

Is the topology of the Internet network really fit to sustain its function?

V. Rosato^{a,b,*}, L. Issacharoff^a, S. Meloni^c, D. Caligiore^d, F. Tiriticco^e

^a ENEA, Casaccia Research Centre, Computing and Modelling Unit, Via Anguillarese 301, 00123 S. Maria di Galeria, Italy

^b Ylichron S.r.l, Roma, Italy

^c Dip.to Ingegneria Informatica, Università di Roma Tre, Via della Vasca Navale, 79 00146 Roma, Italy

^d ISTC, CNR, Via S. Martino della Battaglia 44, 00185 Roma, Italy

^e Dip.to Ingegneria delle Telecomunicazioni, Università di Roma "Tor Vergata", Via del Politecnico, 1 00133 Roma, Italy

Received 13 February 2007; received in revised form 24 September 2007

Available online 1 November 2007

Abstract

The Internet is one of the most interesting realizations of a “complex” network. As a non-supervised growing object, it allows the study of the selective pressure which drives the network to assume its current structure. The DIMES and the ROUTE VIEWS projects are ongoing projects aimed at evaluating the topological structure of the Internet (at the Autonomous System or AS grain-level) on the basis of different types of measurements. The topological analysis of the networks produced by the two projects has allowed us to infer a growth mechanism which has been used to build up synthetic networks with similar properties. These networks have been used as test-beds for the implementation of a model of traffic dynamics, with the aim of assessing the ability of the Internet’s topology to support the basic actions for data traffic handling. Results have been compared with those obtained by using a random network of similar size. The effects of some structural perturbations (arcs and nodes’ removal, traffic localization) have been also evaluated in terms of the induced variations of the network’s efficiency. The resulting scenario is consistent with the hypothesis that the structure of the Internet is only partially fit to host communication processes and that the intelligence of the TCP/IP protocol is partly needed to overcome some “structural” deficiencies.

© 2007 Elsevier B.V. All rights reserved.

PACS: 89.75.-k; 89.20.Hh

Keywords: Complex networks; Topology analysis; Traffic simulation

1. Introduction

The Internet is among the most interesting examples of self-organized complex systems available. Since the appearance of seminal works [1,2] and of more recent assessments [3] which have focused attention to the origin and the selective pressure which brings unsupervised self-organized networks to structure themselves into specific

* Corresponding author. Tel.: +39 06 30484825; fax: +39 06 30486511.

E-mail addresses: rosato@casaccia.enea.it (V. Rosato), limor@vwi.tu-dresden.de (L. Issacharoff), sandro.meloni@gmail.com (S. Meloni), daniele.caligiore@istc.cnr.it (D. Caligiore), fabio.europe@gmail.com (F. Tiriticco).

topologies, a great deal of effort has been spent to unveil their relevant properties. The underlying idea is to capture the specific pressure controlling the growth of complex systems, which is able to drive them toward a specified structure with “optimal” structural and functional properties.

A ticklish question could be raised, at this level. Structural robustness and functional efficiency could be associated, in principle, to conflicting structural properties (or, if not conflicting, at least quite different ones). In that case, it would be reasonable to ask how systems manage to solve such conflicts and to what extent the resulting structure represents an optimal trade-off between the two different instances.

The Internet, more than others, has attracted much attention [4–7] for both the intrinsic relevance of the subject and for the availability of data, thanks to the efforts of a number of projects going on worldwide [8,9]. These projects, albeit with intrinsic differences related to the way in which the Internet network is sampled, pursue the unifying goal of providing a world-wide map of the network, at the level of Autonomous System (AS) routers. This information can be used for an accurate (time-dependent) study of the network’s graph configuration (at the topological level), an understanding of its growth and its growth-rate, and for highlighting the variations of other topological quantities.

This work addresses the problem of understanding the relation between the *topology* of a communication network and its function and efficiency. Our goal is to attempt to answer several questions. Aside from the evaluation of the most relevant topological parameters, which have been also extensively reported in previous studies [5–7], we intend to study the different vulnerabilities of the network, both structural and the dynamical. For structural vulnerability, we intend to study the response of the network to some perturbation (i.e. node’s or arc’s removal) in terms of the number of nodes which result in being disconnected (cannot be reached anymore by the other nodes) upon the removal of a given number of arcs or nodes. For dynamical vulnerability, in turn, we intend to explore the response of the functional behaviour of the system (i.e. how much traffic is affected by the removal of arcs or nodes of the network). Starting from this analysis, we attempt to collect relevant insights to answer the question regarding the ability of the Internet to sustain the traffic of data also in presence of structural perturbations and, ultimately, on the fitness of the Internet’s structure to accomplish its task.

The plan of the present work is the following: we firstly focus our attention on the evaluation of the relevant topological properties of the networks. Then we set up of a simple model for the simulation of traffic dynamics running on a network and evaluate the variation of the traffic properties when the network develops faults (random removal of arcs and nodes) or deliberate attacks (removal of particularly relevant nodes or arcs), as well as in the case of “localized”-type communications. Traffic dynamics are produced by a much simpler model than the TCP/IP stack, in order to better highlight the effect of the topology on the functionality of the transmission of data.

1.1. Internet data, its topological model and properties

Data have been extracted from the repositories of the DIMES and the ROUTEViews projects [8,9]. For a comprehensive review of the methods applied for obtaining the data, see Refs. [10,11] for the two projects, respectively. The DIMES data used for the present study refers to the snapshot taken on July 2005; the ROUTEViews data refers to the snapshot of May 26, 2001. The comparison between the two data sets has been used to assess if the different methods used to sample the Internet space produce similar results.

