Modeling structure and resilience of the dark network

Manlio De Domenico* and Alex Arenas
Departament d'Enginyeria Informàtica i Matemàtiques, Universitat Rovira i Virgili, 43007 Tarragona, Spain
(Received 5 December 2016; published 27 February 2017)

While the statistical and resilience properties of the Internet are no longer changing significantly across time, the Darknet, a network devoted to keep anonymous its traffic, still experiences rapid changes to improve the security of its users. Here we study the structure of the Darknet and find that its topology is rather peculiar, being characterized by a nonhomogeneous distribution of connections, typical of scale-free networks; very short path lengths and high clustering, typical of small-world networks; and lack of a core of highly connected nodes. We propose a model to reproduce such features, demonstrating that the mechanisms used to improve cybersecurity are responsible for the observed topology. Unexpectedly, we reveal that its peculiar structure makes the Darknet much more resilient than the Internet (used as a benchmark for comparison at a descriptive level) to random failures, targeted attacks, and cascade failures, as a result of adaptive changes in response to the attempts of dismantling the network across time.

DOI: 10.1103/PhysRevE.95.022313

I. INTRODUCTION

Since the Internet became a publicly accessible infrastructure and communication network, its resilience to random failures (caused, for instance, by unexpected crashes due to node malfunction or protocol errors) or attack (actions devoted to isolate nodes that play a vital role in the network) has been widely investigated [1–5]. In fact, the Internet exhibits highly nontrivial structural and dynamical properties, from a heavy-tail distribution of connections (known as a scale-free property [6]) to a moderate amount of clustering (proportional to the fraction of nodes that form closed triangles), whose modeling has been the subject of intense research activity [7–11]. In fact, several years after the Internet's first proper crash, in 1980, the focus of many studies has been, and still is, to improve its resilience [10,12–17]. In the late 1990s, about 30 years after the first Internet prototype, the U.S. Defense Advanced Research Projects Agency and the Office of Naval Research started to develop a communication network, at the application layer, based on anonymous connections and, in principle, resistant to both eavesdropping and traffic analysis [18]. This network was based on *onion routing*, a special infrastructure for private communications over a public network that is able to hide the content of a message and the identity of peers who are exchanging it [19]. Nowadays, this infrastructure is better known as the Tor network and represents the backbone of the Darknet, a Web of hidden services that are not reachable from within the Internet. The Darknet turned out to be the most suitable communication network to exchange sensitive information, both licit and illicit, soon becoming the target of governments trying to identify dissidents or of intelligence agencies, such as CIA and GCHQ [20], to contain unauthorized news leaks, distribution of illegal contents, or trade of illegal substances.

Here we characterize the structural properties of the Darknet across time, from 2013 to 2015, and we compare them against the Internet topology. Note that this comparison is performed at a descriptive level, using the structure of the Internet as a benchmark to highlight the salient features of the Darknet. The autonomous system data capture connections at the Internet layer (Internet Protocol packages), while Tor works on the application layer, which means that is built on top of the Internet and transport layers. We propose a model, based on how Tor functions, to reproduce with high accuracy the most salient characteristics of the Darknet. Finally, we perform a thorough analysis, based on simulations, of the resilience of both networks to three different types of failures—static, due to random disruptions or targeted attacks [1], and dynamical, due to the cascade failures induces by attacking a single specific node of the network [21,22]—and show that the Darknet is much more robust than the Internet from any perspective.

II. OVERVIEW OF THE DATA SETS

For our analysis, we use publicly available data sets for both the Internet and the Darknet. The Internet topology, at the level of autonomous systems (ASs), is sampled from historical AS-level topology data derived from Border Gateway Protocol monthly snapshots, consisting of IPv4 and IPv6 links appearing between different end points during that month. The data are hosted by the UCLA Computer Science Department's Internet Research Lab [23].

The Darknet topology is sampled from the data obtained by probing the Tor network to improve its performance [24]. The links between end points are extracted from the chain of circuits built by Tor clients to probe the network. The network is directed, but we will treat it as undirected in the following. This simplification hides the information about the entry and exit point of a circuit, which might be indistinguishably interchanged in our case. We are aware that this choice might slightly affect our study and a deeper analysis, further accounting for this additional complexity, is beyond the scope of the present study.

The full raw data are available upon request and a partial release can be downloaded from a public repository [25]. Although such data were obtained to study the performance of the Tor network and not its topology, they provide the best

*Author to whom correspondence should be addressed: manlio.dedomenico@urv.cat

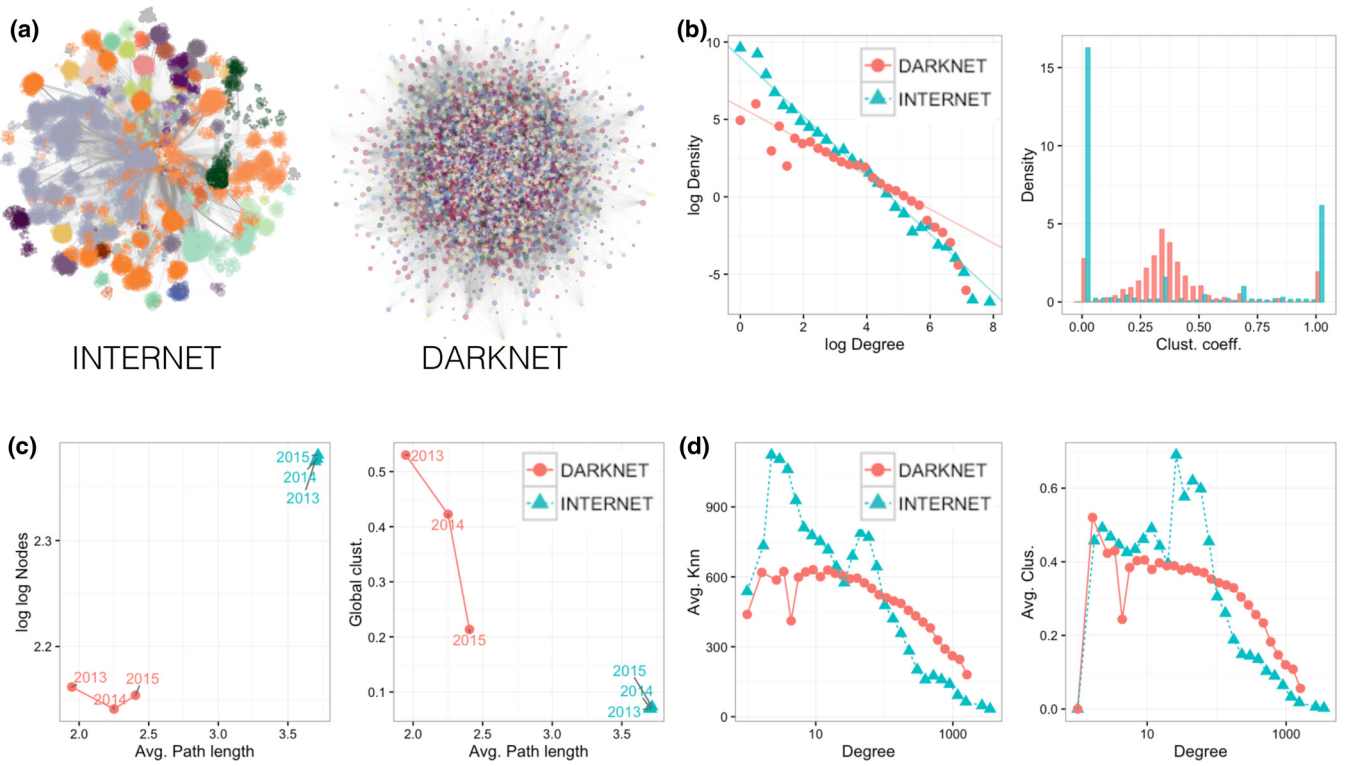

FIG. 1. Structural analysis of Internet and Darknet topologies. Force-directed visualization of the (a) Internet and the Darknet in 2015, with nodes colored to put in evidence the underlying mesoscale structure. (b) Density of the degree (solid lines are for guidance only) and local cluster coefficient for the two networks in 2015. (c) Scatter plot of network sizes, average path length, and global clustering coefficient for the three temporal snapshots considered in this study. (d) Average nearest-neighbor degree (left) and average local clustering coefficient (right) against degree, to characterize higher-order correlations (see Fig. 2 for other structural descriptors and their evolution between 2013 and 2015). The natural logarithm is considered in the figure.

approximation to the underlying topology of the Darknet to date.

