The “Robust Yet Fragile” Nature of the Internet

John C. Doyle[1], David Alderson[1], Lun Li[1], Steven Low[1], Matthew Roughan[2], Stanislav Shalunov[3], Reiko Tanaka[4], and Walter Willinger[5]

[1]Engineering and Applied Sciences Division, California Institute of Technology, Pasadena, CA 91125, USA; [2]Applied Mathematics, The University of Adelaide, SA 5005, Australia; [3]Internet2, 3025 Boardwalk Dr, Suite 200, Ann Arbor, MI 48108, USA; [4]RIKEN, Bio-Mimetic Control Research Center, Nagoya, Japan; [5]AT&T Labs–Research, Florham Park, NJ, USA

Classification: Physical Sciences (Engineering)

Manuscript information: 16 pages, 4 figures, 1 table

Word and character counts: Abstract (154 words), paper (46,966 characters)

Abbreviations footnote: SF, scale free; RYF, robust yet fragile; ISP, Internet Service Provider; HOT, Highly Optimized Tolerance or Heuristically Optimized Tradeoffs; IP, Internet protocol.

Corresponding author: John C. Doyle
Engineering and Applied Sciences Division
California Institute of Technology
MC 107-81
Pasadena, CA 91125, USA
Phone: (626) 395-4808
Fax: (626)796-8914
Email: doyle@cds.caltech.edu

Abstract

The search for unifying properties of complex networks is popular, challenging, and important. For modeling approaches that focus on robustness and fragility as unifying concepts, the Internet is an especially attractive case study, mainly because its applications are ubiquitous and pervasive, and widely available expositions exist at every level of detail. Nevertheless, alternative approaches to modeling the Internet often make extremely different assumptions and derive opposite conclusions about fundamental properties of one and the same system. Fortunately, a detailed understanding of Internet technology combined with a unique ability to measure the network means that these differences can be thoroughly understood and unambiguously resolved. This paper aims to make recent results of this process accessible beyond Internet specialists to the broader scientific community, and to clarify several sources of basic methodological differences that are relevant beyond either the Internet or the two specific approaches focused on here; i.e., *scale-free networks* and *highly optimized tolerance* networks.

A popular case study for complex networks has been the Internet, with a central issue the extent to which its design and evolution have made it “robust yet fragile” (RYF)—that is, unaffected by random component failures but vulnerable to targeted attacks on its key components. One line of research portrays the Internet as “scale-free” (SF) with a “hub-like” core structure that makes the network simultaneously robust to random losses of nodes yet fragile to targeted attacks on the highly connected nodes or “hubs” [1, 2, 3]. The resulting error tolerance yet attack vulnerability has been proposed as a previously overlooked “Achilles’ heel” of the Internet. The appeal of such a surprising discovery is understandable, as SF methods are quite general and do not depend on any details of Internet technology, economics, or engineering [4, 5].

One purpose of this paper is to explore how this SF depiction compares with the real Internet, and explain the nature and origin of some important discrepancies. Another purpose is to suggest that a more coherent perspective on the Internet as a complex network, and in particular its RYF nature, is possible in a way that is fully consistent with Internet technology, economics, and engineering. A complete exposition relies on the mathematics of random graphs and statistical physics [6], which underly the SF theory, as well as on the very details of the Internet ignored in the SF formulation [7]. Nevertheless, we aim to show here that the essential issues can be readily understood, if not rigorously proven, using less technical detail, and the lessons learned are relevant well beyond either the

Internet or SF network models [8, 9, 10].

Power Laws and Scale-Free Models

One widespread focus of attention has been on *power laws* (or *scaling*) in graph vertex connectivity. For a graph having n vertices, let $d_i = deg(i)$ denote the degree of vertex i, $1 \leq i \leq n$. We call $D = \{d_1, d_2, \ldots, d_n\}$ the *degree sequence* of the graph, assumed without loss of generality always to be ordered $d_1 \geq d_2 \geq \ldots \geq d_n$. Let $G(D)$ denote the set of all connected simple graphs (i.e., no self-loops or parallel edges) having the same graph degree D. We will say graphs $g \in G(D)$ have *scaling degree sequence* D (or D *is scaling*) if for all $1 \leq k \leq n_s \leq n$, D satisfies a power-law *rank-size relationship* of the form $kd_k^{\alpha} \approx c$ where $0 < c$ and $0 < \alpha$, are constants and n_s determines the range of scaling [11]. Since scaling implies $\log(k) + \alpha \log(d_k) \approx log(c)$, doubly logarithmic plots of degree d_k versus rank k yield approximately straight lines of slope $-\alpha$. In contrast, exponential rank-size relationships (i.e., $ke^{\lambda d_k} \approx c$) result in approximately straight lines on semi-logarithmic plots.

The most significant SF claims for the Internet are that the router graph has power law degree sequences and these give rise to "hubs", which by SF definition are highly connected vertices that are crucial to the global connectivity of the network and through which most traffic must pass [3]. The SF assertion (later formalized in [12]) is that such hubs holds the network together, giving it "error tolerance" to random vertex failures since most vertices have low connectivity (i.e. are non-hubs), but also "attack vulnerability" to targeted hub removal, a previously overlooked "Achilles' heel." The rationale for this claim can be illustrated using the toy networks in Figure 1, all of which have the identical scaling degree sequence D shown in Figure 1(e). Figure 1(a) shows a graph (size issues notwithstanding) that is representative of the type of structure typically found in graphs generated by SF models, in this case preferential attachment (PA). This graph is drawn in two ways, the left and right visualizations emphasizing the growth process and Internet properties, respectively. Clearly, the highest degree nodes are essential for graph connectivity, and this feature can be seen even more clearly for the more idealized SF graph in Figure 1(b). Thus the SF claims would certainly hold if the Internet looked at all like Figure 1(a) and (b). As we will see the Internet looks nothing like these graphs and is much closer to Figure 1(d), which has the same degree sequence D but is otherwise completely different, with high degree vertices at the periphery of the network, where their removal would have only local effects. Thus, while scaling degree sequences imply the presence of high-degree vertices, they do not imply that such nodes form necessarily "crucial hubs" in the SF sense.

The deeper origins of the claims involving power laws and hubs arise from the SF models' roots in statistical physics, where any particular graph is interpreted as an element from a larger statistical ensemble of graphs, with probability weights that typically arise either implicitly via some underlying stochastic generation process or by a mechanism that explicitly assigns a weight to each element of the ensemble [13, 14]. While there exist a variety of methods for generating ensembles of graphs having scaling degree sequences, including PA, *Generalized Random Graph (GRG)*, *Power Law Random Graph (PLRG)* [15], as well as random degree-preserving rewiring [16]), the resulting models are widely conjectured to be asymptotically equivalent (see for example [6] and references therein).