The available raw data allows us to map the network on a graph $G = G(N, E)$, with N nodes and E arcs connecting the nodes, represented by an Adjacency matrix \mathbf{A} ($A_{ij} = 1$ if nodes i and j are linked, 0 elsewhere). The graph represents the connections present between nodes (AS-level routers). Arcs represent the physical connections between routers (optical fibers, cables etc.). They are bi-directional, as they allow the flux of data in both directions. At this stage, the topological analysis of the network graph allows us to get a first set of information on the properties of the network. In this work we have evaluated:

- (1) the distribution of nodes’ degrees, which allows us to ascribe the graph topology to some specific class;
- (2) the clustering coefficients of each node c_i , their distribution and correlation with each node’s degree:

$$c_i = \frac{n_c(i)}{n_i(n_i - 1)} \quad (1)$$

where n_i is the number of first neighbours of node i and $n_c(i)$ is the number of connections existing between these nodes;

- (3) the network’s diameter and the distribution of nodes’ distances;

Table 1
Major topological properties estimated on the DIMES and the ROUTEVIEWES networks

| | N | E | γ | c | k_{\max} | d | $\langle d_{ij} \rangle$ |
|-------------|--------|--------|----------|------|------------|-----|--------------------------|
| DIMES | 14 154 | 38 928 | 2.41 | 0.41 | 1932 | 9 | 3.343 |
| ROUTEVIEWES | 11 461 | 32 730 | 2.35 | 0.35 | 2432 | 9 | 3.565 |

N is the number of nodes, E the number of links, γ the exponent of the power law, C the average clustering coefficient, d the network's diameter and $\langle d_{ij} \rangle$ the characteristic path length.




Fig. 1. Graph obtained by representing the DIMES network. Plot produced by using the package AiSee (www.absint.com).

(4) the evaluation of the correlation between node distances and degrees [12].

The evaluation of other properties of the DIMES network related to the so-called k -core decomposition has been performed in a recent work [13–15]. In Table 1 we summarize the results obtained for these properties of the networks: aside from the network's dimensions (N , E), we report the fitted values of γ (exponent of the power-law reproducing the node's degree distribution), the average clustering coefficient C , the network's diameter d , the average inter-node distance $\langle d_{ij} \rangle$ (also known as “characteristic path length”). A plot of the network resulting from DIMES data is reported in Fig. 1.

The most relevant information about the network is its scale-free character, indicated by the linear behaviour of its distribution of node degrees in a log–log plot (Fig. 2). The best fit of the linear portion of the distribution allows us to estimate a value of $\gamma = 2.41$ in the scale-free type distribution of node degrees of the type $D(k) = k^{-\gamma}$. Albeit similar, the two values of γ do not refer to equal curves: the log–log plot of the ROUTEVIEWES degree distribution is characterized by a not perfectly linear behavior, with a “depletion” region at intermediate degree values. The estimated value of $\gamma = 2.35$ is a trade-off between the first and the second linear regions of the curve, disregarding the non-linear region.

A further interesting feature is related to the average clustering coefficient C , which is very high for both networks. Fig. 3 reports the distribution of node clustering c_i , which shows a large fraction of nodes having intermediate and high clustering values. This fact has been already noticed: the Internet produces a Scale-Free network with “extreme” properties, such as hubs of very high degree, many leaves and huge clustering coefficients.

It is worth noticing the low network's diameter ($d = 9$ in both cases) and the low characteristic path length $\langle d \rangle$, which in both cases is of the order of $\langle d \rangle \approx 3.4$. The latter quantity is of much concern for communications, as it represents the average distance (in an AS-metric) that a data packet has to travel before reaching its destination, starting from the position of a generic sender node. In Fig. 4 (top) we show the average distance $\langle d \rangle$ of couples with




Fig. 2. Log–log plot of the distribution of node degrees for the DIMES network. The linear section is approximated by a k dependence $k^{-\gamma}$ with $\gamma = 2.4$.




Fig. 3. Distribution of the node clustering coefficients for the DIMES network.

a given degree product $k_i k_j$; this follows a linear behaviour in a semi-log plot, as recently suggested by [12]. In Fig. 4 (middle) we set k_i to 1, thus considering only leaves nodes in the network that are clearly closer to hubs than somewhere between them; in Fig. 4 (bottom), $k_i = k_{\max}$ as we only consider the greatest hub, which turns out to be surprisingly close to the other big nodes in the network (even being a neighbour to many of them).

We have attempted to summarize these observations into a growth mechanism able to reproduce (at least) the two most relevant properties of the network, such as its peculiar scale-free character and its high clustering. The extremely high clustering cannot be properly accounted for by using a straightforward Preferential Attachment [2] (PA) mechanism (which would produce a much lower clustering); such a PA must be complemented by the so-called “Triad Formation” mechanisms (TF) [16], which enhance the resulting clustering. The proposed mechanisms are able to reproduce at least the most relevant topological properties of the Internet and can be stated as follows: wishing to create a network by adding new nodes, each of them with m_0 links, given an initial seed of $N_{\text{init}} = m_0$ nodes, a new node is added with the following rule:

- its first link is added according to a modification of the PA mechanism, where the probability p_i of linking a new node to the node i , belonging to the pool of n already established nodes, is given by [17]:

$$p_i = \frac{k_i^\alpha}{\sum_{j=1}^n k_j^\alpha} \quad (2)$$




Fig. 4. Distribution of the mean distances d_{ij} between pairs of nodes i and j as a function of the product of their degrees k_i and k_j . Top: k_i assumes all degree values in the network. Middle: k_i is set to 1. Bottom: k_i is set to 1977.

- the remaining $m_0 - 1$ links are added by choosing PA or TF with a probability $P(\text{rule})$:

$$P(\text{rule}) = (1 - q)PA + qTF. \tag{3}$$

The proposed growth mechanism provides a network with properties similar to real Internet networks when $\alpha = 1.44$ and $q = 0.93$. This attempt follows a previous one of other authors [18] who identified a growth mechanism

able to reproduce the main features of the Internet's topology based on a weighted PA mechanism where connections between geographically distant nodes were discouraged.

