Tearing Down the Internet

Damien Magoni

► To cite this version:

Damien Magoni. Tearing Down the Internet. IEEE Journal on Selected Areas in Communications, 2003, 21 (6), pp.949-960. 10.1109/JSAC.2003.814364 . hal-00344480

HAL Id: hal-00344480

<https://hal.science/hal-00344480>

Submitted on 15 Apr 2014

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.

Tearing Down the Internet

Damien Magoni

Abstract—Recent advances in scale free networks have claimed that their topologies are very weak against attacks. The inhomogeneous connectivity distribution of large scale current communication networks, such as the Internet, could be exploited by evil hackers in order to damage these systems. However there has not been many studies on the approaches and consequences of such targeted attacks. In this paper, we propose an in-depth study of the Internet topology robustness to attacks at the network layer. Several attacking techniques are presented as well as their effects on the connectivity of the Internet. We show that although the removal of a small fraction of nodes (less than 10%) can damage the Internet connectivity, such a node removal attack would still require a large amount of work to be carried out. To achieve this, we study in detail the interactions between the intra-domain and inter-domain levels of the Internet through the use of an overlay.

Index Terms—Attack, connected component, connectivity, Internet, overlay, topology.

I. INTRODUCTION

SOONER or later everyone will be connected to the Internet. In western Europe, 29% of the households were connected to the Internet in 2001. And this was an increase of 33% over the year 2000. We interact more and more through the Internet and as we get more dependent on it, messages such as "host unreachable" will become more stressful and painful to us. Because the number of people using the Internet and the importance of its use (e.g. for administrative or commercial tasks) will both grow tremendously in the near future, the connectivity of the Internet will become a key factor for productivity. Any connectivity failure will turn into huge profit losses. This characteristic turns the Internet into a potential target for ill-intentioned or terrorist attacks.

Of course there are many levels of connectivity ranging from the physical layer up to the application layer. However a break in any one of these levels will stop all the upper levels from functioning. Thus an attack towards lower levels will do much more damage. In this paper we focus our study of the robustness of the Internet connectivity on the lowest level we can (i.e. given the data at our disposal), namely the Internet Protocol (IP) layer of the Internet. The rest of the paper is organized as follows. In section III, we present the Internet maps that we use in our study. Section V-A discusses the metrics chosen for connectivity measurements. Then we propose several methods for attacking the Internet topology in section IV. Finally, in section V we give measurement results on the Internet connectivity under attack.

II. PREVIOUS WORK

The robustness of the Internet connectivity has not been extensively investigated by the research community. One of

the first study of this kind has been carried out by Albert *et al.* in [1] during the year 2000. They have evaluated the effects of random and targeted node removal upon exponential and scale free networks. They have used generated graphs as well as measured networks. For scale free network examples, they have taken an Autonomous System (AS) level snapshot of the Internet and a World Wide Web (WWW) snapshot. Their metrics for evaluating the effects of node removal include the diameter of the network, the size of the largest cluster (i.e. connected component) and the average size of the isolated clusters (i.e. all clusters excluding the largest one). They have found that under attack (i.e. targeted node removal), scale free networks such as the Internet and the WWW break into small fragments for a threshold fraction of node removal. Under failure, scale free networks keep their largest cluster intact for an unrealistically high removal rate. Another similar work on Internet robustness has been carried out by Tauro *et al.* in [2] during the year 2001 on a snapshot of the AS level topology. They also have used the size of the largest connected component as a metric and have drawn the same conclusions as [1]: the AS network is sensitive to targeted node failures while it is robust to random node failures. Concerning the robustness of the Internet at the router level, Broido *et al.* have studied in [3] the robustness of the IP giant component (a 52505-node map created from the Skitter data) under attack and have measured several properties including the giant component diameter, the largest component size, etc. They have reached the same conclusions for the router level than [1] for the AS level. Also in 2001, Palmer *et al.* tackle in [4] the robustness of the router level map collected by both the SCAN project [5] and the Lucent Internet mapping project [6] which contains approximately 285K nodes. As metrics, they use the number of reachable pairs and the hop exponent. The first metric derives from the neighborhood function. As this function requires a lot of calculation, the authors have proposed an approximate neighborhood function by using data mining analysis. The second metric derives from an approximated power law defined by Faloutsos *et al.* in [7]. With these two metrics, they have investigated random edge deletions and random and targeted node deletions. Consistent with the other studies, they have found that the removal of important nodes (whether by the degree or the hop exponent), severely affects the connectivity of the network while random edge or node removals do not alter it as much and as fast. More recently, Tangmunarunkit *et al.* have investigated in [8] the robustness of an AS map and a router map (both collected in May 2001) by using a graph-theoretic metric called resilience. Our paper is an extension of these previous studies. We go deeper into the Internet robustness analysis in order to find if it is possible to destroy the Internet connectivity and at what cost.

D. Magoni is with the Department of Computer Science, Université Louis Pasteur, Strasbourg, France (e-mail: magoni@dpt-info.u-strasbg.fr)

TABLE I
INTERNET MAPS USED IN EXPERIMENTS

Origin	Nb of nodes	Nb of links	Date
SCAN-LUCENT	284772	449228	1999
LSIIT	188347	235991	4-7/2002
<i>route-views</i>	13529	28060	7/2002

On a more theoretical level, it's worth noticing that several examples of theoretical studies of generated model networks, which give results quite similar to those obtained in this paper for real networks are given by Callaway *et al.* in [9], Cohen *et al.* in [10], which also contains some discussion of the diameter, and Holme *et al.* in [11], which also introduces the effect of correlations. Physicists usually plot the number of cluster of size s , as a cumulative function of s in a log-log scale. This distribution has been studied thoroughly in the physics literature and it is expected that the distribution will follow a power law with slope -2.5 at the transition point and a power law with an exponential cutoff above and below the transition point as shown by Newman *et al.* in [12] and Cohen *et al.* who also discuss in [13] the generalization of this power law for random failure in scale-free networks.