In particular, for a graph g having degree sequence D we define the purely graph-theoretic quantity $s(g) = \sum_{(i,j)\in E(g)} d_i d_j$, where $E(g)$ is the set of edges in the graph. It is easy to check that high $s(g)$ requires high-degree vertices to connect to other high-degree vertices. Normalizing against $s_{max} = \max\{s(g) : g \in G(D)\}$, where $G(D)$ denotes the set of all simple, connected graphs having degree sequence D, we define the measure $0 \leq S(g) \leq 1$ of the graph g as $S(g) = s(g)/s_{\max}$. Although $s(g)$ and $S(g)$ can be computed for any graph and do not depend on any particular construction mechanism, they have a special meaning in the context of ensembles of graphs. Specifically, $S(g)$ has a direct interpretation as the *relative log-likelihood* of a graph resulting from the GRG construction [17], and thus all of the SF model generation mechanisms generate essentially only high S graphs. The S-metric also potentially unifies other aspects of SF graphs, as it is closely related to betweeness, degree correlation [6], and graph assortativity [18], and captures several notions of self-similarity related to graph trimming, coarse-graining, and random rewiring [6].

The focus on ensemble-based methods means that the analysis in SF models has implicitly ignored those graphs that are unlikely to result from such constructions, in particular graphs with small S. Thus while power law degree distributions are unlikely under some traditional random graph constructions (e.g, Erdös-Renyí random graphs [19]), there are a multitude of other model generation mechanisms that give rise to power laws [20]. The SF generating mechanisms are only one kind, but these only generate high S graphs, and this leaves unexplored an enormous diversity of low S graphs, as seen in Figures 1. The graphs in Figures 1(a-b) are relatively likely to result from probabilistic construction while the graphs in Figures 1(c-d) are vanishingly unlikely. The PA-type graph in Figure 1(a) has $S(g_a) = 0.61$ and is typical of the graphs that are likely under a variety of random generation methods. The graph in Figure 1(b) is the $s_{\max}$ graph and thus by definition has $S(g_b) = 1.0$. It can be thought of both as the most likely graph and also (uniquely) as the most “perfectly” scale-free graph with this degree sequence. Of course the sheer enormity of the number of different high S graphs means that any particular one graph, even the *relatively* most likely, is actually unlikely in absolute terms to be selected. The graphs in Figures 1(c-d) have respective values $S(g_c) = 0.33$ and $S(g_d) = 0.34$, furthermore there are relatively few graphs with S values this low, and thus any graphs like these are vanishingly unlikely to arise at random [6]. The remainder of this paper explains in more detail why the underlying forces at work in the evolution of the real router-level Internet avoid the generation of high S graphs and how this feature can be captured in an optimization-based design framework. We also consider what, if anything, this framework has to say about the RYF nature of the Internet.

A Look at the Actual Internet

An obvious starting point for investigating the structure and underlying forces at work in the Internet is to inspect detailed router-level maps from ISPs (Internet Service Providers). Abilene, the backbone for the Internet2 academic network, is illustrated in Figure 2 and is an ideal example for many reasons that will be exploited throughout this analysis [21]. Abilene publishes detailed hardware specifications for each router and link, so Figure 2 is exact, not an approximation based on indirect measurements. Abilene is also a state-of-the-art network with essentially no difference between physical (i.e. layer 2) and IP (i.e. layer 3) connectivity. This simplifies the exposition without loss of generality and also eliminates a source of confusion in measured data from networks that use older legacy technologies. Using regional academic networks and commercial ISPs, we verified that all the inferences and conclusions based on Abilene hold in general. Commercial ISPs do not allow publishing such details because of proprietary considerations, but router-level measurement studies [22, 23, 24] further confirm our analysis [7, 25, 26], although this requires additional statistical and Internet-specific expertise beyond the intended scope of this paper.

Figure 2 shows that Abilene is designed as a sparsely connected mesh of uniformly high-speed (10Gbps), long-range links between routers located in 11 major U.S. cities, with connectivity to regional and local networks provided by some minimal amount of redundancy. These design features are typical of ISP *backbones* (the main connections and routers composing an ISP's national or international network), which can differ in overall size but are qualitatively similar. One of their most obvious features is the complete absence of SF hubs; high-degree vertices can exist, but are found only within the local networks at the far periphery of the network and would not appear anywhere close to the backbone.

While the issue of whether or not the real Internet actually has power laws in its connectivity is beyond the scope of this article, there exists considerable evidence that existing claims of power laws in the router-level Internet may be the result of misinterpretation of available measurements and/or their naive and inappropriate statistical analysis (see [6] for a detailed discussion). It is certainly the case that current router technology could in principle support high variability (possibly scaling), but a closer look at existing router technology confirms that if high connectivity exists at all, it will be found toward the network periphery and not in its core. As we will show next, this is a consequence of the need for a high performance network.

Internet Modeling: An Optimization-Based Approach

Much of the topology of Abilene and other ISP backbones can be understood using annotated graphs with a few technological and economic constraints that provide a simple, yet surprisingly complete model of the essentials of network design [7]. *Highly Optimized/Organized Tolerance/Tradeoffs (HOT)* has been proposed as a conceptual framework for capturing the highly organized, optimized, and RYF structure of complex highly evolved systems [8]. HOT seeks an abstract but unified approach to diverse complex systems through models involving optimization of tradeoffs between multiple functional objectives of networks subject to constraints on their components, usually with an explicit source of uncertainty against which solutions must be tolerant, or robust. Constrained optimization and robustness are the universal themes, but models of function, uncertainty, component constraints, and environment are necessarily domain specific. One consistent result of the HOT framework has been that once functional performance and robustness tradeoffs are considered then in a variety of toy models engineering design [7] or biological evolution [9] easily generates power laws. This can occur in both deterministic and stochastic HOT models, including models motivated by physics [8]. Power laws have been a central focus of the “emergent complexity” view of SF and related methods, which arrive at them in completely different ways than HOT with its focus on “organized complexity”.

Here, we present toy networks reflecting the HOT approach to modeling the router-level Internet, which we will contrast with the corresponding SF network models. To this end, consider again the example network in Figure 1(d), which we argue captures the kind of essential domain-specific tradeoffs that occur in engineering. Accordingly, we refer to this toy network as *HOTnet*, although it is important to underscore that our results do not depend on designs being formally optimal, which is unlikely to occur in practice. Instead, we will argue that any sensible network design process with minimally realistic assumptions would produce something qualitatively similar.