A recent work which has been brought to our attention [19] has introduced a multivariate statistical method to “classify” growth models of the internet as a function of their predictions of the different topological properties really shown by the network. The work defines a number of sets of topological observables which can be used to compare the properties of synthetic networks, grown from different growth mechanisms, with respect to those exhibited by the real data on the Internet's structure. We have performed the analysis proposed by [19] on a synthetic network generated by the mechanism of Eq. (3) and compared the results with those obtained on the DIMES network. The comparison, obtained by using $\langle k \rangle$, $\langle d_{ij} \rangle$ and $\langle c \rangle$ as the set of topological quantities used for the multivariate analysis, has provided a check of the ability of the proposed mechanism to capture the structure of the Internet network. However, the attempt made to define a novel growth mechanism in this work should be only considered as a necessary task to allow the setting up of small synthetic networks, with characteristics as similar as possible to those of the true Internet structure, to perform traffic simulations.

A further insight into the network's topology can be obtained by performing the spectral analysis of its main associated matrices, such e.g. the Adjacency A and the Laplacian L matrices. The latter is defined as

$$L_{ij} = \begin{cases} \sum_{j=1}^N A_{ij} & \text{if } i = j \\ -A_{ij} & \text{if } i \neq j. \end{cases} \quad (4)$$

An interesting result which can be obtained by the spectral analysis of the network's Laplacian matrix L can be stated as follows: the signs of the components of the eigenvector associated with the first non-vanishing eigenvalue of the Laplacian allow us to optimally bisect the network [20–22]. As L is symmetric, the first eigenvalue is always vanishing. The n components of the eigenvector $\mathbf{v}_2^L = (v_1, v_2, \dots, v_n)$ associated with the second eigenvalue solve the one-dimensional *quadratic placement* problem of minimizing the function

$$z = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (v_i - v_j)^2 A_{ij}. \quad (5)$$

The vector is subject to the constraint $|\mathbf{v}| = (\mathbf{v}^T \mathbf{v})^{\frac{1}{2}} = 1$ [20]. The process allows us to partition the graph $G = (N, E)$ into disjoint subsets N_1 and N_2 such that $L_{12}/(N_1 \cdot N_2)$ is minimized. L_{12} is the number of links to be removed and N_1, N_2 are the number of nodes in the two resulting subnetworks. It becomes clear that this procedure allows the “optimal” bisection of the graph, i.e. it forms the closest possible subnetworks N_1 and N_2 with the minimum number of broken links L_{12} , as the function $L_{12}/(N_1 \cdot N_2)$ is minimum when L_{12} is small and $N_1 \simeq N_2 \simeq N/2$.

If we apply this result to the two Internet maps, we obtain: $N_1 = 13\,534$, $N_2 = 620$ and $L_{12} = 675$ for the DIMES network, $N_1 = 11\,187$, $N_2 = 274$ and $L_{12} = 274$ for the ROUTEVIEWS network. This algorithm can be iterated to bisect the resulting subnetworks. If we bisect the largest subnetwork, we will find a similar situation: the largest subnetwork is bisected into two further networks where one is much smaller than the other ($N_{11} = 13\,030$, $N_{12} = 504$ and $L_{112} = 2407$ for the DIMES network, where the index 11 means the first subnetwork obtained by the previous subnetwork of size N_1 etc.). If, in turn, we grow a network of similar size ($N = 14\,154$) by using the growth model presented in Eqs. (2) and (3), we find different bisections depending on the type of networks that we grow. In the case of the bisection of a high-clustering scale-free network, we obtain results qualitatively similar to those obtained for the DIMES network (i.e. one of the two subnetworks is much larger than the other).

The present result opens the way to some possible interpretation of the Internet's data. The low value of the quantity L_{12} for both the networks (or, equivalently, the fact that, upon bisection, one of the two subnetworks is much smaller than the other) indicates that the networks are composed of “tiles” of intermediate sizes (from a few hundreds to some thousands nodes), stuck together by a moderate number of links. This can be either the result of the mapping strategy adopted by the two projects, or a peculiar feature of the Internet's actual structure. The visual inspection of the graphs obtained by the two networks (shown in Fig. 1) does not help to clarify this point, which thus deserves further investigation.




Fig. 5. Schematic layout of the model: circles are the nodes, data packets travel on the links, each node owns a buffer of unlimited size.

2. A dynamical model of traffic flow on the Internet

A relevant literature is available for the dynamical study of communication networks, focussing on both technological and fundamental aspects (see, for instance, [23]). We have also attempted to analyze the functional vulnerability of the network, intended as the impact that several types of perturbation might produce on the network behaviour. To this end, a dynamical model of the communication network has been used in a “what-if” mode in a way so as to answer the following questions: What will happen if some structural or functional perturbation is applied to the network? What is the susceptibility of the system to a perturbation?

To this purpose, we have firstly defined a *dynamical* model of traffic flow, able to reproduce the functional behaviour of the network, seen as a route where packets of data can be exchanged (send/received) among the different nodes. This should be regarded as a qualitative “behavioural” model; far from being a model of all the actions performed by the Internet protocol stack, it wishes to select out those that are the effects of the network’s topology on the process of data transmission from one node to another, simulated with a basic set of actions needed to accomplish this task. The dynamical model will thus be able to reproduce the traffic of packets of data, generated with a given frequency by a node and directed to another node of the network.

We have defined several models of the network’s dynamics, with increasing levels of complexity. The basic one, referred hereafter to as *homogeneous*, is characterized by the fact that all nodes of the network (representing the AS-level routers) are supposed to have equal technological properties *independently* of their degree. In a further model, hereafter referred to as *heterogeneous*, we remove such a constraint: as it would be reasonable to expect, nodes (i.e. the routers) have technological properties *dependent* on their degrees (hubs should be able to perform a larger number of operations per unit time than routers of lower degree).

