2005 IEEE International Conference on Systems, Man and Cybernetics
Waikoloa, Hawaii October 10-12, 2005

Robustness of Large Networks

Piet Van Mieghem
Delft University of Technology
P.O Box 5031, 2600 GA Delft, The Netherlands
P.VanMieghem@ewi.tudelft.nl

***Abstract*— The increasing importance of large networks in our "linked" world necessitates to enhance our understanding of their most characteristic behaviors. We discuss how to describe large networks and argue that a stochastic approach is most appropriate. The "robustness" of large networks is not uniquely defined. Here, the robustness is approached from two angles: we are studying the influence of the topology and of the link weight structure on the network's robustness.**

***Index Terms*— Link weight, graph, network, shortest path, robustness measures.**

I. Introduction

Our society depends more strongly than ever on large networks such as transportation networks, telephone networks, the Internet, and power grids. Many of these networks rely to a large extent on decentralization and self-organization. While decentralization removes obvious vulnerabilities related to single points of failure, it leads to a higher complexity. The Internet is an extreme example: there is no global control center, and obtaining complete information on its global state is an illusion. A more complex type of vulnerability appears in such systems. For instance, denial of service attacks, power failures and computer viruses are imminent threats to all computer networks.

In this article, we discuss how to define – in practical measurable terms – and to understand the robustness of complex networks. We present a stochastic approach to study large networks. Finally, by listing some open research questions, we outline the goal of the project "Understanding Complex Networks".

II. Description of a Network

A. The Topology of a Network

Any network consists of a set of N nodes and L links. The nodes are the individual items in the network connected to other items by links. For example, in transportation networks, telephone networks, power grids, social networks, biological molecules, the nodes are respectively crossings, switches, power generators or transformers, persons, atoms and the links are respectively roads, telephone wires, electrical wires, relationships, chemical bounds. Links can be directed (only from $i \to j$, but not from $j \to i$) or undirected (also called bi-directional).

The topological structure of a network is, in the terminology of graph theory, described by the adjacency matrix A. Each element a_{ij} of A is either zero or one: $a_{ij} = 1$ if there is a link between node i and node j, else $a_{ij} = 0$. Hence, the adjacency matrix expresses how the nodes in the network are interconnected, but it does not distinguish between the several types of links or nodes. In other words, in the adjacency matrix A each link in the network is equally important, whereas in reality, most of the common networks have heterogeneous links. On the other hand, an adjacency matrix specifies the importance of nodes in a network. The degree of a node, which is the number of direct neighbors in the network and equal to $d_i = \sum_{j=1}^{N} a_{ij}$, is a measure for the importance of a node: the higher the degree, the more links to other nodes in the network and, generally, the more important or central that node is.

B. The Link Weight Structure of a Network

1) Concept: In order to specify the heterogeneity of a network, link weights are assigned: a link weight w_{ij} is a real number (in some appropriate units) that quantifies a property of that link. For example, in a transportation network, a typical link weight is the physical distance between two points (nodes) in the network. In a communications network, link weights are also called Quality of Service (QoS) parameters. Examples of QoS parameters are the number of hops, the distance, the monetary cost, the delay, the jitter, the loss rate, the capacity (or bit rate) etc. Depending on the level of detail or the purpose of the network, a link can be specified by more than one link weight or, in short, by a vector of link weights. The components of the link weight vector such as distance, delay, etc. can depend on each other and are generally varying in time.

The importance of link weights is related to the purpose of the network and to an optimality criterion. Since most networks are expensive infrastructures, optimal use of resources is a common driver. For example, a telecommunications operator intends to serve as many traffic flows (e.g. telephone connections, IP packets, etc.) in his network as possible in order to optimize his revenues. This economic driver causes the network to transport flows along *shortest* paths where "shortest" depends on the optimization criterion. A criterion consists of a functional in terms of characteristic measures, the link weights. The functional is most often a linear combination of the involved link weights, although more complex forms are possible. Thus, the link weight structure $\mathcal{W}$ (the set of all link weights) affects the transport in a network if paths are determined in the network based on a shortest path criterion. In the sequel, we will assume that traffic flows from a certain source to a destination along a shortest path. More precisely, a path is a shortest path if it minimizes the sum of the link weights between that source and destination.