III. INTERNET MAPS

Studying Internet robustness involves knowing the Internet topology. In this section we present the data that we use in our experiments and we explain how we build an overlay in order to relate the IP nodes to their owning AS nodes.

A. Sources

As we want to obtain accurate and directly applicable results, we do not use AS level Internet maps for the basis of our study because they are too coarse-grained. Instead, we focus on the IP connectivity and therefore we prefer to work at the router level of the Internet. We use three Internet maps. The first one is a router level anonymous map which is the result of the merging of a map collected by the SCAN projet [5] and another one collected by the Lucent Internet mapping project [6]. It is the biggest router level Internet map currently available to our knowledge. It has been assembled in 1999 and has been used in [4]. The map as-is is not connected. We have removed 33 nodes in order to make this map connected. This is negligible in comparison of the size of this map. Furthermore these nodes were mostly in connected components of size 1 or 2 (i.e. single nodes or pairs of nodes). The second map is a router level map collected from our laboratory (called LSIIT and located in Illkirch, France) by using the Mercator software written by Govindan *et al.* and described in [5]. This map is connected. The collect lasted four months from April to July 2002. Unlike the '99 map, this one contains the IP addresses of the routers' interfaces. The third and last one is an AS level map collected by *route-views* [14] at the beginning of July 2002. We use it mainly to build an overlay with our '02 map but also for comparison with some router level results. Table I contains some information about these maps.

B. Building the overlay

We build a topological overlay in order to relate router level and AS level information. Our overlay creation method is quite different from the methods used in [15] and in [16] because we directly map the routers found by Mercator to the ASes found in the BGP table through the use of the IP interfaces and the BGP prefixes. Thus we do not have to generate the AS graph by a collapsing algorithm such as the one in [15] and we avoid the potential errors brought by the cases where many disjoint clusters of nodes belonging to the same AS have to be reassigned.

We use a BGP routing table dump from *route-views* created in July, 1st 2002 to build this overlay as well as an AS level map of the Internet containing 13529 nodes. For the overlay construction, we associate every prefix found in the table to its advertising AS (i.e. the AS at the right end of the AS path). This AS is not necessarily the originating AS of the prefix because the originating AS can be masked by AS path aggregation [17] (so errors can be introduced here). In the case where a prefix can be associated to more than one AS (because of protocol or database errors), we keep the first AS having the "i" (i.e. internal) flag set if one is found, otherwise we keep the first AS found (11 cases in our table, also sources of errors). The table contains 119814 prefixes (consistent with the results found by Bu *et al.* in [18]).

Then we use our IP level information collected by using Mercator to build a router level map of the Internet. The description of the Mercator software and its limitations can be found in [5]. Mercator can perform interface disambiguation and thus can properly assign multiple interfaces to their corresponding router. The resulting router level map contains 203854 interfaces and 188347 nodes. This yields an incidence rate of multiple interfaces of 8.2% which is nearly half the value observed in [19]. A first explanation for this difference is that our map, with an average degree of 2.5, is probably lacking an important number of redundant links that may potentially be multiple interfaces to any one node. Then for each interface, we search the longest prefix matching it and associate the originating (or advertising) AS of this longest prefix to the interface.

In this process, 1296 interfaces could not be mapped to an AS. 57 of these interfaces were class A addresses, 405 were class B and 834 were class C. Unresolved interfaces represent 0.64% of all the interfaces which is comparable to the 0.48% rate measured in [15]. Unlike their method, we have not used Internet Routing Registries (IRR) as additional sources of information because they are not accurate enough at least for our usage. Indeed Chen *et al.* have shown in [20] that about 62% of the records in the RIPE database are either void or obsolete despite the fact that RIPE is actively maintained up to date. Among the unresolved interfaces, many do belong to ASes (as a few requests to an IRR shows) but some of them such as the 188.1.x.x German research network (called DFN) are configured not to belong to any AS. We mark all the unresolved interfaces as belonging to the AS number 0. We define the meaning of the AS number 0 as: "an IP address with AS number 0 does not belong to any AS". Despite the

Fig. 1. AS size complementary cumulative distribution function.

Fig. 2. Correlation between AS size and AS degree.

Fig. 3. Number of BGP routers per AS CCDF.

Fig. 4. BGP connection size CCDF.

important number of private IP addresses found in [15], we have not found any private address in our IP level information. These addresses are most probably filtered somewhere on the return path to our machine.

Finally for each router, we define its IP address and search for its AS number by using the interfaces associated to the router by Mercator. If the router has only one interface, the process is straightforward: we assign both the IP address and the AS number of the interface to the router. If the router has more than one interface, we search for the interface that has the smallest IP address and an AS number different from 0 and we assign this IP address and AS number to the router (this method can produce errors). At the end of the process, every router has a default IP address and an AS number (which can be 0). We then bind the routers to their respective ASes in our software in order to obtain the overlay data structure. Lastly we use the overlay to determine which routers could be potential BGP routers. For each router we look at its neighboring routers and if one or more of them is located in an AS that has a non zero and different AS number then we mark the initial one as a BGP router. We find that 40316 routers can be considered as BGP routers and this represents 21.4% of all the routers of the IP topology map. We do not currently have the information for assessing the accuracy of this figure to the real value.

