconnection impacts the delivered traffic, or in another way, how the network depends on them to deliver its traffic. For example, a topology that has no redundancy but that can route all traffic, will have $RF = (1, 0)$ indicating that $ADV = 1$ (no losses) and $R = 0$ (no redundancy). If in another network we can draw five links out of it and all traffic continues to be delivered then $RF = (1, 5)$, and so on.

ADV is based on Eq. 1, in fact it is computed as the average of DV in each x different network conditions given by links removals/failures, leading to $DV(x)$. Note that differently from [13] and [14] Eq. 1 is applied to the network as a whole and not to a specific link only. The objective is to evaluate how topology changes due to link failures and reduce its capacity to deliver traffic.

Let us illustrate the computation of $RF = (ADV, R)$ using the example in Fig.1. The topology in Fig.1a has 4 nodes and 5 links, where the capacity of each link is given in Mbps. The computation of R is directly obtained from the topology: $R = 5 - 3 = 2$, i.e. the topology has two redundant links in addition to the links of the spanning tree. The *network order* (R) is used to evaluate all failure conditions (link removals) regarding network redundancies only, and how they impact the delivered traffic.




Fig. 1. Example: obtaining ADV

Hence, according to our approach there are $2^R = 4$ failure conditions to be tested: no failure (Fig.1b), one failure (Figs.1c and 1d) and two failures (Fig.1e). In fact, if more than two failures occur the network will be disconnected and all traffic from and to the disconnected node will be lost. This is an extreme situation that is less likely to happen than the other cases (see Section 2) and will not be considered in this work. Such extreme cases are studied in [12]. It is important to notice that although extreme fault scenarios (nodes disconnection) are not taken into account in the RF , the computation of ADV considers worst case analysis.

First, in Fig.1b it can be seen that the highest network link load is 37%, and so all traffic is delivered since there is no loss in the network, for this first scenario $DV(1) = 1$. The link of highest occupancy is selected to be removed from the topology (worst case) yielding in the scenario of Fig.1c. All traffic crossing that link has to be rerouted to other links, loads on the remaining links increase, however the network has enough spare resources to support such failure and all network demands are still delivered: $DV(2) = 1$. Now the diagonal link is removed to generate the situation in Fig.1d, where still $DV(3) = 1$. Finally, both spare links are removed and the network is connected only by its spanning tree (limiting situation before node disconnection). In Fig.1e there is a link operating at 100%, in fact demands exceed link capacity and there is loss of traffic in the network: $DV(4) = 0.6$.

The parameter ADV is then computed as the mean DV for all failure situations considered:

$$ADV = \frac{DV(1) + DV(2) + DV(3) + DV(4)}{4} = 0.9 \quad (4)$$

and $RF = (0.9, 2)$ for the network in the example. This can be read as: the network has two redundant links and can deliver almost all demanded traffic even in situations involving failures related to the spare capacity (redundancies).

A generalization of Eq. 4 is possible if it is known beforehand the probability of each one of the failure scenarios illustrated in Fig.1 occurs. Thus, the corresponding DV could be weighted, α_i , according to the probability of the specific failure occurs:

$$ADV = \sum_{i=1}^{2^R} \frac{\alpha_i DV(i)}{\beta} \quad \beta = \sum_{i=1}^{2^R} \alpha_i \quad (5)$$

Since such information is not always available and to simplify the approach, we consider in this work all $\alpha_i = 1/\beta$.

The algorithm to calculate the resilience factor is presented below. It takes as inputs the information regarding the network topology (CapM, n and m) and traffic demands (TM), and outputs the resilience factor (ADV and R).

Algorithm 1 RF(In: CapM, TM, n , m ; Out: ADV, R)

```

1:  $ADV \leftarrow 0$ 
2:  $R \leftarrow m - (n - 1)$ 
3:  $DeV \leftarrow$  sum all elements in TM
4: compute  $Q$ 
5: for all elements  $q \in Q$  do
6:   compute  $CapM(q)$  from CapM
7:   compute  $WM(q)$  from WM
8:    $R \leftarrow$  routes calculated using  $WM(q)$ 
9:    $OM \leftarrow$  occupation obtained from TM and R
10:   $LOSS \leftarrow OM - CapM(q)$ 
11:   $DV \leftarrow \frac{(DeV - LOSS)}{DeV}$ 
12:   $ADV \leftarrow ADV + \frac{DV}{2R}$ 
13: end for
14: return ( $ADV, R$ )

```

Legend:

- . n : number of nodes in the topology
- . m : number of links in the topology
- . $CapM_{n \times n}$: adjacency matrix with link capacities (topology)
- . Q : set of all failure scenarios, as illustrated in Fig.1
- . q : element of Q , one particular scenario as in Fig.1d
- . $TM_{n \times n}$: traffic demand matrix, origin-destination pair demands
- . $WM_{n \times n}$: weight matrix for shortest-path routing (e.g. RIP or OSPF [15])
- . R : routing matrix, all path from origin to destination obtained using WM
- . OM: link occupation matrix
- . LOSS: traffic loss matrix

The first step of the algorithm is to initialize ADV to zero. Then in step 2, the *network order* is computed using the number of links and nodes in the topology. Step 3 computes the network demanded value DeV summing up all elements of the TM matrix. Step 4 computes the set Q of all topology combinations obtained from the removal of redundant links as previously illustrated in Fig.1.

In the next step in line 5, the algorithm enters in a loop to check the behavior of each one of the elements in Q (as seen in Fig.1) and accumulates DV for each of them in the average ADV (line 12). After that, the algorithm ends returning the resilience factor.

Inside the main loop the first step in line 6 is to update CapM for the particular element in Q , obtaining $CapM(q)$. Then in line 7 the weight matrix is also updated to $WM(q)$. Routes are compute and stored in R since they will

be used to return OM. In line 10, the loss matrix is calculated by subtracting elements in OM and $CapM(q)$. DV is computed in line 11 and ADV in line 12. After that, a new element in Q (other topology combination) is processed and the loop goes on until the last element in Q .

The occupation matrix OM is obtained using the values in TM, as follows: for each demand in TM the corresponding route is determined between origin to destination, then the amount of demanded traffic is added to each element (link) of OM that takes part of the route. Later in the following step, CapM is applied to compute possible losses: any element in OM above the corresponding element of CapM results in loss of traffic since demands exceed capacity. The matrix LOSS store such values.

The delivered value DV is the difference between DeV and traffic losses due to the lack of capacity on links composing the network to absorb demands. DV equals one whenever all network demands are delivered, i.e. no loss is experienced in the network.

With the use of the resilience factor proposed, it is possible for instance to verify the effect of additions and removals of links connecting two nodes. For each particular situation the network manager may use RF to assess the impact such changes cause on the network. This method provides support for several management decisions the network manager has to take during the operation of the network under his responsibility.

4 Results

In this section we apply the proposed method to evaluate the resilience of a real network topology: the Brazilian National Research Network (RNP) in Fig. (2). The idea is to employ RF and analyze how changes in topology affect the network and its traffic. RF is also used as a tool to provide important information to support decisions about network expansion and protection against failures and attacks.

The Brazilian National Research Network has 27 nodes, 29 links and 7 redundancies that covers all states of Brazil. This network was selected since there is available information on the web (<http://www.rnp.br>) to conduct our experiment, for instance it was possible to obtain source destination peak demands of network traffic. However, information about all links occupation is missing and had to be inferred.

We considered two main traffic patterns to obtain the missing information about links occupation: max-min and proportional fair share of network resources [16]. While max-min fair is widely adopted as an ideal form of network sharing, proportional fair is shown to better resemble “TCP-like” sharing of network links.

We also studied different load conditions to better assess our method under a greater number of scenarios: peak (rnp historical peak demands), 100%, 95%, 90%, 50% and 30% of the maximum link capacities in the topology.

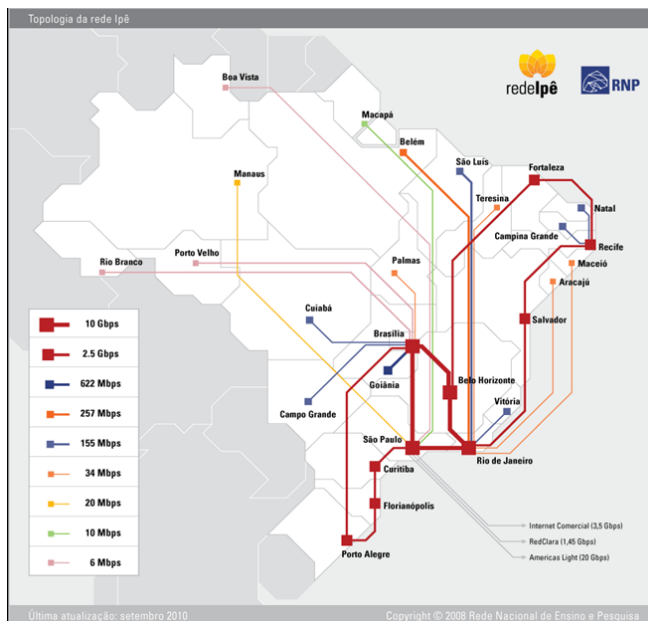


Fig. 2. Brazilian National Research Network

Regarding routing schemes, we considered the two most common forms of routing in our tests: minimum hop count (same link weights as used in RIP) and shortest paths according to link capacity (weights given by $100/\text{capacity}$ (Mbps) as used in OSPF [15]).