In both cases, we have been able to build three temporal snapshots, corresponding to the networks in December 2013, May 2014, and January 2015. The Internet network snapshots have 46 462, 47 626, and 49 635 nodes with 195 446, 204 254, and 221 470 connections in the three periods, respectively. The Darknet network snapshots have 5921, 4953, and 5535 nodes with 2 017 542, 536 287, and 274 831 connections in the three periods, respectively. The structure of both networks for the 2015 period is shown in Fig. 1(a).

III. STRUCTURE OF THE DARKNET

A. Characterizing the Darknet topology

Let us indicate with $A_{ij}^{[\xi]}(t)$ the entries of the adjacency matrix of each network (where ξ denotes the Internet and Darknet) at time t ($t = 2013$, 2014, and 2015), with value equal to one if i and j are connected and zero otherwise (here $i, j = 1, 2, \ldots, N^{[\xi]}(t)$, where N indicates the number of nodes in the network and E the number of edges). For any node i in each network, we calculate the degree $k_i^{[\xi]}(t) = \sum_j A_{ij}^{[\xi]}(t)$, characterizing the number of connections of each node, and the local clustering coefficient $c_i^{[\xi]}(t)$, characterizing the tendency of nodes to form triangles, defined by the ratio between the number of closed triangles involving node i and the maximum number of triangles $\frac{1}{2}k_i^{[\xi]}(t)[k_i^{[\xi]}(t) - 1]$ node i might be part of. The mean degree is defined by $\bar{k}^{[\xi]}(t) = \langle k_i^{[\xi]}(t) \rangle$, whereas the average local clustering coefficient if given by $\bar{c}^{[\xi]}(t) = \langle c_i^{[\xi]}(t) \rangle$. Another macroscopic structural descriptor of interest is the global clustering coefficient $C^{[\xi]}(t)$, defined by the ratio between the total number of closed triplets and the total number of connected triplets of nodes in the network. In general, the values of $\bar{c}^{[\xi]}(t)$ and $C^{[\xi]}(t)$ are different for networks with a nonhomogeneous connectivity. Throughout the analysis no power-law fitting is performed; the power laws indicated in the plots are a simple guide for the eye.

The degree distribution is shown in Fig. 1 and exhibits high inhomogeneity, with evident heavy tails that resemble two different truncated power laws with a cutoff. The distribution of the local clustering coefficient is also different in the two cases, with much more unclustered nodes in the Internet than in the Darknet. In the Internet, a significant fraction of nodes have local clustering equal to 1 and few nodes have intermediate values, at variance with the Darknet, where the local clustering is more uniformly distributed and peaked around 0.3. On average, Darknet nodes have mean degree close to 100 and are more clustered than Internet nodes, which have mean degree close to 10.

To quantify how easy is to transmit information between any two nodes of each network, we calculate the average path length $\ell^{[\xi]}(t)$, obtained by averaging the length of all shortest path connecting any pairs of nodes, and the diameter

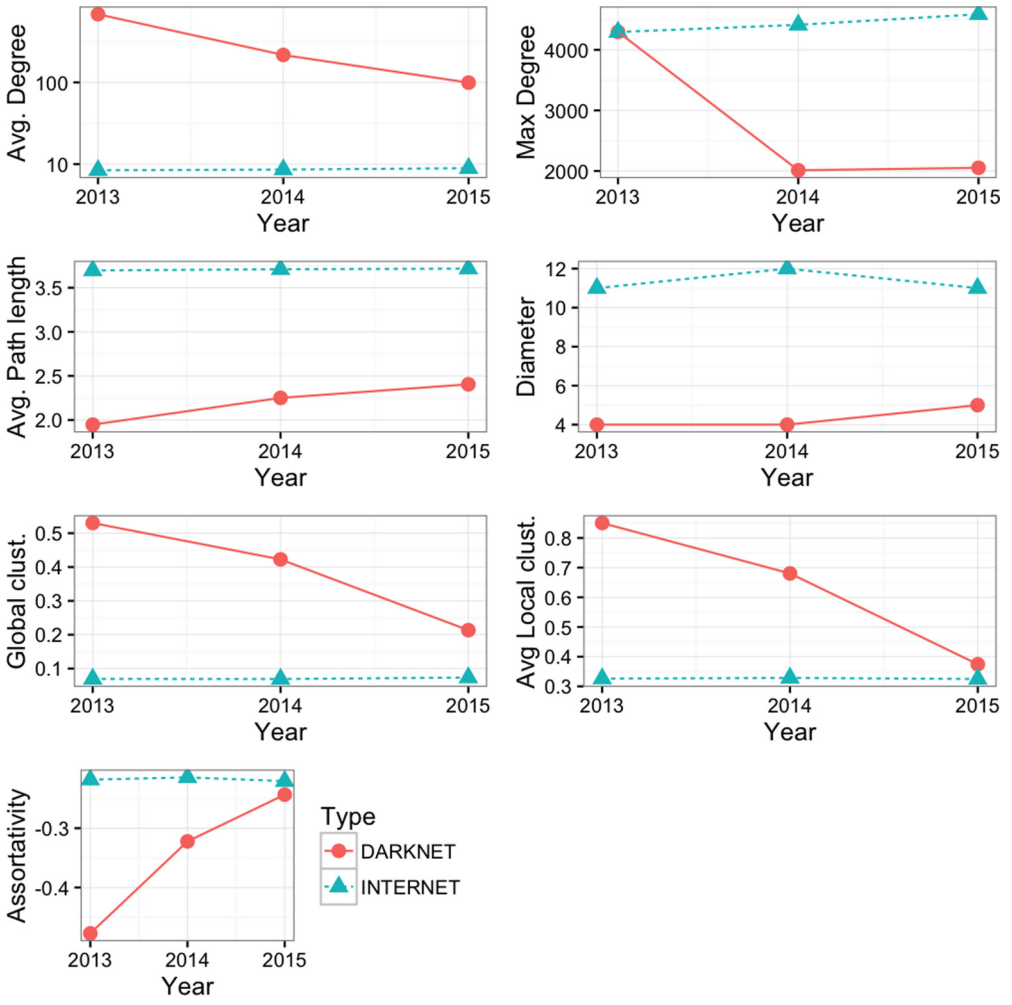

FIG. 2. Variation across time of important structural descriptors.