To assess the accuracy of our overlay we study the AS size distribution. The size of an AS is equal to the number of routers contained in this AS. We find that only 5277 ASes contain 1 or more routers in our AS overlay and that represents only 39% of all the nodes of the AS map. This clearly shows that our router level map is far from being complete. Despite this flaw, we look for the heavy tail distributions found in Internet overlays by Tangmunarunkit *et al.* in [21]. A convenient method already used in [21] and [22] to achieve this is to plot the AS size complementary cumulative distribution function (CCDF) on log-log scales. Figure 1 shows the CCDF of the AS size distribution and we can see that it is heavy-tailed. To follow the analysis of [21], we also plot in figure 2 the correlation between size and degree for the 100 largest ASes. The correlation for these ASes is strong and very similar to the one found in [21].

Figure 3 shows on log-log scales the CCDF of the number of BGP routers per AS distribution and it unsurprisingly exhibits a heavy tail (a linear regression yields a CC of 0.993). Finally thanks to our overlay creation technique, we are able to study the BGP connection size distribution. The size of a BGP connection is the number of IP links going from any router belonging to a given AS to any router located in another given single different AS. Similarly to the ASes,

only 6630 connections are filled with one or more IP links. This yields a 23.6% filled vs empty connection ratio that is much lower than the ASes filling ratio. This again tends to point out that our router level map lacks a significant number of IP links. Fig. 4 shows on log-log scales the CCDF of the connection size distribution and it reveals a heavy tailed distribution (with a CC of 0.995). We can conclude that although our router-AS overlay is incomplete, it displays many heavy-tailed distributions which are typical macroscopic level characteristics of the Internet topology.

IV. ATTACK TECHNIQUES

Before explaining how we attack the Internet connectivity, a few words must be written on the obvious limitations of our study. The most important limitations are:

- The Internet maps that we use are incomplete. The router level maps are obviously incomplete as explained in [5] and in section III. The AS level map from *route-views* is also incomplete especially concerning the number of BGP connections. Chang *et al.* have shown in [23] that using additional sources of AS routing information to *route-views* increases the number of BGP connections seen by up to 47.7%. This has a direct impact on the connectivity robustness of the AS network.
- The router maps produced by Mercator (or any other IP topology measurement tool) should not be seen as containing only real router devices and real point-to-point links but simply as maps containing IP nodes and IP links. Indeed many of these IP nodes and links are purely virtual and thus do not exist physically. This is because, especially in backbone and core networks, the (virtual) circuit switching technologies such as Frame Relay, ATM or MPLS, may not only completely hide their own topology to the IP layer but may also emulate routers or IP links in a completely different topology as theirs. As a result it may not be meaningful to remove one given IP node or its removal may imply the simultaneous removal of several other nodes because they are physically the same node. The same remark does apply for the IP links. Although we speak of "routers" throughout the paper for the sake of simplification, this characteristic must not be overlooked. However we can not avoid this pitfall given the data currently at our disposal.
- We can only quantify the destruction of the network by measuring the size of its connected components. This means that every node has the same value. This is certainly not true in reality. Indeed the quantity of valuable information available on the local area network of a given router will probably not be equal to the other routers. For instance, if a very popular server is cut from the largest part of the network (and no mirrors exist) then the users will complain even if 90% of all the nodes can still communicate. However it is hard to assign an information quality or importance value to the IP nodes without any prior study on this topic.

In this paper, we do not study random node or link failures. This topic has already been studied in [1], [2], and [4] and

all have shown that the overall connectivity of the Internet is very resistant to random failures. We are rather interested in targeted node removal. This means removing nodes that are important to the connectivity of the network in order to tear it down. We do not study targeted link removal for the moment because it is harder to characterize the topological importance of a link and its removal impact is usually less effective. Indeed the removal of a node simultaneously cuts down all its adjacent links. In reality, the removal of a node would surely not be someone coming in the cable room and destroying the device. Tearing down the Internet, as we show later, requires the removal of hundreds to thousands of routers to be effective and this can not be done physically. We envision that an attack would be possible only by the use of specific software tools such as in the case of distributed denial of service attacks. The routers would not be really destroyed but rendered useless by the injection of forged packets that would wreak havoc in their routing information bases or in the code of the routing protocols running on them. This threat has not been overlooked by the IETF which has produced several years ago three request for comments on the authentication of the routing protocol messages of RIP-2 [24], OSPFv2 [25] and BGP [26].

Our aim is not currently how to remove a specific node but how to remove the nodes in the most efficient way in order to minimize the amount of node removal necessary to reach any given level of fragmentation. The attack techniques on nodes can be classified in two broad groups: static and dynamic. In the static group, each node is assigned once and for all an importance value based on one or more criteria. The criteria for each static attack are listed in table II. For instance in the degree-distance attack, the nodes are valued first according to their degree and second according to their average distance. The higher the degree, the higher the importance of the node. If several nodes have the same degree, the node having the minimum average distance is assigned the highest value, etc. We point out that, as we said earlier in the paragraph on limitations, the importance of a node is only evaluated by using topological criteria (e.g. degree, node is a root, node is in mesh, etc.). The nodes are then removed from the network one by one in decreasing order of importance. The advantage of this method is that it is very fast to compute because values are computed only once at the beginning. All the techniques of the static group follow a common algorithm given below:

Static attack group

1. Determine the importance value of each node
2. Place the node names in a stack in decreasing order of node values
3. While (network is not empty)
4. Extract the top node name from the stack
5. Remove the corresponding node from the network
6. Analyze the connectivity of the network

In the dynamic group, each node is assigned an importance value, at the beginning and after each removal of a single node, based on one or more criteria. The criteria for each dynamic attack are listed in table III. For instance in the adaptive attack, the nodes which are in the largest connected component are selected first and are then valued depending on their degree. Thus this attack removes the highest degree node of the largest