A HOT model of the Internet’s router-level topology requires two general elements: constraints and functional objectives. First, the technological and economic constraints on components such as routers and links and their interconnection restrict what topologies are feasible or possible. Second, network backbones, and router-level connectivity more generally, are subsystems in the larger decentralized and layered Internet infrastructure. The consequence is that such subsystems can only be fully understood in terms of the functions they provide to the higher layers of the protocol stack and the rest of the network. The main purpose for building physical network infrastructures at the lower layers of the protocol stack is to carry effectively the expected or projected overall traffic demand generated at the higher layers, which in turn is ultimately driven by users at the application layer. Such teleological explanations are understandably avoided in physics but are essential for engineering networks, and this gap is responsible for much of the difference between the two approaches described in this paper.

A standard metric for *network performance* adopted here is the maximum throughput of the network under a "gravity model" of end user traffic demands [27]. It assumes that every end vertex i has a total bandwidth demand x_i, that two-way traffic is exchanged between all pairs (i, j) of end vertices i and j, the flow X_{ij} of traffic between i and j is given by $X_{ij} = \rho x_i x_j$, where ρ is some global constant, and is otherwise uncorrelated from all other flows. While more elaborate metrics are possible, this notion of network performance is a reasonable measure of the ability to provide a *fair* allocation of end user bandwidths. Our performance measure for a given network g is then its maximum throughput with gravity flows, computed as $P(g) = \max_\rho \sum_{ij} X_{ij}$, subject to $RX \leq B$, where R is the routing matrix obtained using standard shortest path routing. $R = [R_{kl}]$, with $R_{kl} = 1$ if flow l passes through router k, and $R_{kl} = 0$ otherwise. X is the vector of all flows X_{ij}, indexed to match the routing matrix R, and B is a vector consisting of all router bandwidth capacities.

The crucial elements of a design aimed at this notion of performance are realistic router capacities and economic considerations. Hardware technology fundamentally limits the number of data packets that can be processed per unit time, so routers must obey a form of *flow conservation* in the traffic that they handle. Although total router capacity is constantly increasing as hardware improves, this tradeoff in router utilization cannot be avoided. Figure 3(a) shows the router bandwidth-degree limits used in this model. In terms of economics, the cost of installing and operating physical links increases with link distance and can dominate the total budget for the global infrastructure, particularly in the backbone. While routers impose overall bandwidth limits, the backbone cost is primarily dominated by the installation and operation of links. This cost imposes strong incentives to minimize the number and length of deployed links by aggregating and multiplexing traffic at all levels of the network hierarchy, from the periphery to the core. Thus the combination of router technology and link costs *necessitate* that when moving from the periphery to the network core the link capacities, link lengths, and total router throughput generally increase while router degrees *decrease*. The result is possibly highly variable bandwidth and router degrees at the network's periphery, with necessarily a much greater uniformity of high bandwidth and low degree routers in the core.

As noted above, the network *HOTnet* shown in Figure 1(d) was inspired by the real Abilene network, and its overall connectivity was designed to achieve high performance while maintaining the scaling degree sequence in Figure 1(e). This network uses essentially the Abilene backbone as its core (the inner circle of routers in Figure 1(d)), and then assumes that end users (the outer circle) connect through small and greatly compressed single-level regional networks (the middle circle of vertices). This allows us to create a network that uses the same technology as the real Internet, but has a scaling degree sequence. In particular, this scaling vertex degree is achieved in a minimal but technologically plausible way by choosing a gravity model of end user traffic demands and then aggregating these end users with routers that have high variability in their connectivity but must satisfy a particular router technology constraint. While the resulting network in Figure 1(d) is far too compressed to look like the real Internet, it has the same performance objectives, constraints, and design principles, though simplified and shows that a scaling degree sequence is at least plausibly consistent with Internet technology and economics. It could also reasonably be argued that this design-driven toy model grossly oversimplifies real Internet technology and economics, but we next demonstrate that this type of model has superior explanatory power to alternatives which ignore them entirely.

Contrasting HOT and SF Models

In view of the empirical evidence and the engineering arguments against popular SF claims regarding the location and criticality of the highest-connectivity routers, we next quantify more precisely the qualitative observations made above to illuminate the key methodological differences behind these different approaches and their resulting models. In doing so, we consider again the four toy models from Figure 1 along with their most relevant properties. To contrast the features of graphs having the same scaling degree sequence, we first consider the network *HOTnet* in Figure 1(d) alongside the "most preferential" network in Figure 1(b), which we denote in the remainder of this paper as *SFnet*. In computing the performance of these two graphs, we observe that $P(HOTnet) = 5.76 \times 10^{11}$ bps while $P(SFnet) = 4.89 \times 10^{9}$ bps, a difference of more than two orders of magnitude.

This enormous performance difference can be understood by examining the *utilization* of individual routers within each network, as illustrated in Figure 3(a). This figure shows the overall feasible configuration region encapsulating the conservation between router degree and router throughput (measured in bandwidth) as discussed above and represented as B in the computation of performance. While greatly simplified for use here, this abstract representation for router bandwidth is consistent with real router technology [7], and it is adequate for our purposes since the resulting conclusions depend only on the most general features of Figure 3(a) and not on specific details. The unambiguous source of the poor *SFnet* performance is that the high-degree hubs become saturated and create severe bottlenecks, leaving the rest of the network with low overall utilization. In contrast, the connectivity in *HOTnet* is such that the core routers are highly utilized and therefore enable greater overall network throughput.

An additional view into the performance and utilization of these two networks is available by considering the distribution of bandwidth that is actually delivered to the end users in these two networks under maxflow conditions, as shown in Figure 3(b). The distribution of achieved end user bandwidth for *HOTnet* is highly variable spanning four orders of magnitude (as opposed to five or more found in real networks [25]), but is considerably higher than what is received by users in *SFnet* who get uniformly low bandwidth. Another issue not quantified here is that, no matter where the high-degree SF hubs were physically located, the link costs to connect them would be prohibitively high. In contrast, the design aspects incorporated into *HOTnet* ensure that the deployed routers are used efficiently, the network is able to satisfy end user bandwidth demands that are highly variable, and with relatively few long-range links. For network engineers, the combination of superior throughput, high router utilization, low link costs, and realistic end user bandwidth makes *HOTnet* highly desirable but *SFnet* a very poor design choice, even though networking reality dictates the need for some degree of overprovisioning that will result in a slightly less efficient network than *HOTnet*.

Another important comparison between the graphs *HOTnet* and *SFnet* is to investigate the presence of "Achilles' heel" hubs. Here, we will consider robustness to router failures, defining this robustness as the remaining performance of the network after routers are removed and after re-routing of traffic. That is, addressing the issue of network robustness for the Internet requires at a minimum incorporating a simple abstraction of IP routing that accounts for the feedback mechanism by which the real network "sees damage and works around it." Note that the main mechanism by which users improve robustness to network losses is through link redundancy (e.g., multi-homing) but this was not an objective of our heuristic HOT design. However, we can still illustrate the differences between *HOTnet* and *SFnet* in some limited way. Figure 3(c) shows the impact of deleting routers in succession, always taking the worst-case that has not yet been deleted. The measure of performance after deletion of a vertex is the amount of original traffic that can still be served by the remaining network after re-routing, but with routers that still have to adhere to their original bandwidth constraints.