$D^{[\xi]}(t)$, defined as the length of the longest shortest path. The Internet has an average path length close to 3.5, with a diameter between 10 and 12, whereas the average path length in the Darknet ranges between 2 and 2.5, with a diameter between 4 and 5 (see Fig. 2). Therefore, assuming no time constraints in the propagation of the information for both networks, communication in the Darknet is, in principle, much faster than the Internet, with the two most extremal nodes separated by no more than five hops, less than half of the Internet. Given that the traffic in the Darknet is encrypted and the routing is decentralized, at variance with the Internet, the shortest paths for communicating compensate the higher latency and throughput of the channels.

Figure 1(c) shows the relationship between the size of the networks, their average path length, and their global clustering for the three snapshots under consideration. The presence of high clustering and short average path lengths is a strong indicator that the two networks have nontrivial topology and formation mechanisms. In fact, both networks exhibit the property known as the small-world phenomenon [26]. Small-world networks are characterized by high clustering, with respect to random expectation, and characteristic length scaling as $\ell \sim \log N$. The average local clustering of the Internet is more than 1000 times higher than its uniformly random expectation, whereas the clustering of the Darknet is between 7 and 32 times larger than its uniformly random expectation, suggesting nontrivial triadic closure mechanisms underlying both networks. The characteristic length of the Darknet is larger than its uniformly random expectations, whereas this is not the case for the Internet. Although the small-world phenomenon is better understood as a tendency, rather than being quantified by a single number, it is worth remarking that $\ell(t) \approx \log N(t)$, as in small worlds, for the Internet snapshots and $\ell(t) \approx \log\log N(t)$, as in ultrasmall worlds [27], for the Darknet ones.

The Internet and the Darknet also exhibit different types of higher-order correlations, as shown in Fig. 1(d), where the average nearest-neighbor degree and the average local clustering coefficient are scattered against the degree. This kind of analysis is generally used to shed light on the degree-degree correlations of the network: If the degree or the clustering of each node is independent of the nodes in the neighborhood, no trends are expected. Instead, the two systems present highly anticorrelations, with hubs (i.e., nodes with the largest degree) tending to be connected, on average, to nodes with a much smaller degree and local clustering coefficient. This tendency is confirmed by the negative assortative mixing measured in both networks (see Fig. 2), defined by the

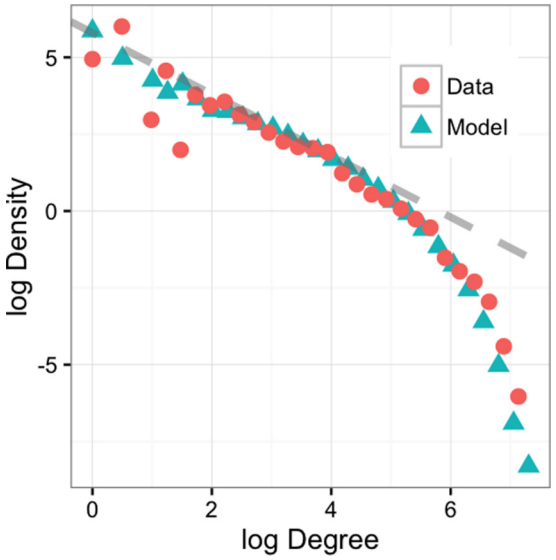

FIG. 3. Modeling the Darknet structure in 2015. The degree distribution obtained from an ensemble of 50 random realizations of our model is compared against the empirical distribution (the dashed line is for guidance only, to show the deviation of the degree distribution from a pure power law with scaling exponent equal to -1). The natural logarithm is considered in the figure.

Pearson's correlation coefficient of the degree of linked pairs of nodes [28,29].

From this structural analysis we find that while structural descriptors of the Internet do not change over time, this is not the case for the Darknet, which is still evolving (see Fig. 2 for other structural descriptors and their evolution between 2013 and 2015). Nevertheless, while the Internet has been widely investigated and several models have been proposed to explain its structure, the peculiar properties of the Darknet and the previous unavailability of data about its structure call for a model that is able to reproduce its most salient characteristics.

B. Modeling the Darknet structure

To model the Darknet, it is crucial to understand how the Tor network functions. Any Tor client initially queries the Directory Authorities to get the consensus, a table providing information about all active nodes in the network and their metadata. The metadata are then evaluated to build a circuit, a chain of three nodes used to connect the client to the server where the (possibly hidden) service is located. The choice of the nodes of the circuit is subjected to severe constraints, by default. For instance, the same node cannot be chosen twice and nodes run by the same operator are usually avoided. The choice of the first node is not performed uniformly at random: Instead, nodes with the largest bandwidth are favored, with priority to long-lived nodes called guard nodes. Guard nodes were previously relay nodes, i.e., common routers in the Tor network, that have been flagged as suitable for the role of entry point according to specific parameters, including high security level and traffic load balancing for end users.

At variance with the Internet's autonomous systems, whose connections are physical, our information about the Darknet is functional because it works at a different layer. Connectivity inferred by probes, as described in the previous section, provides a proxy that is used for our analysis. We model the above procedure with the following simple growing model and instead of analyzing the dynamics of the system, we allow it to grow until it reaches the size of the empirical networks we have. Therefore, we stop the growing process and analyze the resulting static snapshot, as for the data. At time $\tau = 0$ a small random network with $n_0 \ll N$ nodes and e_0 edges is created first. We assign a time stamp $T(n)$ to such nodes ($n = 1,2,\ldots,n_0$), which will allow us to calculate their age $\mathcal{A}(n,\tau)$ later at any time τ, and a property $\mathcal{B}(n)$, not varying over time, whose value is sampled from a heavy-tailed distribution, to mimic the empirical bandwidth distribution. In the following, we will use a log-normal distribution.

At each time step $\tau = 1,2,\ldots,N - n_0$, a new node n and M links enter the network. The time stamp $T(n) = \tau$ and a bandwidth are assigned to n, as previously described. It is crucial to remark that the M links do not necessarily involve

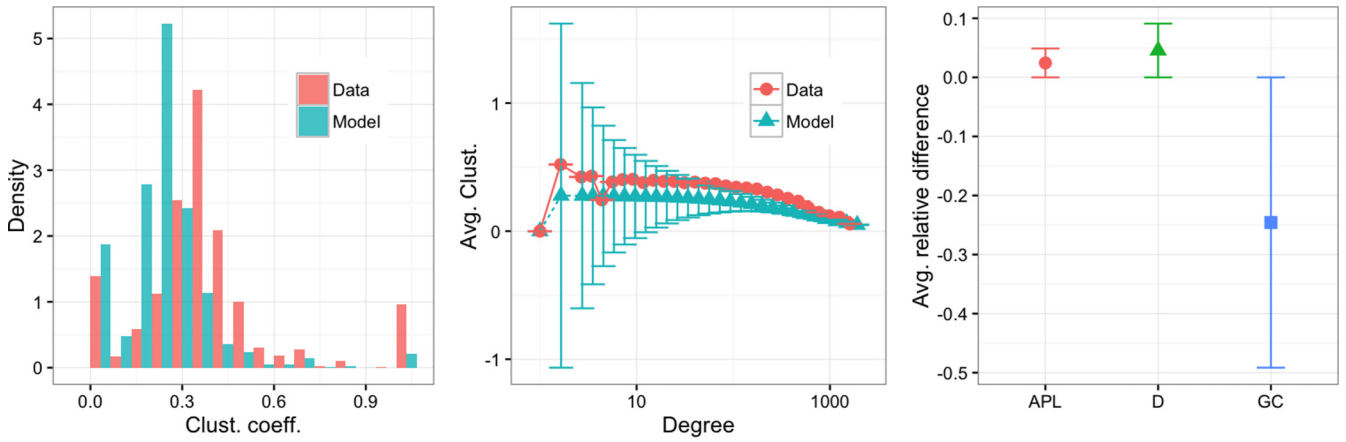

FIG. 4. Reproducing the Darknet (2015) structural descriptors. Different structural descriptors obtained from an ensemble of 400 random realizations of our model are compared against the observed values. In the last panel, we report the mean relative difference between the value calculated from the data and the values calculated from the model (D is for diameter, APL is for average path length, and GC is for global clustering).