TABLE II
STATIC ATTACK CRITERIA FOR NODE SELECTION

Attack technique name	Criteria in decreasing order of importance
Degree	Degree
Degree-topology	Degree, root, mesh, cutpoint
Degree-distance	Degree, average distance

TABLE III
DYNAMIC ATTACK CRITERIA FOR NODE SELECTION

Attack technique name	Criteria in decreasing order of importance
Adaptive	Largest connected component, degree
Adaptive-topology	Largest connected component, cutpoint, degree

connected component (which is not necessarily the highest degree node of the network). The most important node is removed from the network, the connectivity is evaluated and then importance values are calculated again to find the most important node among the remaining ones. The advantage of this method is that it is very accurate in removing important nodes in a given topological situation (as the latter continually evolves with the removal of nodes). However the calculation time is longer. All the techniques of the dynamic group follow a common algorithm given below:

Dynamic attack group

1. Determine the importance value of each node
2. Select the node that has the highest value
3. While (network is not empty)
4. Remove the selected node from the network
5. Analyze the connectivity of the network
6. Recalculate the importance value of each remaining node based on the new connectivity
7. Select the node with the highest value

Although the authors of the previous papers have not detailed exactly how they have carried out the attacks, they have indicated using the following ones:

- [1], [2] and [3] have all used the highest-degree node first attack.
- [3] has also used the lowest-average-distance node first attack.
- [4] has also used the highest-hop-exponent node first attack.

To conclude we use five attacking techniques: the degree, degree-topology and degree-distance static attacks; and the adaptive and adaptive-topology dynamic attacks. The degree attack has already been used in previous work while the four others are new (the degree-distance attack is not the same as the distance attack).

V. EXPERIMENTS

The experiments consist in testing our five attack techniques on the three Internet maps at our disposal with a strong focus on the two router level maps. The experiments show that the static attacks yield very similar results. In the same way, dynamic attacks also yield very similar results. As a consequence, we only present in this section the results concerning the degree (representative of the static group) and the adaptive

(representative of the dynamic group) attacks. Furthermore in some cases the results produced with the SCAN-LUCENT'99 map are very close to the ones produced with the LSIIT'02 map. In this case, we only show the results produced with the SCAN-LUCENT'99 map in order to avoid repetition.

A. Metrics

Finding the proper metric to quantify connectivity is not an easy task. In this section we present existing metrics, define the metrics that we use in our experiments and explain why we choose them. In [1], the authors use the diameter as a metric of interconnectedness of a network. They define it as the "average length of the shortest paths between any two nodes". This is usually called the average distance or average path length of a graph, the diameter being usually defined in graph theory [27] as the maximum length of the shortest paths between any two nodes. Nevertheless this metric may be interesting as the inflation of the average distance shows how the network connectivity is damaged, but even if a path is inflated by 100% because of network element removals, one can still communicate. Furthermore the distance used here is usually the hop count metric and it is not necessarily proportional to the delay or bandwidth metrics. The use of a distance estimation service such as IDMaps [28] could make possible the use of the delay instead of the hop count. Several distance metrics such as the diameter and the average distance are also used in [3] however this still does not inform us on the true loss of connectivity among nodes. We want to investigate what it takes to truly block any communication between the largest number of nodes. The authors of [1] also monitor the average size of the isolated clusters (i.e. all the clusters excepted the largest one). However this is not really an interesting metric because the cluster size distribution is very erratic (i.e. it seems to be heavy tailed). This means that the average or the median values of this distribution do not accurately reflect it. In fact measuring these metrics always yields values between 1 and 2 for the average size and 1 for the median size whereas clusters of hundred or thousand nodes may be present.

Another metric used in [1], [2] and [3] is the size of the largest connected component. This is a very interesting metric because it gives an upper bound on the number of nodes that can communicate, thus we use this metric in our experiments. A scale free network such as the Internet is usually composed of a forest part (i.e. containing trees) and a mesh part (i.e. containing cycles and bridges). The ratio is typically 2/3 of forest for 1/3 of mesh. Moreover the largest part of the mesh is typically composed of a giant biconnected component (i.e. in it there exists at least two disjoint paths between any pair of node). In our '02 map, the biconnected component is 82.9% of the mesh. All this means that breaking the connectivity of this portion of the network is probably difficult and studying it shrinking may provide valuable information.

The most interesting metric to look at would be the connected component distribution itself. However this distribution contains usually too many values to be easily manipulated. That is why we have defined groups of connected component

TABLE IV
DEFINITION OF CLASSES

Class identifier	Cluster size range
class 0	1
class 1	2-10
class 2	11-100
class 3	101-1000
class 4	1001-10000
class 5	10001-100000
class 6	100001 and above

size values in order to have a more concise representation of the connected component distribution. We call these groups *classes* and they are defined in table IV. For instance, any connected component in class 2 contains at least 11 nodes and at most 100 nodes. As a reminder, the class number is the decimal logarithm of the upper bound size of that class. We choose absolute thresholds instead of relative ones (i.e. in % of the network size) because they carry much more information. If a connected component is from class 0 or 1, we have a precise idea of its situation especially concerning the distances. Any two nodes in a class 1 connected component will be at most at 9 hops from each other. If we define a class with a maximum relative size such as 0.1%, this represents an upper bound of 284 nodes in the SCAN-LUCENT map. A connected component of this size can be quite widespread and as the diameter measured in Internet router level maps is usually around 30 hops, we cannot deduce how distance limited are the paths in this component. We compute the connected components by using a modified version of the Tarjan algorithm [29] valid for undirected graphs [30] and having the same complexity.