Consistent with SF claims [3], the *SFnet* network is indeed fragile to the deletion of worst-case vertices (here, worse-case means highest-degree), but resilient to deletions of other vertices. In stark contrast, *HOTnet* is not only robust to worst-case deletions (here, worst-case are low-connectivity core vertices), but also shows high tolerance to deleting other vertices. In particular, loss of high-degree edge routers disconnects only low-bandwidth users and has no other effect on overall connectivity. Because *SFnet* has such poor nominal performance to start with, its performance is worse intact than *HOTnet* after the latter has sustained substantial damage. Thus the "Achilles' heel" claim for SF networks does not appear to hold simply on the basis of having a scaling degree sequence D, and we will provide one possible explanation for this difference in the next section. Fortunately, the actual Internet is more like *HOTnet* than *SFnet* and also has a great deal of additional robustness. Because the real Internet consists of multiple redundant *HOTnet*-type backbones that are moderately loaded, the ability to re-route traffic ensures that end users typically experience no discernible degradation in performance when core routers fail. In particular, the real Internet would never experience the type of separation of the network into disjoint components as claimed by the "Achilles' heel" hub argument unless massive losses occurred.

The *HOTnet* model in Figure 1(d) and the *SFnet* model in Figure 1(b) are just two points in the space $G(D)$ of simple connected graphs having identical scaling degree D (shown in Figure 1(e)). The space $G(D)$ is difficult to visualize, mainly because it is very diverse and has a combinatorially large number of elements. However, some aspects can be explored by projecting this high-dimensional space onto lower dimensions using macroscopic measures. Here, we leverage our previously defined notions of performance $P(g)$ and relative likelihood $S(g)$, and we show the values for our toy networks in Figure 4. The ability of the $P(g)$ and $S(g)$ measures to help differentiate among graphs in the space $G(D)$ is further illuminated by considering the two other networks shown in Figure 1.

Figure 1(c) depicts a graph having a heuristically "poor" engineering design (denoted *BADnet*), and Figure 1(a) depicts a graph having "random" connectivity (denoted *RNDnet*) and is typical of graphs grown by preferential attachment. Even though all four toy networks in Figure 1 are identical as far as their degree sequence D is concerned, three of them occupy completely opposite corners of the $P(g)$ vs. $S(g)$ plane. The *BADnet* network demonstrates that low $S(g)$ does not necessarily imply high performance, and in general graphs having low-$S(g)$ may be completely different from one another. In contrast, the *RNDnet* shows that other graphs resulting from SF models have the same poor qualitative and quantitative features as *SFnet*, and for the same reasons. We also observe that graphs having $S(g) \approx 1$ are much more alike (essentially unique) and our results to date suggest that if D exhibits high variability, it is impossible to have graphs $g \in G(D)$ with both high $S(g)$ *and* high $P(g)$.

Finally, we consider the graph operation of *pairwise degree-preserving rewiring*, whereby two randomly chosen edges are rewired but constrained to preserve the graph degree and network connectivity. It is easily shown that by a finite succession of such rewirings, any graph $g \in G(D)$ can be converted to any other, and thus this process provides a simple mechanism for exploring the space $G(D)$. The additional points in Figure 4 correspond to rewirings of *HOTnet* and *SFnet*, respectively, and demonstrate that while rewiring the former quickly produces networks with poor performance, rewiring the latter produces little improvement in performance. Note that the preferential attachment graph from Figure 4(a) is representative of a large number of graphs resulting from arbitrary pairwise rewiring, thus justifying its name as *RNDnet* and supporting the conjecture that all SF models generate essentially the same ensemble of graphs.

While far from comprehensive, the structural metric $S(g)$ provides some understanding of the diversity of graphs in the space $G(D)$. One striking feature of this view is that some of the most celebrated features of SF models, particularly the Achilles' heel vulnerability, appear to hold only for graphs having scaling degree sequence

D and high $S(g)$. It is not a necessary consequence of scaling alone, as it does not apply to networks like *HOTnet* that have low $S(g)$, even if they have the same scaling degree D. However, recalling that in addition to measuring the "hub-like" nature of a graph, $S(g)$ also has an interpretation as relative graph log-likelihood, the concentration of points in Figure 4 suggests that the vast majority of graphs resulting from SF models have relatively high likelihood of occurring and that the likelihood of recovering a graph like *HOTnet* via probabilistic construction is vanishingly small.

The Real RYF Internet

The preceding discussion suggests that probabilistic constructions are unlikely to capture the true router-level structure of the Internet and also that claims of a vulnerability in high-degree nodes are not supported by either engineering data or theory. The true RYF nature of the Internet is a complex and heavily studied issue, but we will sketch some central features. The perception of the Internet as a simple, robust, and homogeneous resource is the result of a layered architecture that utilizes multiple forms of feedback control which enable robust performance in the presence of frequent disruptions and enormous heterogeneity. The lowest layers of the protocol stack—involving the physical infrastructure such as routers and fiber-optic cables—have hard technological and economic constraints, but each higher layer defines its own often unique connectivity, and the corresponding network topologies become by design increasingly virtual and unconstrained. For example, in contrast to routers and physical links, the allowable connectivity of documents and virtual links in the World Wide Web (WWW) is designed to be essentially completely unconstrained.

An important feature of the Internet's highly organized but largely hidden complexity is to make the full system robust to the perturbations for which it was designed [28], but also potentially quite vulnerable to other perturbations [29]. All components must obey the protocols, but because of extensive feedback regulation, the overall system can tolerate otherwise enormous variability within these constraints and still deliver robust functionality to applications, which are also the least constrained components. Since the complete absence of a component is allowed, the system is thus by design robust to components which "fail off" by removal from the network, whether caused by focused attacks or other failures.

Note that it is protocols and feedback regulation and not simple redundancy per se that enables this extraordinary robustness. Another striking aspect of this robust design is a scalability, evolvability, and adaptability to exactly the kind of radical network change (i.e. in both hardware at the lower layers and applications at the highest layer) that the Internet has undergone in transforming from an academic research network to a critical component of the information infrastructure. Unfortunately, the Internet's strong robustness and adaptability coexists with an equally extreme fragility to components "failing on", particularly via malicious exploitation or hijacking of the very mechanisms that confer its robustness properties at higher levels in the protocol stack. Worms, viruses, spam, and denial of service attacks remain familiar examples [30]. This RYF tradeoff is a critical aspect of the Internet, and much research is devoted to enhancing these protocols in the face of new challenges. Thus, understanding Internet robustness requires a perspective that incorporates protocols, layering, and feedback regulation, and this view suggests that the most essential RYF features of the Internet actually come from aspects that are only indirectly related to graph connectivity.