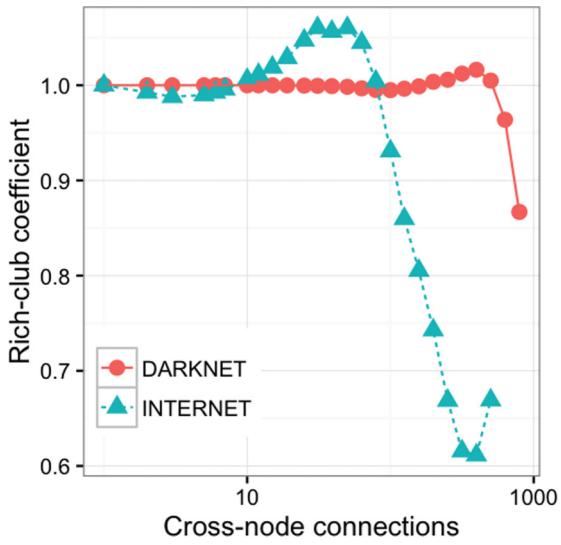

FIG. 5. Rich club structure of the Internet and Darknet in 2015. The ratio between the rich-club coefficient calculated for the empirical networks and its random expectation (see the text for further details) is plotted as a function of the degree threshold. A coefficient is not shown for degree values for which the number of nodes is smaller than 100 (about 2% of the Darknet and 0.2% of the Internet).

node n, at variance with processes based on the traditional preferential attachment. In fact, in a trust network like the Darknet, new nodes have to increase their reputation before being trusted and this is more likely to happen with aging. Nevertheless, links have to be created between trusted nodes at time τ, therefore $2M$ nodes are randomly chosen with probability

$$p_{n'}(\tau) = \frac{\mathcal{A}^{\beta}(n',\tau)\mathcal{B}^{\gamma}(n')}{\sum_{i=1}^{n}\mathcal{A}^{\beta}(i,\tau)\mathcal{B}^{\gamma}(i)}, \tag{1}$$

where $\mathcal{A}(n',\tau) = \tau - T(n')$ is the age of node n' at time τ. The nodes are then randomly linked into M pairs. Crucially, the degree of each node at each time step does not play any role in the growing process, which is completely driven by exogenous node properties such aging and bandwidth. The final step is to rewire nodes randomly while preserving the degree distribution: This last stage destroys possible structural correlations due to the previous stages of the model and it is crucial to introduce a higher level of randomness in connectivity.

This Darknet stochastic model is very general and, varying the exponents β and γ, it is possible to explore different scenarios, e.g., where probability is inversely proportional

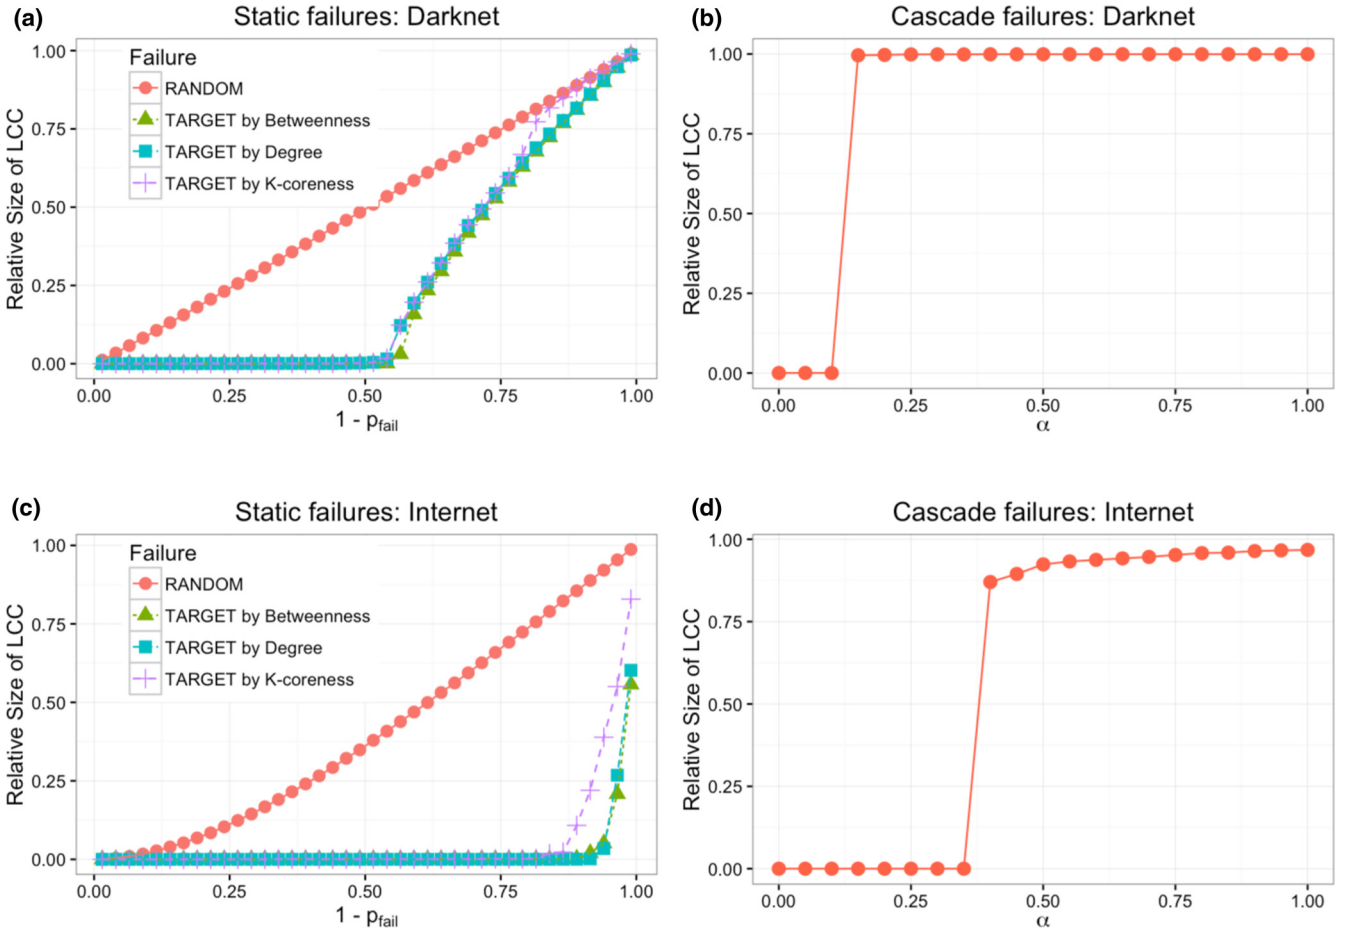

FIG. 6. Resilience of the Darknet and the Internet in 2013 to topological and dynamical attacks. Shown is the relative size of the largest connected component of (a) the Internet and (c) the Darknet after disrupting a fraction p_{fail} of nodes uniformly random or by targeted attacks (with respect to degree, betweenness, and k-coreness). Also shown is the relative size of the largest connected component of (b) the Internet and (d) the Darknet after dynamical disruptions due to induced cascade failures.

to age ($\beta < 0$) and proportional to bandwidth ($\gamma > 0$). In practice, the value of M is fixed by the data as $M = (E - e_0)/(N - n_0)$ and therefore the only free parameters of the model are in general β and γ. In the following, we do not fit the parameters and we just consider the simplest scenario with linear proportionality ($\beta = \gamma = 1$).