Classifying the connected component in classes gives an interesting insight on the network fragmentation level but we propose a metric that catches even better the situation of any given node in the network. This metric is the distribution of the relative number of nodes per class (i.e. the number of nodes per class expressed as a fraction of the network total size). This metric enables us to answer the question: "what is the % of probability that a node can communicate with at most a given number (i.e. class upper bound) of nodes for a given % of node removal?". We compute the distribution by labelling each node belonging to the same connected component with an common identifier. We can then compute the relative size of each class by adding the sizes (i.e. number of nodes) of each connected component belonging to the same class.

To conclude we use three metrics: the largest connected component size which has already been used in previous work, the distribution of the frequency of the connected component classes and the distribution of the relative number of nodes per class which are both new.

B. Core robustness behavior

We study the evolution of the relative size of the largest connected component. The size is expressed as a fraction of the initial total number of nodes in the network. Figure 5 shows this evolution for the degree and adaptive attacks. The

two plots have roughly the same shape. The size abruptly decreases until a threshold where the size is close to zero. This phenomenon has already been observed in [1]. Knowing that the network can be torn down by removing roughly 5% of its nodes may seem frightening. The adaptive attack only significantly differs from the degree attack in the middle area of the graph. It reduces the threshold from 6.5% to 4.5% which does not seem to represent an important gain over the degree attack. The attacks on the LSIIT'02 map shown in figure 6 are more effective, especially the adaptive one. This may be caused by the lack of redundant links in this map which has an average degree of 2.5 compared with the average degree of 3.15 of the SCAN-LUCENT'99 map.

Figures 7 and 8 show the relative size of the largest connected component plotted on a decimal logarithmic scale. This representation is very interesting because it highlights the fact that the adaptive attack is much more potent than the degree attack. Indeed for the SCAN-LUCENT'99 map at 4% of node removal, the size of the largest connected component is already one order of magnitude smaller under adaptive attack than under degree attack. At 5.5% this gap reaches 3 orders of magnitudes. Even at 10%, there is still a gap of one order of magnitude. The results with the LSIIT'02 map show a similar behavior although the degree attack plot is more erratic.

C. Fragmentation of connected components

The fraction of the number of connected components for a given class with respect to the total number of connected components is given in figure 9 for the SCAN-LUCENT'99 map under a degree attack. After 0.5% of node removal, most of the classes do not bear a significant amount of connected components. Only class 0 and class 1 do make the bulk of all connected components. This clearly shows that the degree attack mainly produces single nodes or very small connected components (having less than 10 nodes). This means that even a targeted attack such as removing the highest degree node does not succeed in breaking the network into several parts of equal size. Figure 10 shows a similar behavior for the adaptive attack. We find similar results for the LSIIT'02 map. It would have been very interesting to analyze the diameters of the connected components to have an idea of how far one can go inside a connected component (as a function of its size for instance). As the diameter of the router level Internet is small (around 30 hops), connected components of class 2 (up to 100 nodes) or above could still span the whole network even if they can not reach a lot of other nodes. The results in [4] concerning the hop exponent tend to indicate strongly that this is not the case. The neighborhood size as a function of the hop count does decrease a lot under attack which together with our results tend to show that connected components of class 2 to 4 (middle classes) most probably have a small diameter. Also, due to the computational cost of the diameters for every connected component, we choose to leave this issue for future work.

The fraction of the number of connected components for a given class with respect to the total number of connected

Fig. 5. SCAN-LUCENT'99 largest cluster size evolution.

Fig. 6. LSIIT'02 largest cluster size evolution.

Fig. 7. SCAN-LUCENT'99 largest cluster size evolution (log scale).

Fig. 8. LSIIT'02 largest cluster size evolution (log scale).

components is shown in figures 11 and 12 plotted on a decimal logarithmic scale for the SCAN-LUCENT'99 map. Again this representation highlights the fact that the adaptive attack is much more potent than the degree attack. In the adaptive attack, the components are smashed much faster and in much smaller sizes than under the degree attack. For instance at 4.5% of node removal, no component of class 3 or above remains in the network whereas class 3 components do exist until 9% of node removal under a degree attack. We find similar results for the LSIIT'02 map.

D. Distribution of the nodes

We show here how the nodes themselves are distributed among the classes. Figure 13 shows this distribution on the SCAN-LUCENT'99 map under a degree attack. This figure accurately shows how the reduction of the largest connected component produces nodes in classes 0 to 2. Classes 3 to 5 really seem to catch the transitional state of the largest component towards destruction. The distribution under an adaptive attack is shown in figure 14 and exhibits some differences with the one under a degree attack. Particularly the plots for components of class 1 and 2 do not inflect as much as in a degree attack and they even seem to be rather linear when the node removal fraction increases beyond 5%. As a

result, at 8% of node removal, only class 0 and 1 components remain in the network under an adaptive attack while class 0 to 3 components can still be found under a degree attack. Furthermore the existence of intermediate classes 3 to 5, due to the decomposing largest component, is much shorter. These figures are very interesting in the sense that they provide the % of probability of being in a given fragmentation situation. For instance under an adaptive attack that has removed 3% of the nodes, you have 25% of probability to be still connected to at most 9 other nodes only, whereas at 6.5% this probability reaches 50%! We find very similar results for the LSIIT'02 map.

E. Destruction levels

An interesting question is "what fraction of nodes has to be removed in order to reach a given level of fragmentation?". Figure 15 helps to answer this question by giving the fraction value to remove so that no component is bigger than a given class number. For instance if we want no component of the LSIIT'02 map to have more than 100 nodes, we have to remove 4% of the nodes in the network with an adaptive attack. We can see that for the same technique, the values of the SCAN-LUCENT'99 map and the LSIIT'02 map are very close. We also notice that the plots for the AS'02 map have

Fig. 9. Cluster class fraction (degree attack).