The presentation here has emphasized the HOT framework as an alternate approach to SF models when considering the RYF nature of the Internet, and many other choices of functions and constraints are possible. Other researchers might emphasize alternative features that highlight particular tensions (e.g., design tradeoffs at different levels of the Internet protocol stack) and would be justified in doing so. The main point is the importance of incorporating issues such as performance, constraints, and tradeoffs—all of the things that make engineering different from physics—when considering the "essential" features of a highly evolved system. Here, we denote *highly evolved* systems as those resulting from an iterative design that incorporates tradeoffs between performance and the use of available resources. Thus, the RYF features of the Internet are the result of its highly evolved nature, and a key objective here has been to incorporate some of the most essential features in a simple model that can be used to highlight the potential dangers of ignoring such aspects entirely.

Conclusion

The most important SF claim is that the Internet has "hubs" that form an "Achilles' heel" through which most traffic flows and whose loss would fragment the Internet and constitute its attack vulnerability (i.e., global loss of connectivity). We have shown that there exist technological, economic, and graph theoretic reasons why this cannot be (and is not) the case for the current router-level Internet, even if the initial claim that the Internet's router-level graph has a scaling degree is assumed to hold. The comparison between HOT and SF networks in Figure 4 only distinguishes two dimensions, $P(g)$ vs. $S(g)$, but there are many more dimensions in which such comparisons could be made. Because SF generation mechanisms effectively yield only high S graphs, it is easy to characterize most properties that result and *SFnet* is perfectly representative. Similarly, HOT designs necessarily have highly structured characteristics that are also easily evaluated. Not surprisingly, *SFnet* and *HOTnet* are opposite in essentially every meaningful sense, and the real Internet network is much more like *HOTnet*.

Structure	**SFnet**	**HOTnet**	**Real Internet**
High Deg. Vertices	Core	Periphery	Periphery
Deg. Dist.	Power law	Power law	Highly variable
Generated by	Random	Design	Design
Core vertices	high degree	low degree	low degree
Throughput	Low	High	High
Attack tolerance	Fragile	Robust	Robust
Fragility	high-deg hubs	low-deg core	Hijack network

An additional question not addressed here is whether or not SF models, representing graphs with scaling degree sequence and characteristically high $S(g)$, may be appropriate representations at other layers of the Internet protocol stack. One answer is that since scaling is abundant at higher levels of the protocol stack [26], SF models could conceivably represent some *virtual* graphs associated with the Internet (such as, hypothetically, the WWW or other types of overlay networks). However, it is not clear what such models would have to say about the RYF nature of the Internet. For example, if high connectivity websites such as `Google` or `Yahoo!` are disabled, the rest of the WWW continues to function as it was designed to. At the same time, while previous applications of the HOT framework to the WWW have focused on the tradeoffs between performance and constraints in the design of individual websites [31], the HOT approach has not yet been applied to study the structure of the WWW as a whole.

It is certainly appealing that SF network models can avoid all Internet-specific structure such as protocol stacks, technological or economic constraints, and user heterogeneity yet make interesting and testable predictions. Unfortunately, this yields results that collapse when tested with real data or when examined by domain experts. This raises the more basic question of the applicability to highly evolved systems of unstructured, ensemble-based approaches, of which SF networks is just one example, and a largely parallel story in biology further suggests that the answer may be negative. Here interesting and testable SF claims about metabolic networks [32, 33] contrast sharply with both real data and concrete HOT models [9]. Again, functional descriptions and component constraints, such as conservation of energy and small moieties, the biochemical nature of underlying reactions, and the importance of robustness and evolvability prove essential [34]. However, while the router-level story here may be reflective of a broader debate about methodologies appropriate for complex networks, it is expected to take an even greater effort in domains like biology to reach the same level of clarity.

References

[1] Barabási, A.-L., & Albert, R. (1999) *Science* **286**, 509–512.

[2] Yook, S.-H., Jeong, H., & Barabási, A.-L. (2002) *Proc. Nat. Acad. Sci. USA* **99**, 13382–13386.

[3] Albert, R., Jeong, H., & Barabási, A.-L.. (2000) *Nature* **406**, 378–382.

[4] Albert, R. & Barabási, A.-L. (2002) *Reviews of Modern Physics*, **74**, 47–97.

[5] Newman, M.E.J. *SIAM Review* **45**, 167–256.

[6] Li, L., Alderson, D., Doyle, J.C., & Willinger,W. (2005) Towards a Theory of Scale-Free Graphs: Definition, Properties, and Implications. *Internet Math.*, to appear.

[7] Li, L., Alderson, D., Doyle, J.C., & Willinger,W. (2004) *Proc. ACM SIGCOMM 2004.*

[8] Carlson, J.M., & Doyle, J.C. (2002) *Proc. Nat. Acad. Sci. USA* **99**, Suppl. 1, 2539–2545.

[9] Tanaka, R. (2005) *Phys. Rev. Lett.* **94**, 168101.

[10] Moritz, M.A., Morais, M.E., Summerell, L.A., Carlson, J.M., and Doyle, J.C. (2005) Wildfires: An Example of Highly Optimized Tolerance? *Proc. Nat. Acad. Sci. USA*, to appear.

[11] Mandelbrot, B.B. (1997) *Fractals and Scaling in Finance*. Springer-Verlag, New York.

[12] Bollobás, B. & Riordan, O. (2003) *Internet Math.* **1**, 1–35.

[13] Dorogovtsev, S.N. & Mendes, J.F.F. (2003) *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press.

[14] Pastor-Satorras, R. & Vespignani, A. (2004) *Evolution and Structure of the Internet: A Statistical Physics Approach*. Cambridge University Press.

[15] Aiello, W., Chung, F., & Lu, L. (2000) *Proc. ACM STOC 2000.*

[16] Gkantsidis, C., Mihail, M., & Zegura, E. (2003) *Proc. SIAM Alenex'03.*

[17] Chung, F. & Lu, L. (2003) *Internet Math.*, **1**, 91–113.

[18] Newman, M.E.J. (2002) *Phys. Rev. Lett.* **89**, 208701.

[19] Erdos, P. & Renyi, A. (1959) *Publ. Math. (Debrecen)* **9** , 290-297.

[20] Newman, M.E.J. (2005) Power laws, Pareto distributions and Zipfs law. *Contemporary Physics*, in press.

[21] Detailed information about the objectives, organization, and development of the Abilene network are available from `http://www.internet2.edu/abilene`.

[22] Govindan, R. & Tangmunarunkit, H. (2000) *Proc. IEEE INFOCOM 2000*.

[23] Cooperative Association for Internet Data Analysis (CAIDA), Skitter. Available at `http://www.caida.org/tools/measurement/skitter/`.