There are other types of criteria, which do not necessarily minimize an end-to-end sum of link weights. Perhaps the most important one, is the requirement that along the path the available capacity should be above a certain threshold. All links that do not satisfy the requirement are excluded and, from a network point of view, after removing these links, we obtain a reduced graph in which all links are eligible. Another example are policy rules such as the Internet's border gateway routing: every Internet Service Provider (ISP) announces his preferred paths based on his service level agreements with other providers. The end-to-end path is then a best match [8] between the several advertised set of preferred paths of individual Internet providers, which in about 70% of the cases corresponds to a shortest hop path.

2) Importance: Link weights also play a crucial role as a lever between technology and business. Beside the optimization criterion desired by a network operator to optimally use the network's resources, a business plan for a service provider is likely more geared towards the end-to-end quality of the transported service. The service provider has a pricing scheme for the various types of flows he offers to clients. A typical example is a (future) ISP that offers high quality or interactive video applications beside basic Internet access such as webaccess and email. While the latter services are typically best effort services that are not time critical, the video service is both time *and* capacity critical: if a video stream is not received within some critical time (typical 200 ms), the end-user experiences an unacceptable quality. The objective of a service provider is to offer services subject to the user's end-to-end QoS requirements. For example, a user may be willing to pay more for a higher quality of video, which translates into a higher capacity and a firmer guarantee that the end-to-end delay will not exceed a critical time. In these cases, link weights such as delay and available capacity allow the service provider to select those paths in the network that satisfy the user's QoS requirements. Instead of searching for a shortest path, the service provider's problem becomes a multiple constraints routing problem [13]: finding a path P such that $w_k(P) = \sum_{i \to j \in P} w_k(i \to j) \leq QoS_k$ for all k relevant QoS measures. In addition to routing, the service provider also needs price-coupled scheduling, which is beyond this paper and for which we refer e.g. to [11].

Finally, from a traffic engineering perspective, an ISP may want to tune the weight of each link such that the resulting shortest paths between a particular set of in- and egresses coincides with the desirable routes in his network.

In summary, any network can be described as a graph consisting of N nodes and L links whose topology is specified by the adjacency matrix A. Link differentiation asks for a link weight structure $\mathcal{W}$. Transport in networks is, most often, dictated by a shortest path criterion and, as expected more increasingly in the future, by a multiple constraints path problem. Thus, apart from the topology of the graph, the link weight structure clearly plays an important role. In most of the published work since Watts and Strogatz [16], Albert *et al.* [1] and Strogatz [9] gave birth to the new field of the "physics of networks", merely the underlying structure (i.e. adjacency matrix of the graph) has been taken into account and studied in large graphs. Hardly any paper discusses the importance of a link weight structure.

III. Describing Large Infrastructures

Although the principles outlined in Section II seem straightforward, large infrastructures suffer from a series of artifacts. First, often in large infrastructures, both the topology and the link weight structure are not accurately known due to (a) the size of the network (e.g. the Internet, road infrastructure), (b) the decentralized operation or/and (c) the time varying behavior. In large networks, links and nodes may join or leave the network temporarily or definitely, but there is no central authority that stores all information needed to construct the adjacency matrix. In many cases, a centralized database may not be (a) desirable and (b) feasible (if the network changes too rapidly).

Whereas the topology information is, in principle, simple to understand, the determination of a global link weight structure is more complicated. A first problematic issue lies in the definition of the link weight vector: which components are (a) desirable, (b) unique to determine, (c) stable enough in time to flood to neighboring nodes. Second, even if the components of a link weight vector are standardized, we face maintenance issues such as (a) when do we need updates or refreshments of the link weights, (b) what is the scope or extent of the network to which link weights should be announced. These few complications already illustrate why, in most large networks, the determination of a link weight structure is a difficult problem. For example, we refer to road traffic where even today, many traffic-jams are a daily persisting phenomenon, in spite of the large common cost. If drivers would know the link weight structure (e.g. the available capacity per road), traffic-jams could be reduced significantly.