Our simulations revealed that this model generates networks that are remarkably close to the observed one from different perspectives. The high clustering and small characteristic length are satisfactorily reproduced, although the main finding here is that the degree distribution of simulated networks reproduces with excellent accuracy the empirical one, as shown in Fig. 3.

This result is of particular interest because, in general, networks with degree distribution scaling as power laws with exponent smaller than -2 (and especially close to -1) and a cutoff are very difficult to model. Mechanisms involving degree-based preferential attachment and the influence of some exogenous properties, such as aging and cost, have been proposed [30,31] (see Ref. [32] and references therein for a thorough review), although they are able to reproduce power-law scaling with exponent equal to or larger than -2 with a cutoff.

The ensemble of random realizations of our model is sufficient to reproduce the main topological properties of the Darknet, including its structural resilience, as we will see later. Figure 4 shows the comparison between observed values and their expectation according to our model.

C. Lack of "rich-club" effect in the Darknet

It has been shown that in many complex networks, especially the Internet, nodes that are very central tend to interconnect more with each other. This effect, called a rich club, produces a core of nodes that is really important for the stability and the robustness of the network and can be quantified [33,34]. Let us denote by $E^{[\xi]}_{>k}(t)$ the number of connections among the $N^{[\xi]}_{>k}(t)$ nodes with degree larger than the threshold k. The rich-club coefficient is defined by

$$\phi^{[\xi]}_k(t) = \frac{2E^{[\xi]}_{>k}(t)}{N^{[\xi]}_{>k}(t)\left[N^{[\xi]}_{>k}(t) - 1\right]}. \quad (2)$$

However, to understand to what extent the observed rich-club effect is not due to chance, we generate an ensemble of 1000 random networks preserving the empirical degree distribution and calculate the expected coefficient $\tilde{\phi}^{[\xi]}_k(t)$. Therefore, we study how the ratio $\phi^{[\xi]}_k(t)/\tilde{\phi}^{[\xi]}_k(t)$ changes as a function of k: When this ratio is close to 1 the observed rich club is compatible with random fluctuations, whereas when it is larger (smaller) than 1 it indicates the existence (absence) of a rich core of nodes. We show in Fig. 5 the calculated value of the

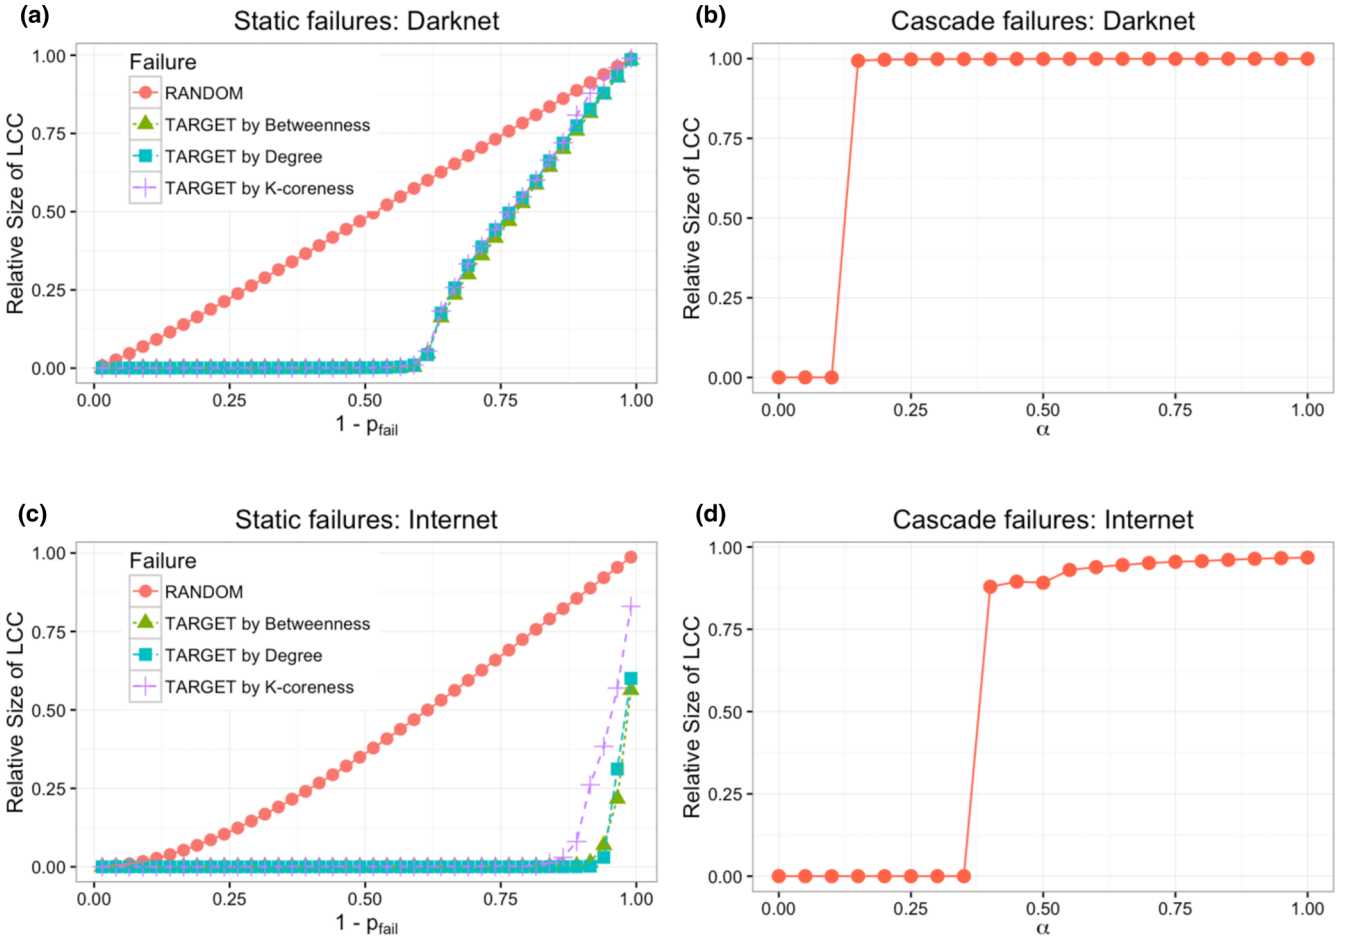

FIG. 7. Resilience of the Darknet and the Internet in 2014 to topological and dynamical attacks.

ratio for the Internet and the Darknet in 2015. The Internet exhibits a clear rich-club effect for nodes with intermediate degree, around $k = 50$, with largest hubs tending to be not interconnected with each other. Conversely, the Darknet does not exhibit a rich core, with a slight tendency of the largest hubs to be not interconnected with each other, an effect that is in magnitude significantly smaller than the case of the Internet.