[24] Spring, N., Mahajan, R., & Wetherall, D. (2002) *Proc. ACM SIGCOMM 2002*.

[25] Alderson, D. (2004) Tech Report CIT-CDS-04-004, California Institute of Technology.

[26] Willinger, W., Alderson, D., Doyle, J.C., & Li, L. (2004) *Proc. ACM SIGCOMM IMC 2004*.

[27] Zhang, Y., Roughan, M., Lund, C., & Donoho, D. (2003) *Proc. ACM SIGCOMM 2003*.

[28] Clark, D.D. (1988) *Proc. ACM SIGCOMM'88*, in: *ACM Comp. Comm. Rev.*, **18**(4): 106–114.

[29] Willinger, W. & Doyle, J.C. (2004) Robustness and the Internet: Design and Evolution. In *Robust design: A Repertoire of Biological, Ecological, and Engineering Case Studies*, E. Jen, Editor, Oxford University Press.

[30] Bank, D. & Richmond, R. (2005) ”Where the dangers are” *The Wall Street Journal*, July 18, 2005; page R1.

[31] Zhu, X., Yu, J. & Doyle, J.C. (2001) *Proc. IEEE Infocom 2001*.

[32] Jeong, H., Tombor, B., Albert, R., Oltavi, Z.N., & Barabasi, A.-L. (2000) *Nature* **407**, 651–654.

[33] Ravasz, E., Somera, A.L., Mongru, D.A., Oltvai, Z.N., & Barabasi, A.-L. (2002) *Science* **297** 1551–1555.

[34] Csete, M., & Doyle, J.(2004) *Trends in Biotechnology*, 22, 446-450.

Figure 1: Diversity among graphs having the same degree sequence D. **(a) RNDnet** A network consistent with construction by “preferential attachment”. The two networks represent the *same* graph, but the figure on the right is redrawn to emphasize the role that high-degree “hubs” play in overall network connectivity. **(b) SFnet** A graph having the “most preferential” connectivity again drawn both as an incremental growth type of network and in a form that emphasizes the importance of high-degree nodes. **(c) BADNet** A poorly designed network, with overall connectivity constructed from a chain of vertices. **(d) HOTnet** A graph constructed to be a simplified version of the Abilene network shown in Figure 2. **(e)** Power-law degree sequence D for networks (a)–(d). Only $d_i > 1$ shown.

Figure 2: **Router-level topology of Abilene.** Each vertex represents a router, and each link represents a physical connection, however each physical connection can support many virtual connections, giving the appearance of greater connectivity at higher layers of the Internet protocol stack. End user networks are shown in white, peering networks are shown in blue, and high-degree routers can only be found at the network periphery (not shown here).

Figure 3: **(a) Achieved router utilization:** HOTnet (circles) is close to the “efficient frontier”, SFnet (diamonds) operates significantly below this frontier, with the highly connected “hub” core router representing a glaring bottleneck. **(b) Achieved distribution of end user bandwidths:** HOTnet (circles) delivers a wide range of realistically different bandwidths to end users; SFnet (diamonds) delivers uniformly low bandwidth to all users. **(c) Apropos the “Achilles’ heel” of the Internet:** Robustness of HOTnet (SFnet) to deletion of worst-case nodes (deleting the worst 20 vertices corresponds to removing about 20% of the routers).

Figure 4: The diversity of graphs having the same degree sequence D. Despite having identical budget, technology constraint, degree sequence, and traffic demand model, when computing and plotting for the four network models in Figure 1 both their $S(g)$ (x-axis) and network performance $P(g)$ (y-axis) metrics, the four models occupy completely different areas in the $S(g)$ vs. $P(g)$ plane.