This uncertainty about the precise structure leads us to consider both the underlying graph and each of the link weights as random variables. This choice is natural in the sense that it has proved to be successful in statistical physics in the study of large ensembles of particles [7, Chapt. 40]. While reasoning and computing with random variables is generally harder than with ordinary deterministic numbers, a probabilistic approach has several advantages. First, instead of determining the precise structure of the topology and link weight, we content ourselves to studying particular classes of random graphs and certain types of link weight distribution. In large networks, the link weights are hardly correlated and can be considered as independent to a good approximation. For example, the simplest class of random graphs are the Erdös-Rényi random graphs (see [4]) where $a_{ij} = 1$ with probability p for any pair (i, j), independent of the other links. The Erdös-Rényi random graphs are extremely well suited for computer simulations. The simplest distributions for the link weights are the exponential or uniform distribution.

The second advantage of probability theory lies in the power of the law of large numbers, or more generally, in limit laws. When the number of nodes in the graph is large, many of the properties in a graph such as the distance between two nodes, the weight of a shortest path between two arbitrary

chosen nodes etc. can be approximated by asymptotic laws when $N \to \infty$. In many problems, asymptotic behavior allows analytic modeling, which results in formulas that seem reasonably close to reality. A nice example of such a stochastic modeling is the shortest path computation in the Erdös-Rényi random graph with exponential links where an almost complete mathematical modeling is possible [12, Part III]. Another fascinating result of asymptotic laws in a network is that many properties of the graph become deterministic if $N \to \infty$. For example, although an arbitrary member of the class of the Erdös-Rényi random graphs exhibits a very random structure, the number of neighbors of an arbitrary node is almost equal to the average (which is a deterministic number). In some sense, these observations in graphs illustrate that the small local randomness hardly effects the macroscopic global behavior of the network. Other examples are given by nature itself where biological molecules such as proteins are so complex and seemingly all different, yet, their macroscopic function is almost deterministic.

In summary, probability theory with graph theory seems an adequate analysis tool for large networks to estimate the network's main properties. A deterministic approach is often not feasible due to the large size and the insufficiently known structure of large networks.

IV. Robustness of Large Networks

A. Very Brief Overview of the Literature

In recent years the study of "real-life networks" has attracted a lot of attention in the theoretical physics research community (see e.g. [2], [6] and [15]). The availability of large databases has disclosed the structure of complex evolving networks in such diverse fields as biology, sociology, engineering and computer science. What was for a long time purely the domain of mathematical graph theory has broadened to include also empirical sciences. The interest in the "physics of large networks" is now spreading to a much larger scientific community. About 5 years ago, the insight that many very different complex networks (e.g. the Internet, the world wide web graph, proteins, social relations networks, etc.) evolve and behave according to more general common rules has provoked a due surprise. For example, some remarkable characteristics, such as a power-law connectivity distribution ("scale-free networks"), and a small diameter in spite of a highly local structure ("small-world networks") are present in many of these complex networks.

The most important related mathematical and physics theoretic work is discussed in a two page section in [10, Sec. 1.4]. At a more general level, in data communications conferences (general such as INFOCOM[1] as well as measurement specific such as PAM[2]), in networking journals on both Ad-hoc networks and peer-to-peer networks, the number of papers on network robustness and security is increasing, quite likely pushed by the current social feeling of unsafety due to terrorism and, consequently, the larger amount of available funding. While some articles motivate their work by robustness arguments, we believe that, currently, there is no substantial body with scientific methods that address network robustness appropriately. Perhaps the only large programs, which gain a reviving interest are the RAND programs, which were set up in the 1960s (main time of the cold war) with the purpose to create a robust network against Russian attacks (see e.g. Baran [3]).