Simulations have been carried out over a fixed number of *Time-Steps* (TS). The basic actions that a node might perform, at a given time, are:

- send one or more (depending on the adopted model) packets of data to a neighbour node.
- receive packets of data from neighbour nodes.

A node cannot send a packet of data to itself. A packet of data is supposed to be infinitesimal in size; this constraint has been imposed to avoid dealing with the finite capacity of links. A data packet will contain the values of two different quantities: the time of emission and the destination node. Both these pieces of information will be used to direct the packet throughout its journey until the destination is reached (see Fig. 5). At each TS, a packet can traverse only one link.

Both models (*homogeneous* and *heterogeneous*) share basic features which might be summarized as follows:

- infinite buffers; each node has a buffer able to contain an infinite number of packets (the infinitesimal size of the data packets is also relevant for ensuring this property) received by a node from its neighbours, at different times. In the buffer, the data packets are queued before being dispatched, according to the given routing strategy;
- routing tables (RT); each node hosts a routing table $RT_i = R_{ik}$ which associates, with each destination node k , a pair of adjacent nodes j_1 and j_2 , to which the transit packet could be directed (if the i node is a leaf, then $j_1 \equiv j_2$).

RT's are defined according to the minimum-distance path between nodes, evaluated (once and forever) via the Dijkstra algorithm. The choice between the direction j_1 or j_2 is made according to the rule [24]:

$$P(j_1) = \frac{e^{-\beta X_{j_1}}}{e^{-\beta X_{j_1}} + e^{-\beta X_{j_2}}} \quad (6)$$

$$P(j_2) = \frac{e^{-\beta X_{j_2}}}{e^{-\beta X_{j_1}} + e^{-\beta X_{j_2}}} \quad (7)$$

where X_k , ($k = j_1, j_2$) is the number of data packets sent to node k and $P(j_1) + P(j_2) = 1$. The parameter β can be used to obtain several types of routing strategies: a **probabilistic** routing if $\beta = 0$ (the two directions have equal probability to be chosen); a **deterministic** routing if $\beta = 1$ (the two directions are chosen according to the previous traffic registered on them) and a **fixed** routing if $\beta = \infty$ (only the first direction is chosen).

- routing policy. We have implemented a FIFO policy for buffer management (FIFO stands for *First-In First-Out*, the first arrived packet will be the first to be dispatched).

The two dynamical models differ in the technological properties associated with each node: in particular

- in the *homogeneous* model, the number of operations that a node is able to perform in a TS is equal for all the nodes, independently of their degrees: at a given TS, a node can receive an unlimited number of packets (which are then queued in the buffer) and send one (and only one) packet to a neighbour node. This packet can be the result of an emission event (a new packet is generated from that node) or of a dispatching event (a packet is removed from the buffer, according to the FIFO policy, and sent to a neighbour, according to the RT prescription). The *homogeneous* model will be always simulated with $\beta = \infty$ (the same direction is always chosen). It will thus represent the very basic model of data transmission.
- in the *heterogeneous* model, the number of operations which a node is able to perform in a TS depends on its degree, allowing nodes to send more than one packet in the same TS. To relate the number of packets to the topological properties of the node, we have adopted the following rule: if the node i is characterized by the degree k_i such as

$$k_i > k_m b \quad (8)$$

where k_m is the average degree of the network and b is a tunable parameter, we assume that it can perform more than a single send operation per TS (it will thus be called a “power router”). In that case, the number of send operations θ_S it can perform per TS is

$$\theta_S = \frac{k_i}{k_m} a. \quad (9)$$

The free parameters a and b presented in Eqs. (8) and (9) can be suitably adjusted to realize different network capabilities: with $b \ll 1$ all routers are considered “power routers”, while for $b \gg 1$ only a very limited number of hubs will have enhanced capabilities. The value of a will, in turn, specify the extent of the increased power attributed to the router (the larger a , the higher the number of send operations which the power router will be able to perform). The *heterogeneous* model will be simulated by using different values of β , in order to check the improvements obtained by adding further “intelligence” to the data transmission strategies.

The amount of traffic present in the network is measured by the variable λ , which measures the frequency with which nodes emit a packet of data. Each packet is generated by a randomly chosen node and directed to a randomly chosen destination node. With this definition, for instance, $\lambda = 0.1$ represents a level of traffic where, at each TS, 10% of the N nodes of the network generate packets of data directed toward an equal number of destination nodes.

Simulation starts with an “empty” network, i.e. with no traffic and with all buffers empty. As soon as the simulation progresses, buffers start filling and emitted packets take a certain time (i.e. a certain number of simulation TSs) to reach the destination node. In fact, once emitted, a packet of data jumps from one node to another (according to the Routing Tables of each traversed node and subjected to the adopted routing strategy) before reaching the destination node. In the best possible case, the packet is immediately re-emitted by a (non-destination) node the TS after its arrival at that node. However, if the buffer of that node is already filled by packets that arrived at earlier times, the packet must remain on that node for a number of TSs (which is equal to the number of packets waiting in the buffer that arrived at earlier times), according to the FIFO buffer policy.




Fig. 6. Average delivery time for the DIMES network, when the *homogeneous* model ($\beta = \infty$, deterministic routing) is applied.

In the end, the packet will reach the destination node. The measure that provides a significant estimate of the effects of the traffic level in the network is the average time $\langle T \rangle$ needed for travelling from the emitted node i to the destination node j , via the minimum path connecting node i to node j . If t_k is the time needed for the packet k to travel from its emitting node its the destination node, then the required value of $\langle T \rangle$ is

$$\langle T \rangle = \frac{1}{M} \sum_{k=1}^M t_k \tag{10}$$

where M is the number of packets which have been effectively *delivered* within the simulation time (the units of $\langle T \rangle$ are simulation time steps TS). This definition considers the fact that, as we assume a finite simulation time, at its completion, a fraction of emitted packets will have been not yet delivered and will still be travelling toward their destination nodes. Therefore the average of Eq. (10) is evaluated only for packets which have completed their route up to their destination nodes during the simulation time.