Summing up, the Internet consists of a backbone of high-centrality nodes, whereas the Darknet does not. This result is compatible with the fact that nowadays the Internet is a very centralized network providing an easier way to manage and search for online services, whereas the Darknet is very decentralized (as the Usenet, the ancestor of the Internet) but it is more difficult to manage and search for hidden online services.

IV. RESILIENCE OF THE DARKNET

A. Resilience to static failures

Here we investigate how the structural properties of the Darknet and the Internet are reflected in their resilience to perturbations. We consider three different types of disturbances based on topological and dynamical perturbations. Topological perturbations are static removals of nodes that might mimic either random disruptions or targeted attacks [1]. Dynamical perturbations start with the disruption of a single node, generally the one with the highest degree, which triggers a cascade of failures [21,22].

In random disruptions, a fraction p_{fail} of nodes is chosen uniformly at random in the network and removed. In targeted disruptions, the fraction p_{fail} of nodes is chosen according to their ranking with respect to a measure of centrality. Usually, the degree is used, but also the betweenness [35], quantifying centrality with respect to the communication flow, and k-coreness [36], based on the core decomposition of a network and characterizing which nested shell a node belongs to. It is common to quantify the resilience of a network to such perturbations by observing how the relative size of the largest connected component changes as a function of $1 - p_{\text{fail}}$, i.e., the fraction of surviving nodes. This method allows us to quantify if the survived nodes are all together on the same component or if they form small disconnected components that hinder the network's function.

Nonhomogeneous random networks are known to be very robust to random disruptions but very sensitive to targeted attacks. In fact, our findings confirm that both the Internet and the Darknet are fairly robust to random failures, whereas they are more damaged by targeted attacks (see Figs. 6–8). It is worth remarking that the critical point, i.e., the fraction of disruptions for which the largest connected component of the network is minimum, is very different for the two networks. In fact, while it is enough to target 10% of the Internet nodes to reach the critical point, in the case of the Darknet much more

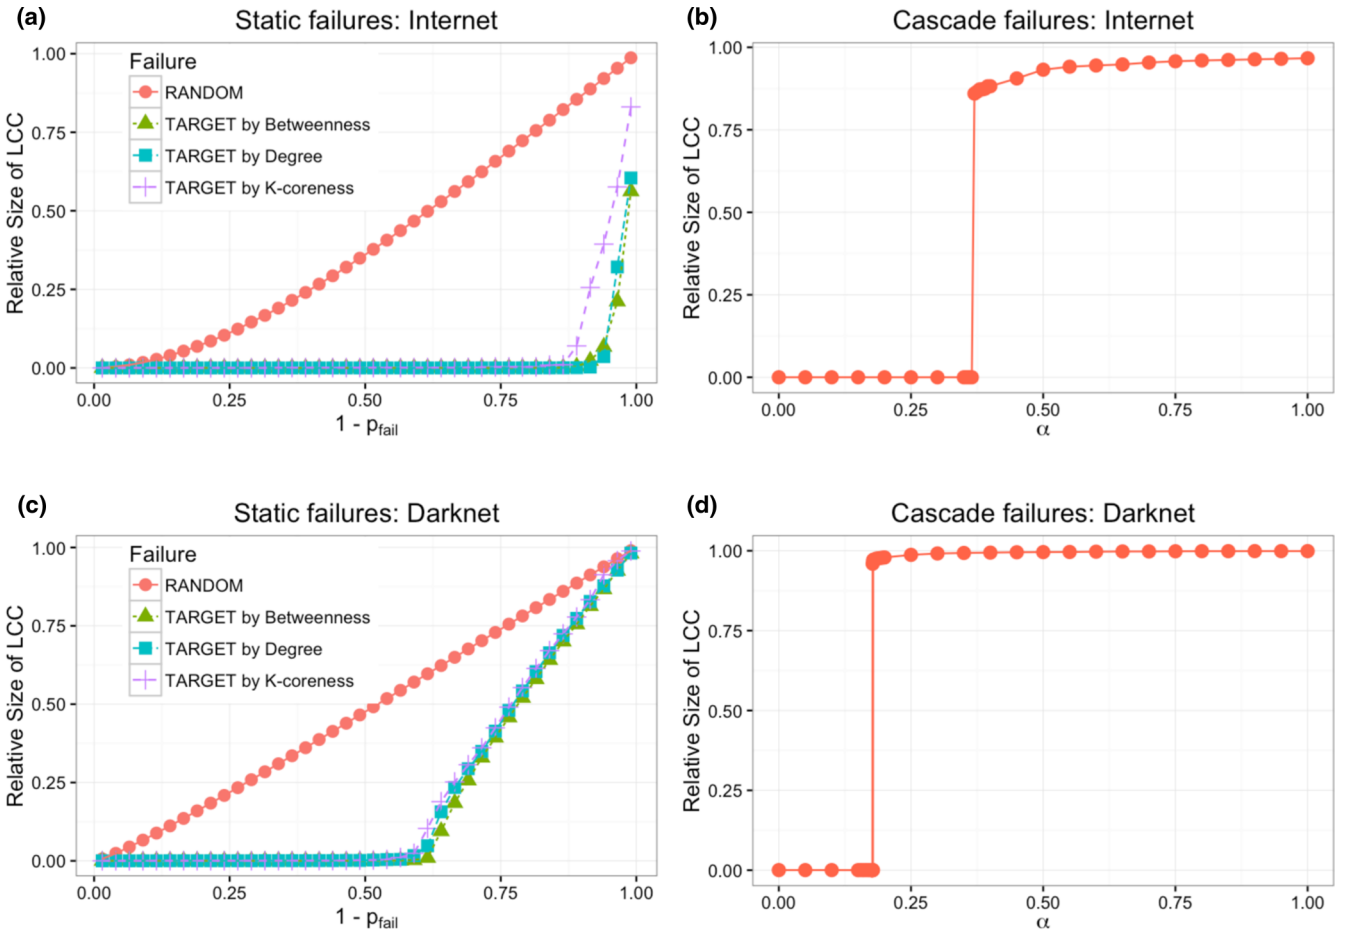

FIG. 8. Resilience of the Darknet and the Internet in 2015 to topological and dynamical attacks.

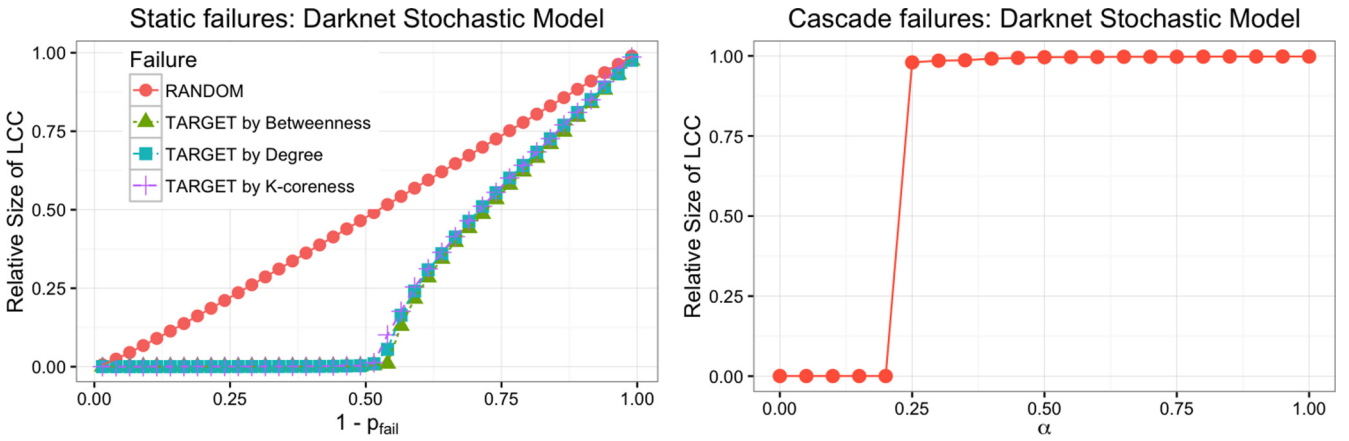

FIG. 9. Topological and dynamical resilience of simulated networks obtained from our model of the Darknet in 2015.

effort is needed, requiring 40% of the disruptions (this result is in excellent agreement with expectation from our model; see Fig. 9). The corresponding relative differences between the two networks are explicitly shown in Fig. 10(a), where it is evident that the Darknet is, by orders of magnitude, more resilient than the Internet, even with respect to random disruptions.