B. Robustness of the Topology

Intuitively, most people have some understanding of what robustness may mean. Given an arbitrary network (with specific topology and link weight structures), the basic question is: "What is the *robustness* of that network?" A first and natural way is to resort to graph theory in an attempt to define precisely the robustness of a network. There exist many metrics that characterize a graph [12, Chapt. 15] such as e.g. the degree distribution, the hopcount distribution of an arbitrary path, the diameter and the complexity of a graph, the clustering coefficient, the expansion, the resilience, the distortion, the betweenness and some more. Since robustness is closely related to disconnectivity and to the number of disjoint paths between two arbitrary points, we may consider the edge and vertex connectivity as relevant metrics. The edge connectivity $\lambda(G)$ of a connected graph G is the smallest number of edges (links) whose removal disconnects G. The vertex connectivity $\kappa(G)$ of a connected graph (different[3] from the complete graph K_N) is the smallest number of vertices (nodes) whose removal disconnects G. Unfortunately, both $\lambda(G)$ and $\kappa(G)$ are not so easy to compute.

So far, most of the above metrics of the graph are defined in the "topology domain" of the graph. The dual domain is the graph's spectrum. The spectrum of a graph is the eigenstructure of the adjacency matrix A and the related Laplacian $Q = \Delta - A$, where $\Delta = diag(d_1, \ldots, d_N)$ and d_j is the degree of node j. It can be shown [12, Theorem B.3.1] that the multiplicity of the smallest eigenvalue $\lambda = 0$ of the Laplacian Q is equal to the number of components in the graph G. Further, the second smallest eigenvalue μ_Q of Q has many interesting properties that characterize how strongly a graph G is connected. It is interesting to mention the inequality [5, pp. 265]

$$\kappa(G) \geq \mu_Q \geq 2\lambda(G)\left(1 - \cos\frac{\pi}{N}\right) \qquad (1)$$

This theory suggests us to consider both $\lambda(G)$ and $\kappa(G)$ as well as μ_Q as measures for the robustness of a graph. We are currently studying the distribution of the second smallest eigenvalue μ_Q of the Laplacian in the class of Erdös-Rényi random graphs. We consider the understanding of properties in the class of Erdös-Rényi random graphs as a first step to understand more complex and more realistic classes of graphs. Most often holds that, if we cannot solve the problem for Erdös-Rényi random graphs, the chance is very small that we can do it for other classes (except for regular classes such as lattices or k-ary trees). Secondly, if the three proposed

[1] http://www.ieee-infocom.org
[2] http://www.pam2005.org

[3] The complete graph K_N cannot be disconnected by removing nodes and we define $\kappa(K_N) = N - 1$ for $N \geq 3$.

measures $\lambda(G)$, $\kappa(G)$ and μ_Q seem appropriate and effective to understand robustness in a graph, efficient methods to compute these measures need to be derived. A logical next step is to propose rules to increase the degree of robustness in a given infrastructure.

V. Link Weights and Shortest Paths

Here, the influence of the link weights on the shortest path tree (SPT) is discussed and an intuitive relation to the robustness is indicated.

A. Theory

The SPT rooted at some node consists of the union of shortest paths from that node to any other node in the network. Since the SPT problem is mainly sensitive to the smaller, non-negative link weights, the probability distribution of the link weights around zero will dominantly influence the properties of the resulting SPT. A *regular* link weight distribution $F_w(x) = \Pr[w \leq x]$ has a Taylor series expansion around $x = 0$,

$$F_w(x) = f_w(0)\, x + O\left(x^2\right)$$

since $F_w(0) = 0$ and $F'_w(0) = f_w(0)$ exists. A regular link weight distribution is thus linear around zero. The factor $f_w(0)$ only scales all link weights, but it does not influence the shortest path. The simplest distribution of the link weight w with a distinct different behavior for small values is the polynomial distribution,

$$F_w(x) = x^\alpha 1_{x\in[0,1]} + 1_{x\in[1,\infty)}, \quad \alpha > 0, \tag{2}$$

where the indicator function 1_x is one if x is true else it is zero. The corresponding density is $f_w(x) = \alpha x^{\alpha-1}$, $0 < x < 1$. Figure 1 illustrates the three major $\alpha-$regimes. The exponent α is called the *extreme value index* of the probability distribution of w and $\alpha = 1$ for regular distributions. By varying the exponent α over all non-negative real values, any extreme value index can be attained and a large class of corresponding shortest path trees, in short α-trees, can be generated.