The results of our simulations on the DIMES network are reported in Fig. 6, where the *homogeneous* model ($\beta = \infty$ deterministic routing) has been assumed.

The behaviour of a network can be resumed as follows: if we identify the $\langle T \rangle$ value as the relevant functional outcome of the network, its behaviour as a function of the data traffic presents two distinct phases: the first, hereafter called the *normal* phase, is characterized by a linear increase in the value of $\langle T \rangle$ as a function of the traffic λ . A second phase, called hereafter *congested*, occurring above a given traffic threshold λ_c , is characterized by a marked non-linearity in $\langle T \rangle$ versus λ . This is a general behaviour which has been detected in several types of networks.

The overall behaviour of data traffic on the DIMES network of Fig. 6 is qualitatively similar to that detected on other network types. The value of λ_c defining the onset of the congested phase is $\lambda_c = 2.5 \times 10^{-4}$ (also known as the point of *jamming* transition [23–25]). The *normal* phase is, in turn, characterized by a very small value of $\langle T \rangle = 3.5$ which is of the order of the characteristic path length of the network (see Table 1), which hardly increases with the traffic value for $\lambda < \lambda_c$. There are several factors which influence the position of λ_c : the first is the network’s size (in terms of number of nodes), the second is its topology. We have attempted to assess these dependencies by performing two types of simulation: the first, for measuring the size effect and the second to compare the results on the Internet network with those of a network of equal size belonging to the topological class of random networks.

2.1. Size effects

We have applied the recipe (Eqs. (2) and (3) above) with $\alpha = 1.44$ and $q = 0.93$ to grow “synthetic” networks of sizes $N = 500$ up to 10 000, on which we have applied the *homogeneous* dynamic model (deterministic routing). These networks share with the real DIMES network the values of γ and c (see Table 1). The curve expressing $\lambda_c = f(N)$ can be suitably fitted to a function $\lambda_c \sim 2/N$ (Fig. 7). Note though that the same behaviour holds also if other routing strategies are adopted with different values of the λ_c range.




Fig. 7. Variation of the value of λ_c as a function of the network size N for synthetic networks sharing the values of the parameters γ and c with the DIMES network.




Fig. 8. Average delivery time for a DIMES-like (circles) and a random network (squares) with $N = 5000$ nodes. In both cases, the *homogeneous* model (deterministic routing) has been adopted.

2.2. Comparison with random networks

In order to compare the behaviour of the DIMES network with that of a random network of similar size, we have grown a synthetic random network and compared its behaviour to that of a synthetic network having the same topological properties as DIMES. We have used two $N = 5000$ nodes networks and used the *homogeneous* ($\beta = \infty$ deterministic routing) model. Results are reported in Fig. 8. It becomes clear that the random network (squares) has a higher value of λ_c than the DIMES-like network, while it shows a slightly higher value of $\langle T \rangle$ in the *normal* phase. These two sets of data can be interpreted as follows: the largest hubs in DIMES-like network represent a dramatic bottleneck for communications. As long as they are treated on an equal footing with respect to nodes of lower degrees (equal functions), because they are also *central* nodes and most of traffic should travel through them, they tend to “trap” most of traffic in their buffers, which start filling at higher rates than they can empty themselves. This effect is, in turn, less relevant in random networks, where nodes are nearly equivalent (in terms of degree) and, more importantly, also each node’s *centrality* is similar. As such, differently from what happens in DIMES-like networks, the traffic makes use of a larger number of nodes; this enables it to avoid congestion until higher values of λ . The DIMES-like network, in turn, displays a lower $\langle T \rangle$ in the *normal* phase with respect to *random* networks, as the average distance in the former is lower than in the latter. As such, in low traffic conditions, data packets can be delivered faster (in average) on the DIMES network than in a *random* one.




Fig. 9. Top: comparison of the average delivery time $\langle T \rangle$ for DIMES-like synthetic networks ($N = 5000$): *homogeneous* model (circles, $\beta = \infty$ deterministic routing) and *heterogeneous* model ($\beta = \infty$) with different values of the b parameter ($b = 6$ squares, $b = 20$ diamonds, the value of the a parameter is fixed at $a = 1.5$ in both curves) see Eqs. (7) and (8). Bottom: comparison between the $\langle T \rangle$ for DIMES-like synthetic networks ($N = 5000$) with three different *heterogeneous* models with $\beta = 0$ (triangles), $\beta = 1$ (diamonds) and $\beta = \infty$ (squares). All the *heterogeneous* models have been produced by setting $a = 1.5$ and $b = 6$ in Eqs. (8) and (9).

2.3. Comparison of different models

To produce a more complete assessment of the properties of the network when hosting data traffic, we have simulated the time behaviour of the model representing the DIMES network in different technological conditions. To this end, we have grown synthetic models of DIMES-like structures with $N = 5000$ nodes and performed traffic simulations with the following dynamical models:

- (1) **homogeneous routers** (i.e. they can send just one packet of data each TS) and **fixed routing** ($\beta = \infty$);
- (2) **heterogeneous routers** (i.e. the routers can send a number of data packets out at each TS, which is a function of their degree, see Eq. (9)) and **deterministic routing** ($\beta = 1$);

Fig. 9 reports the behavior of $\langle T \rangle$ with λ for a DIMES-like network ($N = 5000$) in the two cases. From Fig. 6, we could predict that the critical traffic value for a $N = 5000$ nodes network would be at $\lambda_c \sim 4 \times 10^{-4}$. In fact, once the DIMES network has been used to model the network graph, the resulting behaviour of T with respect of the traffic value λ , reported in Fig. 6, shows the onset of the congestion transition at $\lambda_c = 4.5 \times 10^{-4}$.