Fig. 10. Cluster class fraction (adaptive attack).

Fig. 11. Cluster class fraction (degree attack, log scale).

Fig. 12. Cluster class fraction (adaptive attack, log scale).

the same behavior as the router level maps: only the values are different. Although these values depend on the map, a quick extrapolation can give us an idea of what fraction would be required to tear down the Internet. The Internet had around 200 million hosts in June 2002 according to [31]. If we assume that 1% of it are routers, this gives us 2 million routers which is around 10 times the number of routers in the LSIT'02 map. If we assume again that there is a linear relationship between a network size and its fraction threshold we find that one has to remove 14.8% of the routers of the Internet in order to leave only components of class 2 or less (this would represent nearly 300000 routers to attack!). We calculate this example with the class 2 (i.e. 100 nodes max) because this value is what is considered close to zero in [1] for a large enough network (i.e. larger than a few thousand nodes). It's also worth noticing that [1] have found a threshold of 5% for their AS network snapshot which tends to indicate that they have used a static degree attack (indeed the threshold for our snapshot under a degree attack is 4.7% and only 2.8% under an adaptive attack). Figure 16 has a similar purpose than the previous one excepted that the fragmentation levels are given as the size of the largest cluster expressed as a fraction of the initial network size. For instance, to obtain a maximum cluster size of 10% (of the network size) with an adaptive attack on the LSIT'02 map, one must remove 3.4% of the nodes. We can clearly

TABLE VI
NUMBER OF NODES TO REMOVE TO REACH A GIVEN LEVEL UNDER
ADAPTIVE ATTACK

Level	SCAN-LUCENT'99	LSIT'02	AS'02
Damaged	4272	2825	108
Destroyed	11391	7534	379
Annihilated	21358	14126	582

see in this figure that the benefits of using an adaptive attack instead of a degree attack dramatically increase with the level of fragmentation wanted.

All previous results provide a lot of information but they are not easy to summarize. In order to concisely evaluate the amount of work to tear down a network connectivity, we define (quite arbitrarily) three destruction levels in table V. They provide a basic and well-defined distinction of the various thresholds of fragmentation of the network. The fraction of nodes to remove to reach a given destruction level does depend on the size/kind of network (i.e. router level or AS level) to attack as was already noticed in [1]. Table VI shows the fractions translated into the number of nodes to remove using an adaptive attack in our studied maps in order to reach a given destruction level.

Fig. 13. Distribution of the nodes among classes (degree attack).

Fig. 14. Distribution of the nodes among classes (adaptive attack).

Fig. 15. Node removal fraction as a function of the highest cluster class.

Fig. 16. Node removal fraction as a function of the largest cluster size in %.

TABLE V
DESTRUCTION LEVELS UNDER ADAPTIVE ATTACK

Level	Effect	(Router / AS) fraction to remove
Damaged	Any node sees at most 50% of the network	1.5% / 0.8%
Destroyed	Any node sees at most 1% of the network	4.0% / 2.8%
Annihilated	Any node sees at most 9 other nodes	7.5% / 4.3%

F. Identifying targets

It is possible with the destruction levels to define how many nodes have to be destroyed in order to put the network in a given state of fragmentation but we still do not know who these nodes are excepted for their topological properties. The use of our overlay gives us a lot of valuable information on these target nodes. Table VII lists the first 10 targets chosen by an adaptive attack on the AS'02 map. Unsurprisingly, this top 10 target ranking exactly matches the degree ranking of the nodes (i.e. the target node list corresponds to the nodes sorted in decreasing order of degree). However this is not always true: the target and degree ranks match only for 29 of the first 100 largest ASes. We know from the previous section that 108 ASes have to be removed in order to damage the network and 379 in order to destroy it. This seems nearly

impossible just by looking at table VII. Nobody has already defined the destruction of an AS. One method would be to cut all the IP links connecting the target AS to its neighbor ASes by attacking the corresponding interfaces (recall that the size of a BGP connection follows a power law, the largest connections will probably belong to the largest ASes and will be difficult to tear down). Another method would be to cut inner IP links inside the AS in order to prevent the traffic from passing through it (again the size of an AS follows a power law, the largest ASes will contain a lot of routers and inner links thus they will probably be very difficult to tear down). So despite the fact that the AS removal fraction to destroy the AS network is very low at only 2.8%, this still looks like an impossible mission to carry out in reality.

Table VIII lists the first 20 targets chosen by an adaptive attack on the LSIIT'02 map. As above, the target ranking

Fig. 17. Proportion of the target routers per AS at the damaged level.

Fig. 18. Proportion of the target routers per AS at the destroyed level.

TABLE VII
TOP 10 AS TARGETS

Target rank	AS no	Name	Degree
1	701	Alternet	2602
2	1239	SprintLink	1638
3	7018	AT&T	1343
4	209	Qwest	841
5	3561	Cable&Wireless	806
6	1	Genuity	596
7	702	UUNET	584
8	3549	GlobalCrossing	579
9	6461	Abovenet	551
10	2914	Verio	476

exactly matches the degree ranking of the nodes (however the target and degree ranks match only for 64 of the first 100 largest routers). This table shows four interesting things:

- No router in the top 20 list belongs to any of the top 10 ASes! This is most probably an artifact of our LSIT'02 map which may not only be incomplete but also biased by the starting point of our measurements (i.e. Illkirch, France).
- As another effect of this bias, the first target router (which is also the largest degree router, recall from section IV that it may not be a real router) is in Austria and the second target router is in Switzerland. It's worth noticing that Illkirch is located close to the German border and is very close to Switzerland (100 km) and Austria (700 km).
- 18 of all the top 20 target routers belong to different ASes. Only ASes 3112 and 3320 have 2 routers in the list each. This tends to show that the targets are well distributed among the ASes which makes the task of attacking them more difficult.
- Surprisingly only 13 out of 20 routers are BGP according to our BGP determination method explained in section III.