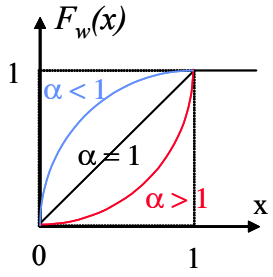

Fig. 1. The three different α-regimes are shown: **(a)** $\alpha > 1$ for which $f_w(0) = 0$, the influence of the link weight structure on the SPT is less important than the underlying topology. In fact, if $\alpha \to \infty$, all link weights are 1 almost surely. **(b)** $\alpha = 1$ corresponds to a regular link weight distribution. **(c)** $\alpha < 1$ for which $f_w(0) \to \infty$, the link weight structure contains a larger amount of very small link weights than the regular distribution and it significantly dominates the shortest path.

We have shown in [14] that the class of polynomial link weights possesses interesting properties, which we summarize here. A polynomial link weight structure tunable by one parameter α on e.g. the complete graph, a broad class of shortest path trees are found as function of α. Instead of worrying about the precise topology of the graph, a tunable link weight structure thins out the complete graph to the extent that a specific shortest path tree can be constructed. The approach thus eliminates the precise knowledge of the underlying graph by immediately concentrating on the tree properties induced by polynomial link weights.

Secondly, relatively small variations in the link weight structures cause large differences in the properties of the SPT. In particular, the average hopcount in a graph with N nodes follows a different scaling: $\mathbb{E}[H_N] = O(\ln N)$ for α around 1 while $\mathbb{E}[H_N] = O\left(N^{1/3}\right)$ if $\alpha \to 0$. The logarithmic $O(\ln N)$-scaling corresponds to "small world" networks that are densely interconnected such that typical paths only possess a few hops. A well-known example of a small world network [15] is the graph whose nodes are persons and whose links are generated by the acquaintance relations between persons. The algebraic $O\left(N^{a/b}\right)$-scaling corresponds to sparse networks where paths contain generally many hops. For example, the hopcount between two random points in a 2-dimensional lattice with N nodes scales as $O\left(N^{1/2}\right)$.

Third, for $\alpha \to 0$, all SPTs coincide with the minimum spanning tree (MST). Hence, for $\alpha \to 0$, all traffic in the graph routed along SPTs traverses precisely the same $N-1$ links that form the "critical backbone". The $\alpha \to 0$ regime corresponds to a strong disorder regime since $\frac{\sqrt{var[w]}}{\mathbb{E}[w]} \sim \frac{1}{\sqrt{\alpha}} \to \infty$. In the other regime ($\alpha$ around 1), the SPT has the structure of a uniform recursive tree (URT) whose properties are described in [12, Chapt. 16].

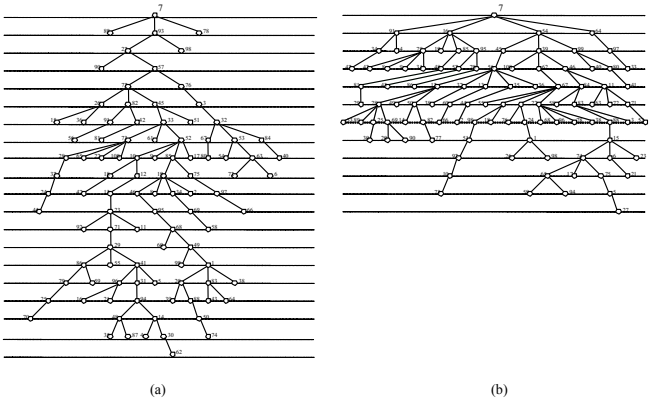

Fig. 2. An example in the graph with $N = 100$ nodes of (a) the MST which is the SPT for $\alpha = 0$ and (b) the URT which is the SPT for $\alpha = 1$. Both trees are structured per level sets where each level shows the number of nodes at different hopcount from the root (here node with label 7).