Fig. 9 (top) shows the effects of the parameters a and b of the heterogeneous model Eqs. (8) and (9) in the improvement of traffic delivery when compared to the homogeneous case (shown for comparison purpose). In all cases $\beta = \infty$. The figure shows that, among the *heterogeneous* cases, that the one with the smaller b value (squares, in the figure), outperforms with respect to the other case ($b = 20$, diamonds). In fact, when b is small (case $b = 6$), many hubs become “power” routers and the effect is a dramatic upward shift of the critical traffic value λ_c (of almost




Fig. 10. Behaviour of $\langle T \rangle$ in a synthetic $N = 3000$ node network ($\beta = 1$, homogeneous model) upon removal of the most used links. Points: no removal. Circles, squares and diamonds: removal of 10, 20 and 100 most used links. The small inset shows a magnification of the behaviour of $\langle T \rangle$ in the non-congested phase (values of $\langle T \rangle$ are reported on the right axis of the inset).

ten) as many *central* nodes of the network have a larger power available (they can send more than a single data packet per TS). This has the overall effect of allowing a slower filling of the buffers and, ultimately, of delaying the transition to the *congested* phase. The same effect occurs when, b is kept fixed while a increases: in such a case, the power routers become more and more powerful (i.e. they can perform more send operations in a given TS).

Fig. 9 (bottom) shows the possible benefits attainable by the introduction of the more complex routing strategies given by the different values of β . In that figure, aside from the $\langle T \rangle$ behaviour in the *homogeneous* ($\beta = \infty$ deterministic) model chosen as our reference, we report the behaviour of the synthetic $N = 5000$ network, for given values of the a and b parameters of Eqs. (8) and (9) (namely $a = 1.5$ and $b = 20$), when varying the β value.

2.4. Vulnerability assessment

In order to perform an assessment of the *functional* vulnerability of the network (intended as the effects of a structural perturbation on the network's functioning), we have resorted to using small, synthetic networks generated via the growth mechanism described in the previous section.

A vulnerability assessment consists in generating, for a given perturbation strength (measured in terms of simultaneously removed network's elements ξ), a large number of network configurations (each characterized by a different choice of the removed elements) and, for each configuration, performing a functional simulation (at different values of traffic level λ) in order to observe the network's response in terms of the critical value λ_c , the value of the average delivery time for the normal and the congested phases, etc. The estimated effects of the perturbation on the network will be provided by an average over the results obtained for the different simulated configurations. To achieve significant statistics we have used a large (several thousands) number of perturbed configurations; the use of synthetic networks (of dimensions smaller than those of DIMES) has been thus driven by the goal of reducing the total computing time. The use of larger networks would have produced qualitatively similar results, the only variation being in the value of λ_c .

In order to assess the networks' vulnerabilities, we have performed two types of "experiments", with different types and numbers of removed elements.

In the first, we have perturbed the structure of the network by removing (either randomly or through a deliberate choice) nodes or arcs.

In order to simulate a deliberate attack to the network, we first observed its dynamic behaviour with no perturbations, in order to have a list of the most used links in the packet routing process. We then gradually proceeded with the removal of these nodes in a subsequent traffic simulation. Under any configurations of model and routing, the network shows an improvement in terms of λ_c , which moves toward higher values in Fig. 10 as the number of removed links increase: this means that the network can handle more traffic before entering the congested phase. As its counterpart, the average delivering time in the normal phase increases.




Fig. 11. Behaviour of $\langle T \rangle$ in a synthetic $N = 5000$ node network ($\beta = 1$, heterogeneous model, $b = 6$) upon removal of the most used links. Circles: no removal. Triangle, diamonds and squares: removal of the 10, 20 and 50 most used links respectively.

The results are quite astonishing but largely comprehensible if we look to the routers' buffer usage: in the congested phase, hubs actually play the role of *bottlenecks*, with their buffer displaying long queues of packets to be delivered, while lower degree nodes' buffers keep the same size. The removal of highly central links (i.e. those where most of the traffic of data flows) produces a sizeable increase of $\langle T \rangle$ at low traffic values (which is to be expected as, in equilibrium conditions, hubs and central links do not constitute a bottleneck but a structural advantage). For high traffic values, in turn, the elimination of central links provides a relief to the network which can route the traffic differently, by using many more links for the communications. This allows a clear shift in the value of λ_c to higher values. This would represent a sort of increase in efficiency after a fault. Although this could be considered counter-intuitive, it is the logical conclusion of what has been noted above concerning the bottleneck effect produced by high central nodes and links. More insights and detailed graphs can be found in [26].

Such behaviour changes direction when the heterogeneous model is used. Fig. 11 shows the anticipation of the transition phase point λ_c and the decrease of the performance of the network in the normal phase. In fact, when the heterogeneous model is used, the arcs that interconnect hubs are much more stressed than others: their removal forces the network to exploit alternative paths that pass through non-enhanced routers, thus creating a greater bottleneck problem.

The random removal of links doesn't produce relevant effects on the dynamic behaviour of the network, unless a high number of arcs is removed. This proves the robustness of this topology against random faults from the functional point of view as well.

The process of node removal is similar to that of link removal, as removing a node is equivalent to cutting its links with the neighbour nodes. The targeted node removal does not bring in any interesting results: due to the scale-free nature of the network, the removal of a high-degree node always brings the disconnection of the network, thus confirming its characteristic poor robustness to targeted attacks. At the same time, the random removal of nodes is as ineffective as the random removal of links.