The objective of the test is to assess how the network (RNP) behaves when subjected to changes in its topology. How some network expansions (link additions) and removals affect the resilience of the network. This can be of great value to support network manager decisions.

To perform this test, we selected five additions and withdrawals of links, always seeking the best results: increase in network robustness for the case of additions and verification of worst cases of loss when links are removed.

We consider the following perception the network manager may have when analyzing the impacts to the network of a link addition or removal.

- a. (link addition): a reasonable increase in robustness may be achieved whenever a link is added to the network core – i.e. when it has the potential to serve a large number of origin-destination demands.
- b. (link removal): a network does not enjoy high levels of robustness whenever it heavily depends on certain core links – i.e. when a problem occur on those links it affects a large number of demands

Considering the principles described before, tests were taken according to Table 1, where the addition of a link is represented by the sign + and removal




Fig. 3. Occupation method proportional fair

is represented by -. The link connecting nodes A and B is represented by (A, B). In the case of additions, the value of the new link is presented. The first line represents the baseline before any change in the topology. In Figs. 3(a), 3(b), 4(a) and 4(b) the first green bar represents the baseline with no change in network, the blue bars represent the additions and the red bars represent the removals corresponding to Table 1, respectively.

Table 1. Tests in RNP

Change	Gbps	Change	Gbps
Initial	no change	Initial	no change
+(RJO,BSB)	10	-(RJO,SPO)	X
+(BSB,CWB)	10	-(BSB,SPO)	X
+(FZA,SDR)	2.5	-(SPO,CWB)	X
+(BSB,REC)	2.5	-(POA,BSB)	X
+(BHZ,REC)	2.5	-(FZA,REC)	X

The bar graphs presented from Figs. 3(a), 3(b), 4(a) and 4(b) correspond to the ADV obtained for each one of the eleven situations described in the table, respectively. Each figure represents a certain network scenario (bandwidth sharing and routing scheme used) under different load conditions (seven groups of bars). For instance, it was considered in Fig.4 a max-min fair share of link resources and a routing scheme following OSPF. The first bar in the figure indicates that $ADV = 0.6$ for the topology without changes (first line of the table) when links can be fully occupied (100%).

Some general conclusions can be made from the figures. First, bandwidth link sharing methods did not have a great impact on the results. For the same routing strategy and network loads, max-min and proportional fair results were close to each other.

The same does not apply to different routing methods, ADV varies significantly according to the routing scheme adopted. This can be explained given that routing has a direct effect on the distribution of traffic loads among links





Fig. 4. Occupation method maxmin fair

and then the delivered traffic also suffers from modifications on routing schemes. However, the general behavior was preserved for most of the cases, a change in topology that provides an increase in resilience under a certain network setting was also observed as positive in all other situations. For instance, the second bar in the graphs, which involves an addition of a 10Gbps link between Rio and Brasília (second line in the table). This addition conforms to the perception of a network manager as discussed before.


The first insertion between Rio de Janeiro (RJO) and Brasília (BSB) has the best results in terms of increased robustness to all situations of traffic profile and routing methods, noting this addition as most indicated. This result was confirmed in testing all of combinatorial additions, as shown in Fig. 5(a). The other additions have increased the RF below the insertion between Rio de Janeiro and Brasília, and in some cases, such as adding a link between Brasília and Recife, the value obtained by the ADV is less than the original. Regarding removals, we can see that removal of the link between Rio de Janeiro and São Paulo (SPO) is the most critical and leads to a more significant decrease in network robustness. In possession of this information, the network manager can take actions to ensure the connection between these two cities is preserved.

A not intuitive result regarding removals can be seen when the link between Porto Alegre (POA) and Brasília, or Recife (REC) and Fortaleza (FZA) is removed and an increase in ADV is obtained. This situation does not conform with the normal perception as discussed before but it indeed has an explanation. Such removals provide other ways to route traffic between those cities and this new configuration is preferred in terms of resilience. Those links were heavily used and the removal of them cause the traffic to be better distributed over other network paths yielding in a favorable resilience condition.