Figure 2 visualizes the different structure of a typical MST (a) and a typical URT (b) of the same size $N = 100$.

Fourth, between both regimes, we found a fascinating phase transition around a critical α_c as illustrated in Figure 3. Figure 3 shows the probability that the union of all shortest paths

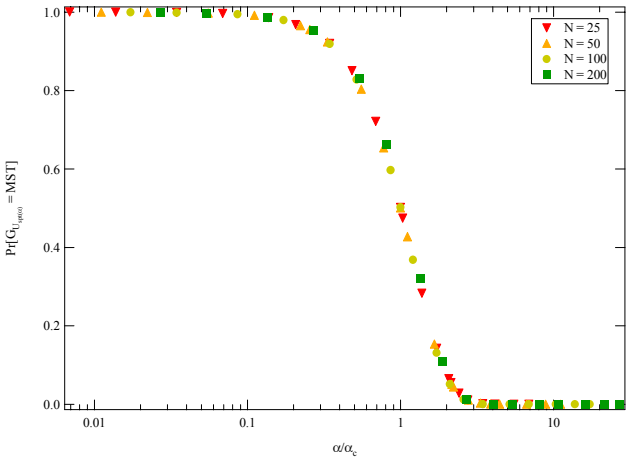

Fig. 3. The probability that the union of all shortest paths $G_{\cup spt(\alpha)}$ in the complete graph with polynomial link weights is a minimum spanning tree as a function of the normalized α/α_c.

$G_{\cup spt(\alpha)}$ between all node pairs in the complete graph with polynomial link weights is a minimum spanning tree. In real networks where almost all flows follow shortest paths through the network, the union of all shortest paths is the observable part of a network. For example, the union of all trace-routes (Internet paths) between all node pairs in the Internet, would represent the observable graph of the Internet. The real Internet is larger because it also contains dark links for back-up paths needed in case of failures. We observe in Figure 3 a phase transition around $\frac{\alpha}{\alpha_c} = 1$, where α_c is defined as $\Pr\left[G_{\cup spt(\alpha_c)} = MST\right] = \frac{1}{2}$. For $\alpha < \alpha_c$, most graphs $G_{\cup spt(\alpha)}$ are trees with high probability while for $\alpha > \alpha_c$ hardly any graph $G_{\cup spt(\alpha)}$ is a tree. Moreover, Figure 3 shows that the phase transition obeys a same function in the normalized $\frac{\alpha}{\alpha_c}$ for any size of the network N. Finally, initial simulations indicate that both the critical extreme value index α_c and the width $\Delta\alpha$ of the phase transition[4] obey a curious $O\left(N^{-\beta}\right)$-scaling law with $\beta \approx \frac{2}{3}$.

B. Implications

Although eccentric at first glance, a number of interesting conclusions can be drawn. From a topological view, the trees in large networks such as Internet indeed seem to consist of a critical bearer tree (corresponding to the $\alpha \to 0$-tree) overgrown with URT-like small trees (influence of $\alpha = 1$). The latter cause that the hopcount in the Internet still scales logarithmically in N, rather than algebraically as for the $\alpha \to 0$-tree (MST). This effect is similar to the small world graphs: by adding a few links in a 'large average hopcount graph', the hopcount may decrease dramatically [15].

Our theoretical study inspires some new research. If trees in large networks can be modelled via α-trees, insight about an *effective* link weight structure may be gained. As mentioned before in Section III, it is generally difficult to determine a realistic link weight structures because that effective link weight structures arises as a combination of shortest path routing and possibly not-shortest path routing such as e.g. policy routing. In a next stage, the modeling of multicast trees in terms of α-trees may be interesting since little about Internet multicast trees is known, while future applications such as pay-tv would considerably benefit from multicast over the current unicast.