In the second type of perturbation experiment, in turn, we have carried out a different type of vulnerability assessment, not related to element's removal but consisting in simulating the behaviour of traffic under "pathological-type" functioning. To this end, we applied a "functional" perturbation, consisting in forcing the network to sustain a "localized" traffic. In normal operational conditions, traffic can be assumed to be "delocalized": each node sends a message to a randomly chosen node belonging to the whole set of $(N - 1)$ nodes. When traffic "localizes", in turn, nodes are allowed to send a message to a node randomly chosen as belonging to a predefined pool of nodes (which in general, is a small subset of the set of all nodes). These pools, as defined for the different simulations which have been performed, consist of an intermediate-size hub with all its neighbours (to simulate the "geographical" localization of communications). This could mimic, somehow, the effect of a deliberate cyber-attack on a given pool of resources, or the situation encountered in communications when specific regions (or countries) undergo some catastrophic event that breaks the homogeneity of traffic flow, which mostly directs toward the routers of those regions. We have simulated three different events: (a) localization on a large hub (we choose one of the major hubs present in the network, having




Fig. 12. Behaviour of the average delivery time (T) of a synthetic $N = 5000$ nodes network (heterogeneous model, $\beta = 1$ with $a = 1.5$ and $b = 6$) in cases of traffic localization. Circles: traffic localized toward a large hub with $n_1 = 206$ nodes ($S = 0.0103$). Squares: traffic localized on a smaller hub containing $n_2 = 77$ nodes ($S = 0.00385$). Diamonds: traffic localized on two separate hubs, one containing $n_3 = 109$ and the other $n_4 = 100$ nodes ($S = 0.0207$, this result could be compared to the case of the larger hub, in terms of the number of nodes). Points: traffic not localized.

a number of receiving nodes of $n_1 = 206$, (b) localization on a smaller hub (with $n_2 = 77$ nodes), to increase even more the localization on a small number of nodes, and (c) localization on two separate hubs, whose total number of receiving nodes are similar to the large-hub case (in fact, in this case, the two hubs have $n_3 = 109$ and $n_4 = 100$ nodes each: two nodes are shared, resulting in a total number of nodes $n_{(3+4)} = 207$). The difference between the case (a) and the case (c) consists in the specifics of their spatial localization; in fact, in case (a) the nodes are all contiguous and share the same central node; in case (c) nodes are in two separate pools, thus reducing the spatial localization.

The extent of localization of the traffic could be measured by defining a quantity S as follows

$$S = \frac{n(R)}{N} \frac{d(R)}{d} \quad (11)$$

where N and d are the total number of nodes and the network's diameter, respectively; $n(R)$ and $d(R)$ are the total number of receiving nodes and the diameter of the receiving pool of nodes. Note that it is always possible to define the value of $d(R)$, even if the pool of receiving nodes is not connected (i.e. it is composed by two not-connected subpools, as in our case). It is, in fact, possible to define a distance between nodes even if a part of the path is formed by links which do not belong to the pool (under such a definition, the lower the value of S , the higher the localization).

Results are reported in Fig. 12. It is evident how the higher the localization, the lower is the network's performance. In fact, the buffer of the hub involved in the reception of all the messages will rapidly fill, and despite its ability to dispatch more packets per TS (i.e. the values of $a = 1.5$ and $b = 6$ mean that it is a power router able to route a certain number of packets per TS), the values are $\lambda_c = 0.03$ and $\lambda_c = 0.015$ for the cases with $n_1 = 206$ and $n_2 = 77$, respectively. λ_c increases linearly with S .

3. Conclusions

This work has addressed the problem of understanding the relationship between the *topology* of a communication network and its function and efficiency.

Recent works have highlighted the impact that the structure of a network (formed by a self-assembly process in an unsupervised regime) might have on its functional properties (the way it works, its efficiency etc.). In many of them (protein interaction networks in living cells, in social networks etc.), the network's function seems to be *supported* by the structure. In other words, the network's topology helps the system to behave in a highly efficient way. The topological structures of these networks, moreover, allow the achievement of a high *robustness* (small propensity of being damaged by random faults).

We have investigated, in some detail, these issues on a model system representing the Internet, using the results of the DIMES project, one of the most recent efforts to map the global structure of the Internet.

In the first part of the present work, we have analyzed the Internet map under the topological point of view, in order to extract relevant insights on its structure. This analysis has allowed us to underline two main facts:

- (1) the “extreme” Scale Free character of its structure (coexistence of large hubs and many loosely connected nodes called *leaves*), including a high clustering (presence of a large number of three-vertex structures), leads the Internet network to display a low average inter-node distances (functional to transmission of data), providing at the same time good resilience to random (i.e. un-targeted) faults.
- (2) the actual map of the Internet, issued from the DIMES project, has a “tiled” structure. Its spectral analysis and the min-cut theorem have shown that the DIMES network is not a single, highly connected, structure but rather a number of large, highly clustered regions which merge one into the other through weak boundaries (i.e. with a low number of interconnecting links). This fact, which should be carefully investigated, can be either a specific tract of the Internet structure or a consequence of the method employed, in the DIMES project, to create the map. Further measurements performed on a different world-wide Internet map, provided by the US-funded RouteViews project, have confirmed the possibility of a “tiled” structure.

What can be inferred from the dynamical analysis (through the traffic flow simulation) and vulnerability tests, is that whereas the Internet’s topology ensures robustness at the structural level (there is a low probability that a random fault produces a high level of structural damage to the network), it does not allow the attainment of an equal efficiency for traffic flow. A major result of this work is the fact that, under the simple flow model implemented to reproduce data traffic, the critical traffic value λ_c for an Internet-like network results in being *lower* than that a random network of the same size. It means that an Internet-like network, under the action of a simple model of traffic flow, reaches the congested phase prior to a random network of equal dimensions. This has been ascribed to the specific topology of the Internet: if large hubs, on one side, are responsible for the high robustness and for the shorter mean inter-node distances, that ensures shorter delivery times in the normal phase; on the other hand, hubs constitute a bottleneck for communications. Nodes and links, in scale-free type networks such as the Internet, are highly non-homogeneous (in terms of centrality, for instance), while in random networks all nodes are practically equivalent over all points of view. This asymmetry is the principal cause of the fact that, under the hypothesis of the equal technological power of all the nodes independently of their degree, the routers of the most central nodes start filling at rates higher than those with which they can unload.