To further illustrate the spreadability of targets among ASes, we plot in figures 17 and 18 the ASes in proportion of their number of target routers for the damaged level and destroyed

level respectively. We know from the previous subsection that 2825 target routers have to be removed in order to damage the LSIT'02 router network and 7534 in order to destroy it. Only the 20 ASes having the largest number of target routers are plotted, the others are grouped under the label "others". We notice several significant points:

- The results shown by both figures are very similar which tends to prove that the distribution of the targets among ASes does not depend on the fraction of removed nodes.
- As already suggested by the results in table VIII, the target routers are not concentrated in specific ASes. For the damaged level, the 2825 routers divide themselves among 583 ASes and for the destroyed level, 7534 routers are split among 1102 ASes. The 20 first ASes only contain around 40% of all the target routers. Furthermore, the first AS only contains 9.2% and 6.7% of the targets for the damaged and destroyed levels respectively, the second AS contains 7.9% and 5.7% and the following ASes much less than that.
- The first 6 largest target-containing ASes do match with the first 6 largest ASes although not in the same order. The order of magnitude of the target routers (several thousands) do restore what was initially presupposed: the top target ASes do contain the largest number of target routers although these routers may not be the largest ones with respect to their degree.
- Finally 47.2% and 48.9% of the target routers are BGP for the damaged and destroyed levels respectively as opposed to the 65% found in the top 20 list. This means that looking to remove BGP routers exclusively or in priority may not bring a significant benefit for an attack.

Of course all these values have to be taken as indications only. The accuracy of our LSIT'02 router map may not be enough to draw definite conclusions about the Internet robustness but they illustrate where and how to look for potential weaknesses in the network.

It is to our knowledge the first time that an attempt is made to identify potential targets for connectivity attacks and although it may not be a complete success because of the

TABLE VIII
TOP 20 IP TARGETS

Target rank	IP address	Network name	Country	Degree	BGP	AS no
1	193.171.13.10	UDNVIE-VIE	AT	1436	yes	1853
2	192.65.185.2	CERN-BLK2	CH	1197	yes	513
3	65.172.70.18	FON-110180915289817	USA	1172	yes	7066
4	198.151.130.225	RUTGERS-XUNET	USA	1024	yes	46
5	66.54.144.41	YIPES-BLK3	USA	764	no	6517
6	202.8.94.2	SINGAREN	SG	632	yes	7610
7	192.88.191.222	OAR-BLK3	USA	601	no	3112
8	128.206.130.253	MONET	USA	588	no	2572
9	206.196.177.58	UMD-NOCNETS	USA	522	yes	10886
10	129.118.4.85	TTUNET	USA	489	yes	10421
11	62.154.17.194	DTAG-BB16	DE	406	yes	3320
12	164.58.12.254	ONENET	USA	372	no	5078
13	146.97.40.82	JANET-IP	GB	362	no	786
14	62.154.5.89	DTAG-BB16	DE	324	yes	3320
15	131.91.200.2	FAU	USA	244	yes	12013
16	192.88.191.234	OAR-BLK3	USA	240	yes	3112
17	193.63.74.233	DARESBUY1	GB	228	no	786
18	62.154.5.121	DTAG-BB16	DE	225	yes	3320
19	164.58.2.174	ONENET	USA	224	no	5078
20	142.227.1.49	EDNET-NS-CA	CA	185	yes	855

lack of accuracy of our data, our findings still give a novel insight on the cost of carrying out an attack. Although the router removal fraction to destroy our Internet map is only 4%, the widespread distribution of the target routers will probably make the task difficult.

VI. CONCLUSIONS

In this paper we proposed several new techniques to attack the IP level connectivity of the Internet as well as new metrics to evaluate the amount of fragmentation in the network and the distribution of the routers among the clusters. We showed that under adaptive attack, the removal of only 4% of the routers leaves clusters of at most 100 routers in a map being 3 orders of magnitude larger. In the same way, the removal of 1.5% of the routers makes the largest cluster less than half the initial size of the network. Despite these worrying results, we also showed that these thresholds are dependent on the size of the network which means that tearing down the Internet itself would require simultaneous attacks on hundreds of thousands of routers. Furthermore we discovered that the targets are not concentrated into specific ASes but spread over thousands of them and that the quality of the router level map is crucial to accurately aim the best targets. We point out that the router level topology of the Internet is very hard to obtain and is still the focus of many ongoing research projects. Finally, thanks to our overlay method, we showed that our Internet router level map was incomplete despite the fact that it was containing more than 180k nodes and collected in several months. We showed that attacks rely heavily on the partial view they maintain for the topology, thus making a focused attack on a regime effective, while an overall world-wide attack almost impossible.

All these factors lead us to think that undertaking a massive attack on the Internet connectivity may not be currently feasible. Moreover the results of our study could be used to increase the robustness of the network by determining how to

place redundant links or routers in order to raise the destruction thresholds. They could also be used in the context of mirror placement [32] to define where to locate mirrors to make them reachable by a maximum number of hosts when the network is at a given level of fragmentation. Although the study of the Internet connectivity robustness is a recent issue, the growing concern about security will probably make this topic play an important role in the near future.