From a control point of view in networks, operators may wish to steer flows by tuning link weights. Our study indicates that large variations in the link weights (α small) will result in overall properties close to the $\alpha \to 0$-tree (MST): many flows will traverse over a same set of links and the overall hopcount will increase. From a robustness point of view, choosing α around 1 will lead to the use of more paths and, hence, a more balanced overall network load. In other words, possible failures of a small set of nodes or links are unlikely to effect the global transport in the network. In the $\alpha \to 0$ regime, all flows are transported over the minimum possible fraction of links in the network. Any failures in a node or link disconnects the MST into two parts and may result in the obstruction of transport in the network. In summary, from the view point of robustness, the $\alpha \to 0$ regime may constitute a weak regime although it is highly efficient: only $N-1$ links are used, which means that a minimum of links need to be controlled and/or secured.

While phase transitions are natural phenomena (e.g. the freezing of water into ice below zero degree Celsius and liquid above), our finding illustrates that phase transitions can be constructed in large infrastructures where link weights can be controlled independently from the underlying topology. In other words, if the link weight structure can be considered as a property of the network orthogonal to the graph's topology, we may switch the traffic in the network between two extreme transport profiles. Since also the width of the phase transition is narrow, from a control point of view, a network operator may choose an independent link weight structure with extremal value index near to α_c: by changing the link weight's extreme value index with $\Delta\alpha$, he can switch traffic over two entirely different patterns. The link weight phase transition is similar to electrical conductivity in superconductive solids. Above a critical temperature T_c, the normal conduction consists of the ensemble of electrons that travel over different paths while below T_c, a superconducting state is formed in which electrical current flows as a kind of super-wave through the solid with non-measurable resistance. The analogy with nature shows that above α_c, a collection of small and seemingly unordered and local flows traverse the network, while below α_c, transport is transformed into a large and global network phenomenon, comparable with a macroscopic quantum effect (such as laser-light and superconductivity).

VI. Agenda for Research

Whereas some understanding of the structure of large networks has already been gained, many open questions remain, in particular for the project "Understanding Complex Networks":

[4]The width of the phase transition is defined by $2\Delta\alpha = \alpha_h - \alpha_l$ where $\Pr[G_{U_{spt}}(\alpha_h) = MST] = 0.05$ and $\Pr[G_{U_{spt}}(\alpha_l) = MST] = 0.95$.

1) What are the measurable or observable quantities that specify a large and continuously evolving network?
2) Can we define robustness of a network such that it can be measured easily and rapidly?
3) How does a complex network degrade under the influence of failures or virus attacks? The spreading of viruses follows largely the usual communication routes in the network, rather than attacking random or selected nodes. What does this imply for the breakdown of a network?
4) What is the influence of the link weight structure (that determines the shortest paths and the main traffic flows in the networks) on the throughput and end-to-end behavior (QoS) of individual flows? How can we by tuning the link weight structure[5] enhance the robustness of a network? What is the effect of heterogenous link weights[6] (e.g. a network may consist of a proportion of red/reliable and blue/unreliable links).
5) Regarding the sensitivity of the link weight structure, what is the maximum amount Δw in link weight change in order not to modify the set of shortest paths in a network. This insight is important to estimate the topology update overhead and the major time scale for updates in networks (e.g. in road traffic or the Internet where the link weights may be associated with the traffic traversing the link). An accurate view of the updated topology is crucial in any network that wishes to provide some end-to-end quality to carried traffic flows. The possible insensitivity of link weight changes (stability of a network) may be regarded as a robustness property.
6) What are the possible threats to a network, their effect on the network behavior and how can we classify them? These threats can range from denial of service attacks at a terroristic scale, to simple overloading of the network at peak times. These threats will be modeled, and studied in combination with a range of underlying topologies such as random, scale-free, and small-world graphs and of link weight structures (tunable by the extreme value index α). We intend to compare these models with measurement data on real systems.