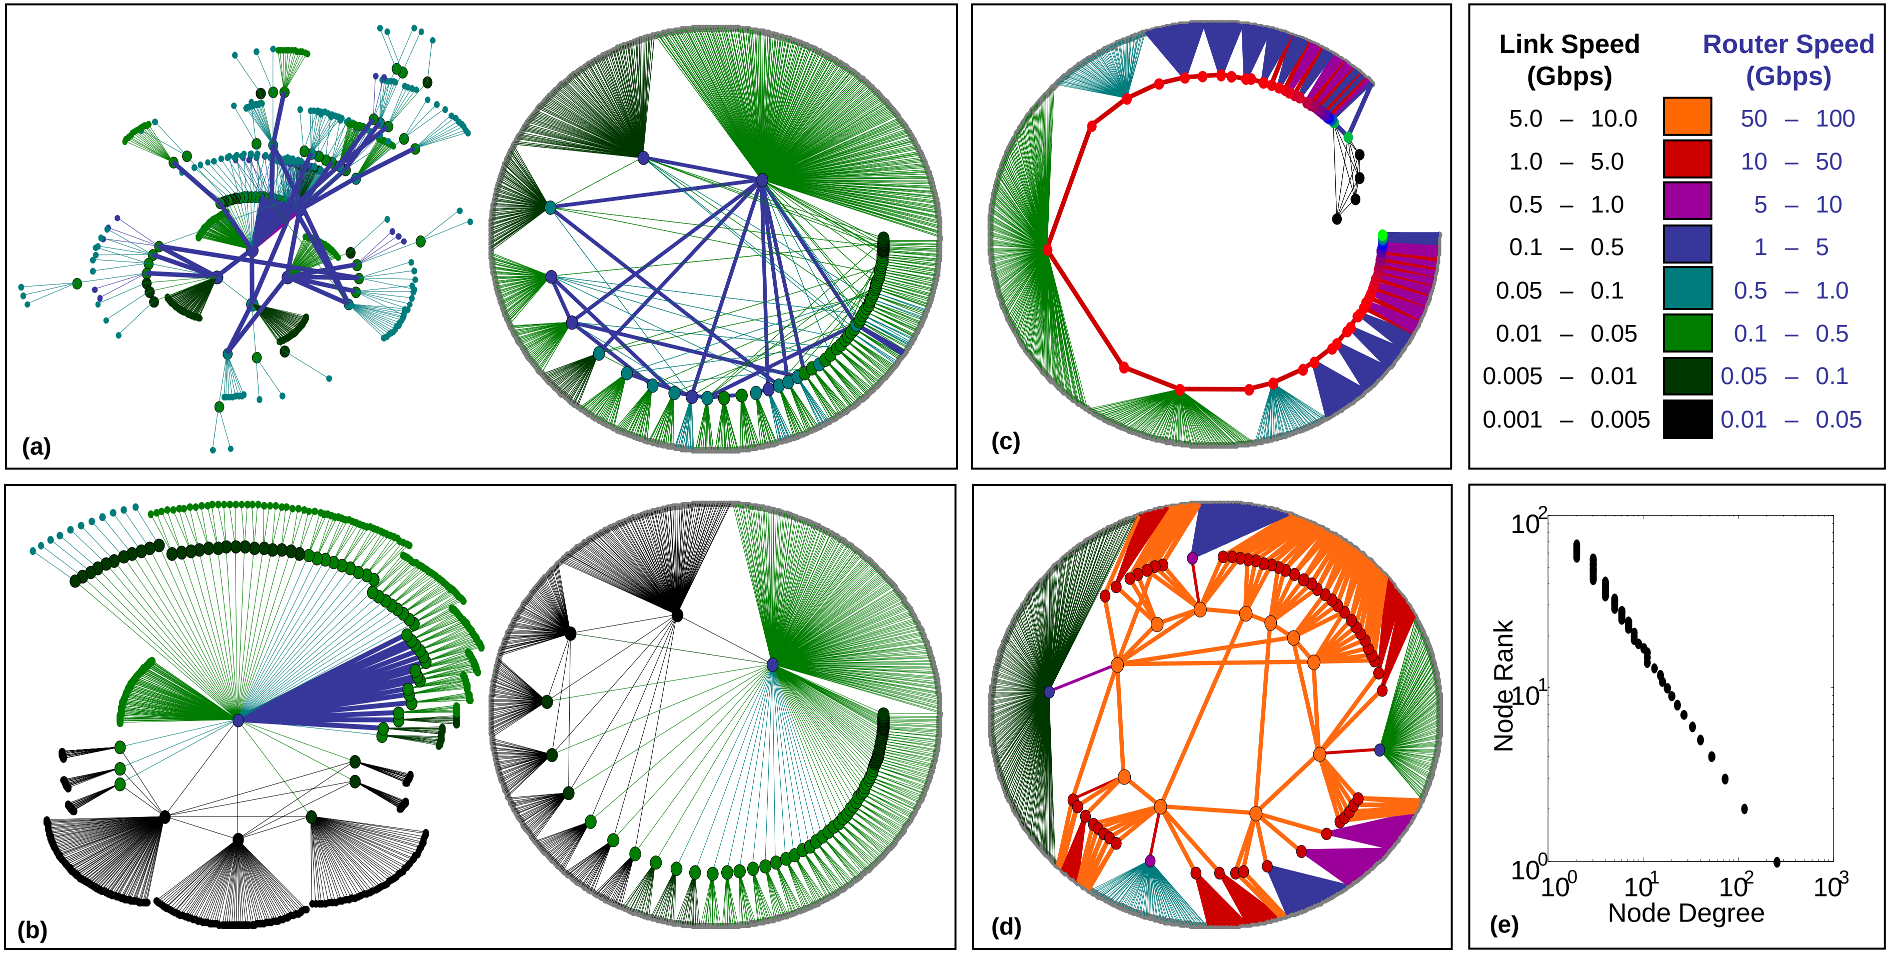
(a)
(b)
(c)
(d)
(e)
Link Speed (Gbps)
Router Speed (Gbps)
5.0 – 10.0 🟧 50 – 100
1.0 – 5.0 🟥 10 – 50
0.5 – 1.0 🟪 5 – 10
0.1 – 0.5 🟦 1 – 5
0.05 – 0.1 🟦 0.5 – 1.0
0.01 – 0.05 🟩 0.1 – 0.5
0.005 – 0.01 🟩 0.05 – 0.1
0.001 – 0.005 ⬛ 0.01 – 0.05
Node Rank
Node Degree
10^2
10^1
10^0
10^0
10^1
10^2
10^3

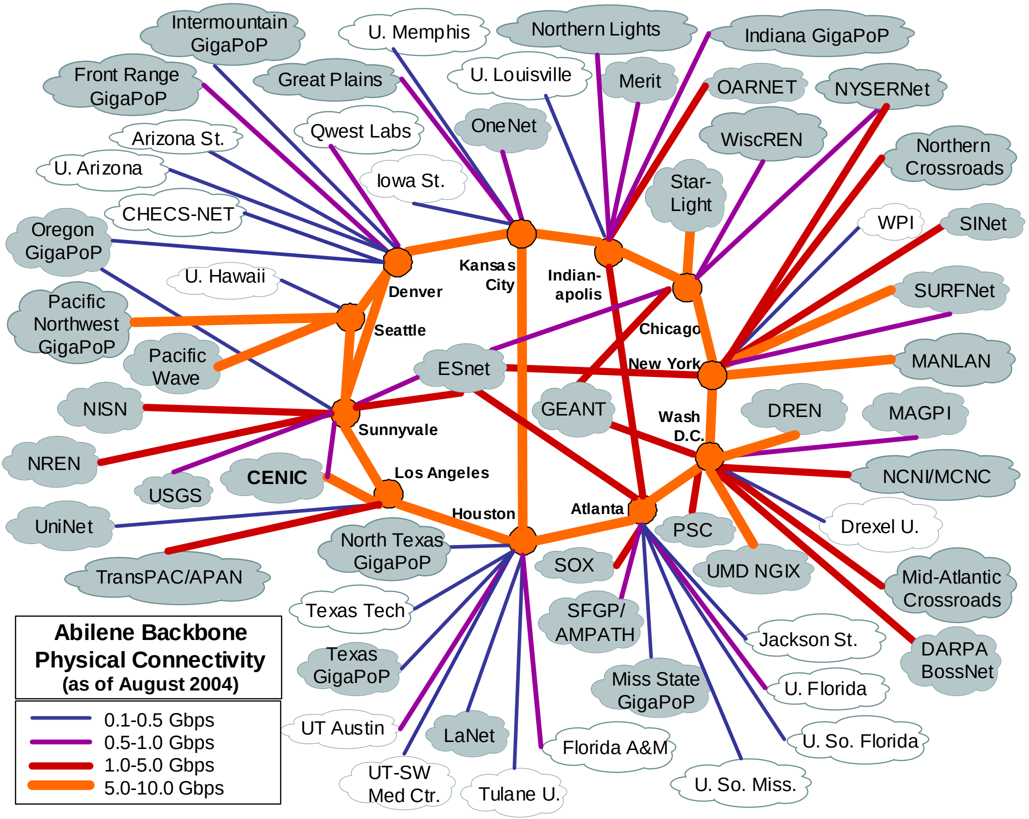

Abilene Backbone Physical Connectivity (as of August 2004)