Combinational tests of additions and removals are shown in Figs. 5(a) and 5(b). For additions, the best result is the addition between cities of Rio de Janeiro and Brasília, with $RF = (0.7223, 7)$. For removals, the worst result occurs when the link between Rio de Janeiro and São Paulo is lost, with $RF = (0.4588, 6)$. The best result among all removals was between the cities of São Paulo and Curitiba (CWB), with $RF = (0.8229, 6)$. Also it is shown in Figs. 5(a) and 5(b) the results of other possible additions and removals, indicating a complete scenario to support the decision of network managers.



(a) Resilience Factor for all possible additions



(b) Resilience Factor for all possible removals

Fig. 5. Possible additions and removals

5 Conclusions

This paper investigated the problem of resilience in computer networks under a network manager perspective. A resilience factor (RF) was proposed to evaluate resilience according to the number of redundancies and also the traffic delivered by the network.

The method was tested using a real network backbone scenario and was shown to be consistent and useful to support decisions performed by network managers about expansions, additions and removals of links. General results conform with the general perception of managers, however the methodology highlighted other important aspects that usually need a deeper understanding of networking operations. The use of RF also indicates which are the most advantageous insertions and removals, the ones that provide greater impact to network operation.

Finally, it is important to mention that the proposed factor and methodology can be used in other scenarios where different types of networks operate: electric power transmission, oil flow, gas distribution, water, sewage, etc.

References

1. Aggelou, G. *Wireless Mesh Networking*. McGraw-Hill Professional (2008).
2. Iannaccone, G., Chuah, C.-n., Mortier, R., Bhattacharyya, S., and Diot, C. Analysis of link failures in an ip backbone. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*, IMW '02, pages 237–242, New York, NY, USA. ACM. (2002).
3. Nucci, A., Schroeder, B., Bhattacharyya, S., Taft, N., and Diot, C. Igp link weight assignment for transient link failures. In *in Proc. International Teletraffic Congress* (2003).
4. Markopoulou, A., Iannaccone, G., Bhattacharyya, S., Chuah, C.-N., and Diot, C. Characterization of failures in an ip backbone. In *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 4, pages 2307 – 2317 vol.4. (2004).
5. Li, V. and Silvester, J. Performance analysis of networks with unreliable components. In *Communications, IEEE Transactions on*, 32(10):1105 – 1110. (1984).
6. Menth, M., Duelli, M., Martin, R., and Milbrandt, J. (2009). Resilience analysis of packet-switched communication networks. In *IEEE/ACM Trans. Netw.*, 17(6): 1950–1963. (2009).
7. Najjar, W., Gaudiot, J.-L. Network resilience: a measure of network fault tolerance. In *Computers, IEEE Transactions on*, 39(2):174 –181, feb 1990. ISSN 0018-9340. (1990).
8. Centre, C. C., Dekker, A. H. and Colbert, B. Scale-free networks and robustness of critical infrastructure networks. In *Proceedings of the 7th Asia-Pacific Conference on Complex Systems*. (2004).
9. Dekker, A. H., Colbert, B. D. Network robustness and graph topology. In *Proceedings of the 27th Australasian conference on Computer science - Volume 26, ACSC '04*, págs. 359–368, Darlinghurst, Australia, Australia, 2004. Australian Computer Society, Inc. (2004).
10. Barabasi, A.-L., Albert, R. and Jeong, H. Scale-free characteristics of random networks: The topology of the world-wide web. In *Physica A: Statistical Mechanics and its Applications*, 281(1-4):69–77. (2000).
11. Faloutsos, M., Faloutsos, P., and Faloutsos, C. On power-law relationships of the internet topology. In *SIGCOMM Comput. Commun. Rev.*, 29:251–262. (1999).
12. Salles, R. M. and Marino, D. A. Strategies and metric for resilience in computer networks. In *The Computer Journal*. (2011).
13. Omer, M., Nilchiani, R., and Mostashari, A. Measuring the resilience of the global internet infrastructure system. In *Proc. 3rd Annual IEEE Systems Conf*, 156–162. (2009).
14. Omer, M., Nilchiani, R., and Mostashari, A. Measuring the resilience of the trans-oceanic telecommunication cable system. In *IEEE Systems Journal*, 3(3):295–303. (2009).
15. Moy, J. T. *OSPF Complete Implementation*. Addison-Wesley Professional. (2000).
16. Kelly, F., Maulloo, A., and Tan, D. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, volume 49. (1998).