B. Resilience to dynamical failures

Another type of disruption, very suitable for communication networks, is based on inducing cascade failures. The rationale behind this method is that a node i in a communication network is characterized by a certain capacity $\mathcal{C}_i$, a fixed feature quantifying the maximum amount of load they can operate with, and a load $\mathcal{L}_i(\tau)$, a dynamical feature depending on the state of the network. Nodes with higher degree are assumed to be the ones with higher capacity, and at any time the total load of the network is constant, i.e., $\mathcal{L} = \sum_i \mathcal{L}_i(\tau)$. If a node with high capacity is disrupted, its load must be redistributed among the other nodes of the network; however, if the new loads exceed their capacities, a new set of nodes will suffer a disruption, redistributing the loads through the remaining nodes and so on, thus generating a cascade of failures that can paralyze the system. The dynamics of cascade failures and the resilience of the network can be studied as a function of a parameter α, which improves the capacity of each node to $(1+\alpha)\mathcal{C}_i$. By varying α and calculating the relative size of the largest connected component at the end of the cascade, we can estimate the required enhancement in capacity to make the network resilient to this attack. The detailed results are shown in Figs. 6–8, with the differences shown in Fig. 10(b). Again, the Darknet is much more resilient than the Internet to this catastrophic cascade of failures, requiring just $\alpha \approx 0.2$ to remain fully operative, whereas the Internet requires at least $\alpha = 0.28$ to keep almost 90% of its nodes operative (full operation is guaranteed for values of α close to 1). This result is, one more time, in excellent agreement with expectation from our model (see Fig. 9). The relative differences between the resilience of the two networks clearly indicate that before and close to the critical point the Darknet is more resilient than the Internet. This property has direct economic implications, because the larger value of α increases the costs to make the network more robust.

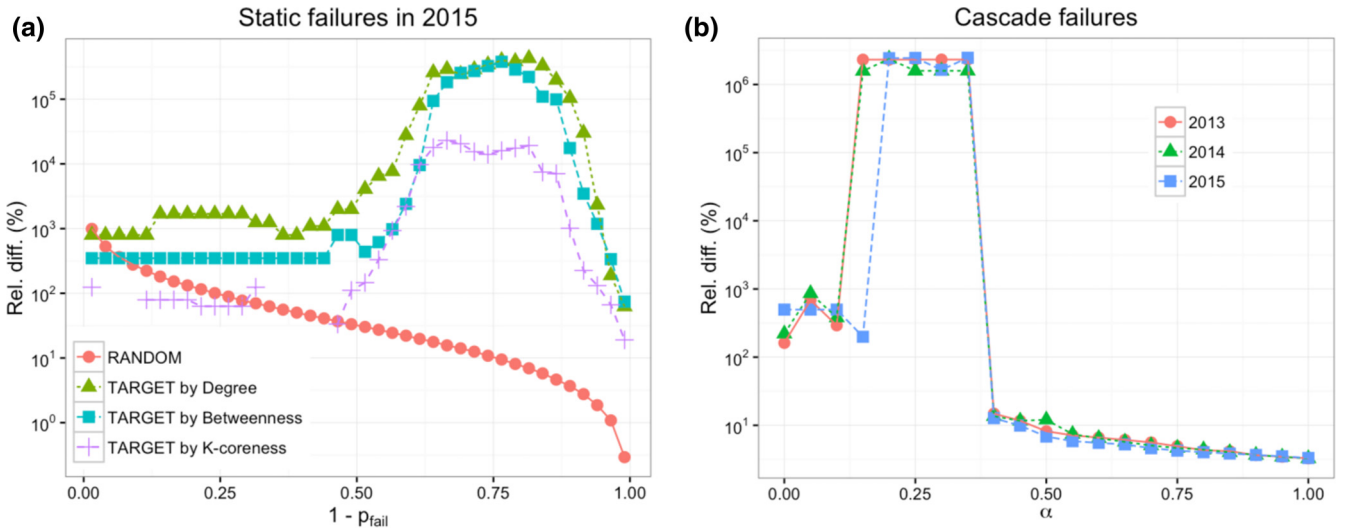

FIG. 10. Relative differences in resilience: (a) differences between the Darknet and the Internet in resilience to random and targeted attacks (see Fig. 11 for the differences in topological resilience in 2013 and 2014) and (b) induced cascade failures.

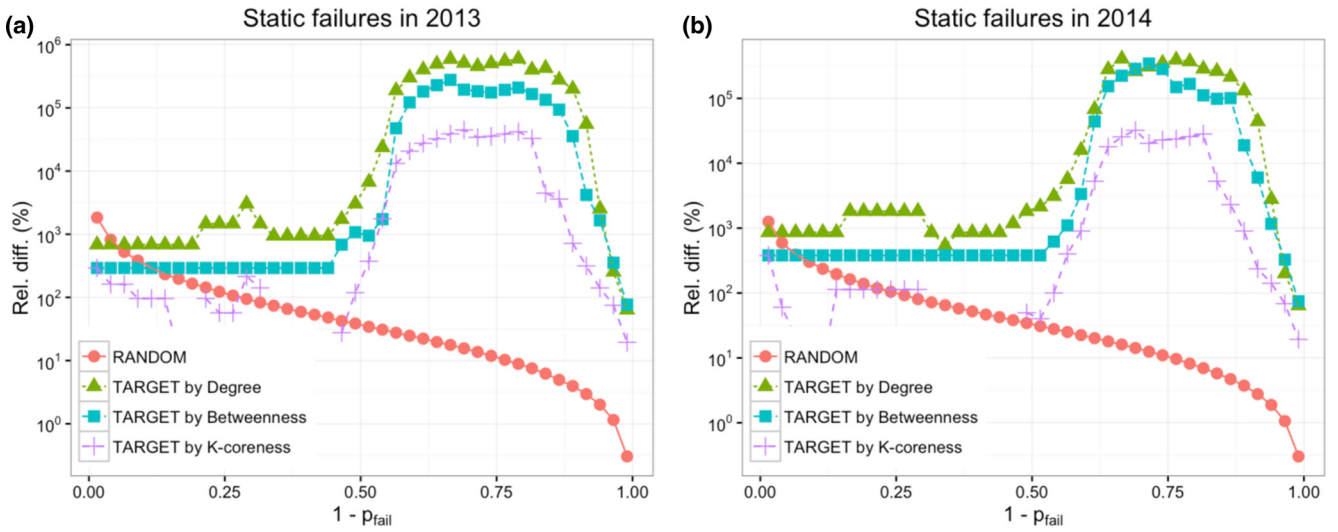

FIG. 11. Relative difference between the resilience of the Internet and the Darknet (we used the curves shown in Figs. 6 and 7 corresponding to topological resilience).