REFERENCES

- [1] R. Albert, H. Jeong, and A.-L. Barabási, "Error and attack tolerance of complex networks," *Nature*, no. 406, pp. 378–382, 2000.
- [2] S. Tauro, C. Palmer, G. Siganos, and M. Faloutsos, "A simple conceptual model for the internet topology," in *Proceedings of IEEE GLOBECOM'01*, San Antonio, Texas, November 2001.
- [3] A. Broido and K. Claffy, "Internet topology: connectivity of ip graphs," in *Proceedings of SPIE ITCOM'01*, Denver, CO, USA, August 2001.
- [4] C. Palmer, G. Siganos, M. Faloutsos, C. Faloutsos, and P. Gibbons, "The connectivity and fault-tolerance of the internet topology," in *Proceedings of ACM SIGMOD/PODS Workshop on Network-Related Data Management (NRDM'01)*, Santa Barbara, May 2001.
- [5] R. Govindan and H. Tangmunarunkit, "Heuristics for internet map discovery," in *Proceedings of IEEE INFOCOM'00*, Tel Aviv, Israël, March 2000.
- [6] H. Burch and B. Cheswick, "Mapping the internet," *IEEE Computer*, vol. 32, no. 4, pp. 97–98, 1999.
- [7] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proceedings of ACM SIGCOMM'99*, Cambridge, Massachusetts, USA, September 1999.
- [8] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: Degree-based vs. structural," in *Proceedings of ACM SIGCOMM'02*, Pittsburgh, Pennsylvania, USA, August 2002.
- [9] D. S. C. nad M. E. J. Newman, S. H. Strogatz, and D. J. Watts, "Network robustness and fragility: Percolation on random graphs," *Phys. Rev. Lett.*, no. 85, pp. 5468–5471, 2000.
- [10] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin, "Breakdown of the internet under intentional attack," *Phys. Rev. Lett.*, no. 86, pp. 3682–3685, 2001.
- [11] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han, "Attack vulnerability of complex networks," *Phys. Rev.*, no. E 65, 2002, 056109.
- [12] M. E. J. Newman, S. H. Strogatz, and D. J. Watts, "Random graphs with arbitrary degree distributions and their applications," *Phys. Rev.*, no. E 64, 2001, 026118.

- [13] R. Cohen, D. ben Avraham, and S. Havlin, "Percolation critical exponents in scale-free networks," *Phys. Rev.*, no. E 66, 2002, 036113.
- [14] *BGP data from route-views*, University of Oregon, Advanced Network Technology Center, <http://www.routeviews.org/>.
- [15] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The impact of routing policy on internet paths," in *Proceedings of IEEE INFOCOM'01*, Anchorage, Alaska, USA, 2001.
- [16] H. Chang, S. Jamin, and W. Willinger, "Inferring as-level internet topology from router-level path traces," in *Proceedings of SPIE ITCOM'01*, Denver, CO, USA, August 2001.
- [17] Y. Rekhter and T. Li, "A border gateway protocol 4 (bgp-4)," Internet Engineering Task Force, Request For Comments 1771, March 1995.
- [18] T. Bu, L. Gao, and D. Towsley, "On routing table growth," in *Proceedings of IEEE GLOBECOM'02*, 2002.
- [19] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the marginal utility of network topology measurements," in *Proceedings of ACM SIGCOMM Internet Measurement Workshop'01*, July 2001.
- [20] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "The origin of power laws in internet topologies revisited," in *Proceedings of IEEE INFOCOM'02*, New York City, NY, USA, June 2002.
- [21] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Does as size determine degree in as topology?" *ACM Computer Communication Review*, vol. 31, 2001.
- [22] T. Bu and D. Towsley, "On distinguishing between internet power law topology generators," in *Proceedings of IEEE INFOCOM'02*, New York City, NY, USA, June 2002.
- [23] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Towards capturing representative as-level internet topologies," University of Michigan, Tech. Rep., 2002.
- [24] F. Baker and R. Atkinson, "Rip-2 md5 authentication," Internet Engineering Task Force, Request For Comments 2082, January 1997.
- [25] S. Murphy, M. Badger, and B. Wellington, "Ospf with digital signatures," Internet Engineering Task Force, Request For Comments 2154, June 1997.
- [26] A. Heffernan, "Protection of bgp sessions via the tcp md5 signature option," Internet Engineering Task Force, Request For Comments 2385, August 1998.
- [27] F. Harary, *Graph Theory*. Addison-Wesley, 1969.
- [28] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang, "Idmaps: A global internet host distance estimation service," *IEEE/ACM Transactions on Networking*, vol. 9, no. 5, pp. 525–540, 2001.
- [29] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, *Data Structures and Algorithms*. Addison-Wesley, 1983.
- [30] S. Baase, *Computer Algorithms*, 2nd ed. Addison-Wesley, 1988.
- [31] *NetSizer: Internet Growth Forecasting Tool*, Telcordia Technologies, <http://www.netsizer.com/>.
- [32] E. Cronin, S. Jamin, C. Jin, A. Kurc, D. Raz, and Y. Shavitt, "Constrained mirror placement on the internet," *IEEE Journal on Selected Areas in Communications*, April 2002.

Damien Magoni received the M.S. degree in telecommunications from the Ecole Nationale Supérieure des Télécommunications de Paris, France, in 1995 and the Ph.D. degree in computer science from the Université Louis Pasteur, Strasbourg, France, in 2002.

He is currently an assistant professor in the Department of Computer Science at the Université Louis Pasteur, Strasbourg, France. His research interests include Internet topology measurements, multicast protocols and application layer overlays.