The picture which emerges from our results can be summarized as follows.

The Internet, as with many of the complex systems which self-assemble in an unsupervised-growth condition, is *compelled* to adopt a growth mechanism which inevitably produce a scale-free type network. In the Internet’s case, it is indeed a combination of several mechanisms (the basic one being the Preferential attachment one), whose global result is the realization of an “extreme” scale-free network. This specific growth mechanism depends on the type of forces which produce the growth: the creation of communities (large hubs), the presence of triangles which settle down the communities etc. This structure, however, although it is robust, is not always appropriate for traffic flow.

This probably represents one of the major driving forces for the development of efficient intelligence strategies, which have been developed to overcome the functional limitations introduced by the Internet’s topological structure.

Intelligent strategies are indeed widely used in the Internet to reduce congestion. Among them, the TCP’s traffic congestion mechanism is one of the best countermeasures the network’s intelligence offers against its bottlenecks problem. If we imagine the protocol being applied to our networks, the strict control imposed by TCP provides a “global” threshold of traffic. This countermeasure is equivalent to the “perception” of an un-sustainable traffic level (which is indeed similar to the phase transition point that has been discussed in this work); when the traffic intensity goes behind that, TCP reduces the network emission of packets, thus restoring the traffic to a level under the threshold, where data can be delivered with the best possible performance.

The limitations introduced by the network’s topology are, indeed, partially overcome by introducing more efficient routing strategies. The “central-links” bottleneck (most minimal paths pass through the same nodes and links, which unavoidably collapse under high traffic conditions) is partially removed by wiser routing strategies, such as that allowing one to choose between several paths for forwarding a packet and by relating this choice to the density of traffic previously delivered along those paths. Recent works have addressed the issue of designing adaptive routing strategies to allow the network to be *aware* of the approaching congestion and to take specific countermeasures [23]. We are also going to explore adaptive routing schemes, whose efficiency in preventing congestion will be reported in a future work.

Acknowledgments

This work has been performed in the frame of the ongoing EU project “IRRIIS” (Integrated Risk Reduction of Information-based Infrastructure Systems) under the IST programme of the Sixth Framework Programme (FP6-2005-IST-4). All partners of the project are gratefully acknowledged for useful discussions and comments. The authors also acknowledge useful discussions with Paolo Palazzari (ENEA and Ylichron), Ingve Simonsen (Technical University of Dresden) and technical support from Rocco Casilli (Ylichron).

References

- [1] D.J. Watts, S.H. Strogartz, *Nature* 393 (1998) 440.
- [2] R. Albert, A.-L. Barabasi, *Rev. Modern Phys.* 74 (2002) 47.
- [3] S. Boccaletti, et al., *Phys. Rep.* 424 (2006) 175.
- [4] R. Pastor-Satorras, A. Vespignani, *Evolution and Structure of the Internet: A Statistical Physics Approach*, Cambridge University Press, Cambridge, 2004.
- [5] M. Faloutsos, P. Faloutsos, C. Faloutsos, *Comput. Commun. Rev.* 29 (1999) 251.
- [6] R. Pastor-Satorras, A. Vazquez, A. Vespignani, *Phys. Rev. Lett.* 87 (2001) 2587011.
- [7] V. Rosato, F. Tiriticco, *Europhys. Lett.* 66 (2004) 471.
- [8] www.netdimes.org.
- [9] www.routeviews.org.
- [10] Y. Shavitt, X. Sun, A. Wool, B. Yener, *IEEE J. Sel. Areas Commun.* 22 (2004) 67.
- [11] Z.M. Mao, J. Rexford, J. Wang, R.H. Katz, *Comp. Comm. Rev.* 33 (2003) 365.
- [12] J.A. Holyst, J. Sienkiewicz, A. Fronczak, P. Fronczak, K. Suchecki, *Phys. Rev. E* 72 (2005) 026108.
- [13] S.N. Dorogovtsev, A.V. Goltsev, J.F.F. Mendes, *Phys. Rev. Lett.* 96 (2006) 040601.
- [14] [arXiv:cs.NI/0504107](https://arxiv.org/abs/cs/0504107) v2 12 Oct 2005.
- [15] [arXiv:cs.NI/0511007](https://arxiv.org/abs/cs/0511007) v2 3 Nov 2005.
- [16] P. Holme, B.J. Kim, *Phys. Rev. E* 65 (2002) 026107.
- [17] P.L. Krapivsky, S. Redner, F. Leyvraz, *Phys. Rev. Lett.* 85 (2000) 4629.
- [18] S.H. Yook, H.W. Jeong, A.-L. Barabasi, *Proc. Natl. Acad. Sci. USA* 99 (2002) 13382.
- [19] L. da F. Costa, F.A. Rodriguez, G. Travieso, P.R. Villas Boas, *Adv. Phys.* 56 (2007) 167.
- [20] L. Hagen, A.B. Kahng, *New spectral methods for ratio cut partitioning and clustering*, *IEEE Trans. Comput.-Aided Des.* 2 (9) (1992) 1074.
- [21] B. Mohar, *Discr. Math.* 109 (1992) 171.
- [22] A. Pothen, H.D. Simon, K.P. Liou, *SIAM J. Matrix Anal. Appl.* 11 (1990) 430.
- [23] P. Echenique, J. Gomez-Gardenes, Y. Moreno, *Europhys. Lett.* 71 (2005) 325.
- [24] T. Ohira, R. Sawatari, *Phys. Rev. E* 58 (1998) 193.
- [25] P. Echenique, J. Gomez-Gardenes, Y. Moreno, *Phys. Rev. E* 70 (2004) 056105.
- [26] F. Tiriticco, *Traffic simulation on a complex topology telecommunication network: analysis of congestion states and vulnerability*, Master Degree final dissertation, available at the University of Rome Tor Vergata.