V. CONCLUSION

We have investigated the structural properties of the Darknet, the communication network developed in the past two decades to guarantee safe and anonymous navigation. The Darknet exhibits some interesting features that are not shared by the structure of the Internet. Triadic closure in the Darknet is more likely than in the Internet, with communication paths much shorter in the former. Like the Internet, the Darknet is characterized by a nonhomogeneous connectivity distribution and the presence of higher-order degree-degree correlations. However, the topology of the Darknet is more interesting because of the peculiar heavy-tailed scaling of the degree distribution, with a scaling exponent close to -1 and a cutoff, at variance with the Internet, appearing more like a power law with a scaling exponent close to -2 and no evident cutoff.

The rich-club analysis has revealed the lack of a core of highly central nodes interconnected each other, at variance with the Internet, where this effect is remarkable. We argue that such topological differences are responsible for the different resilience exhibited by the two communication systems in response to random disruption, target attacks, and induced cascade failures. We have thoroughly shown that the peculiar topology of the Darknet, characterized by highly clustered communication circuits, a small characteristic distance between hops, and lack of a rich core, makes this network much more resilient than the Internet as a result of adaptive changes in response to the attempts of dismantling it across time.

While the resilience of the Internet has not significantly changed over time, the resilience of the Darknet is still changing. In fact, its resilience to topological disruption slightly decreased between 2013 and 2014, remaining unchanged in 2015, whereas in the same year the Darknet became slightly less resilient to induced cascade failures. Together with the trends revealed by other structural descriptors, such as decreasing clustering, slightly increasing characteristic length, and increasing assortativity, we argue that the Darknet might be undergoing a transition from decentralization to centralization of its services. It will be interesting to confirm this prediction in the future, when more historical data will be available.

By mimicking how the Darknet actually works, we have proposed a model, based on a preferential attachment mechanism depending on exogenous properties such as aging and bandwidth and independent of endogenous properties such as node degree, to reproduce the empirical degree distribution with remarkable accuracy. The analysis shows that our model is sufficient to understand the structural correlations and the robustness of the Darknet.

We recognize that there are many possible flaws in the representation of both systems, the Darknet and the Internet, to claim to have the true topological networks. Nevertheless, from the data available it is indeed possible to start thinking about properties of the structure and its implications. The comparison between both structures has been developed at a descriptive level from the underlying graphs, with no other goal than using the Internet as a benchmark to better highlight some salient features of the Darknet. Finally, the proposal of a model driven by the main features of the Darknet seems to be enough to capture its more prominent topological features.

Summing up, the main result of this work is the indication that the mechanisms adopted to guarantee anonymous traffic in the Darknet and to improve the cybersecurity of its users could be the main reason for the peculiar topology of the Darknet observed so far and its resilience.

ACKNOWLEDGMENTS

The authors thank Robert Annessi and Martin Schmiedecker for providing the Tor raw data and support on data processing. M.D.D. acknowledges financial support from the Spanish program Juan de la Cierva (Grant No. IJCI-2014-20225). A.A. acknowledges financial support from ICREA Academia and James S. McDonnell Foundation and Spanish MINECO Grant No. FIS2015-71582-C2-1-P.

[1] R. Albert, H. Jeong, and A.-L. Barabási, Nature (London) **406**, 378 (2000).
[2] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, Phys. Rev. Lett. **85**, 4626 (2000).
[3] D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts, Phys. Rev. Lett. **85**, 5468 (2000).
[4] R. Pastor-Satorras, A. Vázquez, and A. Vespignani, Phys. Rev. Lett. **87**, 258701 (2001).
[5] R. Pastor-Satorras and A. Vespignani, *Evolution and Structure of the Internet: A Statistical Physics Approach* (Cambridge University Press, Cambridge, 2007).
[6] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).
[7] A. Vázquez, M. Boguñá, Y. Moreno, R. Pastor-Satorras, and A. Vespignani, Phys. Rev. E **67**, 046111 (2003).
[8] M. A. Serrano, M. Bogñá, and A. Díaz-Guilera, Phys. Rev. Lett. **94**, 038701 (2005).
[9] M. Á. Serrano, M. Boguñá, and A. Díaz-Guilera, Eur. Phys. J. B **50**, 249 (2006).
[10] M. Boguñá, F. Papadopoulos, and D. Krioukov, Nat. Commun. **1**, 62 (2010)
[11] F. Papadopoulos, M. Kitsak, M. Á. Serrano, M. Boguñá, and D. Krioukov, Nature (London) **489**, 537 (2012).
[12] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker *et al.*, IEEE Micro **19**, 50 (1999).
[13] B. Fortz and M. Thorup, Comput. Opt. Appl. **29**, 13 (2004).
[14] J. C. Doyle, D. L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger, Proc. Natl. Acad. Sci. USA **102**, 14497 (2005).
[15] M. Boguna, D. Krioukov, and K. C. Claffy, Nat. Phys. **5**, 74 (2009).
[16] M. Boguná and D. Krioukov, Phys. Rev. Lett. **102**, 058701 (2009).
[17] B. Briscoe and J. Manner, Internet Eng. Task Force **RFC**, 7141 (2014).
[18] P. F. Syverson, M. G. Reed, and D. M. Goldschlag, J. Comput. Security **5**, 237 (1997).
[19] D. Goldschlag, M. Reed, and P. Syverson, Commun. ACM **42**, 39 (1999).
[20] The Guardian, NSA and GCHQ target Tor network that protects anonymity of web users, accessed 20 May 2016, http://goo.gl/UiZyPd
[21] A. E. Motter and Y.-C. Lai, Phys. Rev. E **66**, 065102 (2002).
[22] A. E. Motter, Phys. Rev. Lett. **93**, 098701 (2004).
[23] Internet Research Lab, Computer Science Department, UCLA, Internet AS-level Topology Archive, http://irl.cs.ucla.edu/topology/
[24] R. Annessi and M. Schmiedecker, *Proceedings of the First IEEE European Symposium on Security and Privacy, Saabrücken, 2016* (IEEE, Piscataway, 2016).
[25] R. Annessi and M. Schmiedecker, NavigaTor, https://navigator.github.io/
[26] D. J. Watts and S. H. Strogatz, Nature (London) **393**, 440 (1998).
[27] R. Cohen and S. Havlin, Phys. Rev. Lett. **90**, 058701 (2003).
[28] M. E. J. Newman, Phys. Rev. Lett. **89**, 208701 (2002).
[29] M. E. J. Newman, Phys. Rev. E **67**, 026126 (2003).
[30] L. A. N. Amaral, A. Scala, M. Barthelemy, and H. E. Stanley, Proc. Natl. Acad. Sci. USA **97**, 11149 (2000).
[31] P. L. Krapivsky, G. J. Rodgers, and S. Redner, Phys. Rev. Lett. **86**, 5401 (2001).
[32] S. N. Dorogovtsev and J. F. Mendes, Adv. Phys. **51**, 1079 (2002).
[33] V. Colizza, A. Flammini, M. A. Serrano, and A. Vespignani, Nat. Phys. **2**, 110 (2006).
[34] T. Opsahl, V. Colizza, P. Panzarasa, and J. J. Ramasco, Phys. Rev. Lett. **101**, 168702 (2008).
[35] L. C. Freeman, Sociometry **40**, 35 (1977).
[36] S. B. Seidman, Social Networks **5**, 269 (1983).