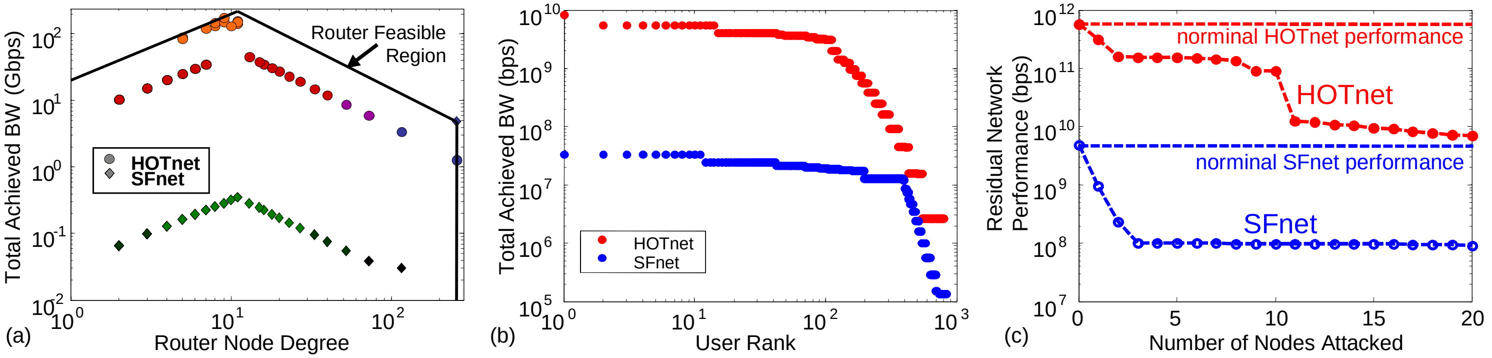
(a)
Total Achieved BW (Gbps)
10^{2}
10^{1}
10^{0}
10^{-1}
10^{-2}
10^{0}
10^{1}
10^{2}
Router Node Degree
Router Feasible Region
HOTnet
SFnet
(b)
Total Achieved BW (bps)
10^{10}
10^{9}
10^{8}
10^{7}
10^{6}
10^{5}
10^{0}
10^{1}
10^{2}
10^{3}
User Rank
HOTnet
SFnet
(c)
Residual Network Performance (bps)
10^{12}
10^{11}
10^{10}
10^{9}
10^{8}
10^{7}
0
5
10
15
20
Number of Nodes Attacked
norminal HOTnet performance
HOTnet
norminal SFnet performance
SFnet

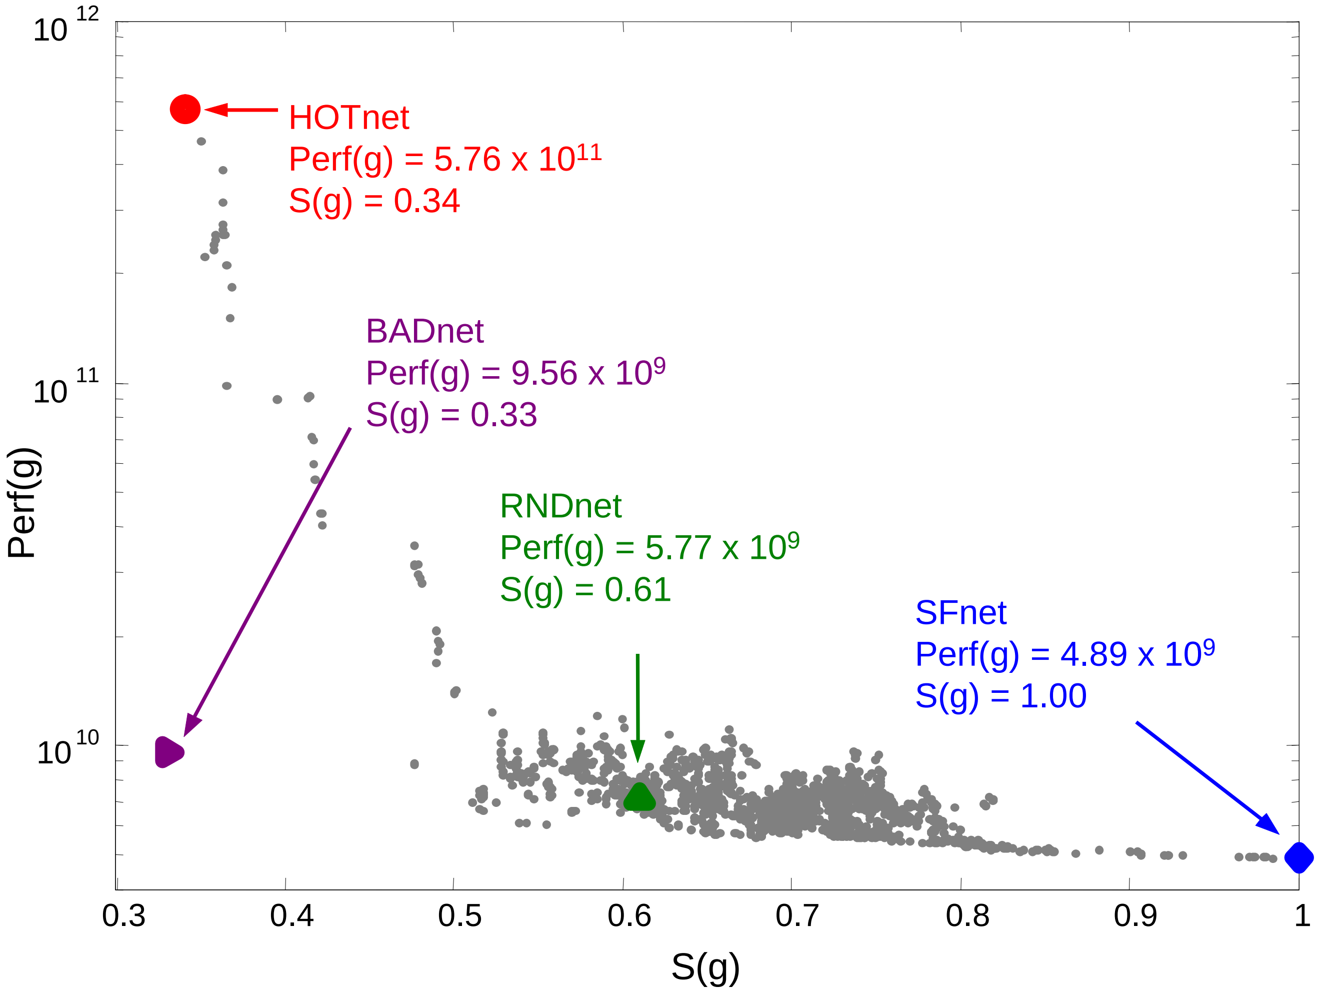
HOTnet
Perf(g) = 5.76 x 10^{11}
S(g) = 0.34
BADnet
Perf(g) = 9.56 x 10^{9}
S(g) = 0.33
RNDnet
Perf(g) = 5.77 x 10^{9}
S(g) = 0.61
SFnet
Perf(g) = 4.89 x 10^{9}
S(g) = 1.00
10^{12}
10^{11}
10^{10}
Perf(g)
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
S(g)