VII. Conclusions

The scientific challenge is to understand the vulnerabilities of complex, self-organized networks, which are exposed to a dynamic and potentially hostile environment. Technically, the challenge is to adapt traditional methods from mathematical graph-theory and statistical physics to study complex networks and, in particular, to model their behavior under (over)loading of the system with benign or malicious traffic.

[5]Specifically, for example, what are the properties of a link weight structure to have between two arbitrary points in a graph with high probability two different shortest paths? Such a link weight structure is likely to exist because in a graph with unit link weights, all equal hop paths are equal weight paths. A solution that includes link differentiation (i.e. not all link weights equal to 1) and provides two possible shortest paths is more robust and favours load balancing in networks.

[6]How many blue (i.e. weak links) are in a shortest path? For example, wired links can be considered as red (high capacity, especially optical links) while wireless links as blue (low capacity especially mobile ad-hoc network links).

The innovative nature of the project "Robustness of Complex Networks" can be best understood by considering the vulnerability of the Internet against denial of service attacks and computer viruses. The usual way to combat these threats is at the level of operation systems: by creating anti-viruses, applying firewalls, securing "holes" in operating system etc. Rather than taking this reductionistic view focusing on single network nodes, our approach is more "holistic": we intend to consider the vulnerability of the system as a whole.

The project leads to a better understanding of network vulnerabilities, and to design rules for improving robustness. For instance the "holistic" approach could lead to immunization strategies for the Internet, which complements rather than replaces the production of curative software such as anti-viruses. Apart from some generic work on virus-spreading on scale-free and small-world networks, not much work on the type described above has been done.

Finally, from an end-to-end point of view (the user's network experience as well as the network operator's control), the study of the influence of link weights is new.

Acknowledgments.

We would like to thank Stijn van Langen for initiating the work of Section V-A, Serena Magdalena for providing the Figures 2 and 3 and Almerima Jamakovic for the constructive input.

References

[1] R. Albert, H. Jeong, and A.-L. Barabási. Error and attack tolerance of complex networks. *Nature*, 406:378–382, 27 July 2000.
[2] A.-L. Barabasi. *Linked, The new science of networks*. Perseus, Cambridge, MA, April 2002.
[3] P. Baran. The beginnings of packet switching - some underlying concepts: The Franklin Institute and Drexel University seminar on the evolution of packet switching and the Internet. *IEEE Communications Magazine*, pages 2–8, July 2002.
[4] B. Bollobas. *Random Graphs*. Cambridge University Press, Cambridge, UK, 2nd edition, 2001.
[5] D. M. Cvetkovic, M. Doob, and H. Sachs. *Spectra of Graphs, Theory and Applications*. Johann Ambrosius Barth Verlag, Heidelberg, third edition, 1995.
[6] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks, From Biological Nets to the Internet and WWW*. Oxford University Press, Oxford, 2003.
[7] R. P. Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics*, volume 1. Addison-Wesley, Massachusetts, 1963.
[8] T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, April 2002.
[9] S. H. Strogatz. Exploring complex networks. *Nature*, 410(8):268–276, March 2001.
[10] R. van der Hofstad, G. Hooghiemstra, and P. Van Mieghem. Distances in random graphs with finite variance degree. *Random Structures and Algorithms*, to appear 2005.
[11] J. A. Van Mieghem and P. Van Mieghem. Price-coupled scheduling for differentiated services: Gcμ versus GPS. *International Journal in Communications*, 15:429–452, 2002.
[12] P. Van Mieghem. *Performance Analysis of Communications Systems and Networks*. Cambridge University Press, 2005.
[13] P. Van Mieghem and F. A. Kuipers. Concepts of exact quality of service algorithms. *IEEE/ACM Transaction on Networking*, 12(5):851–864, October 2004.
[14] P. Van Mieghem and S. van Langen. Influence of the link weight structure on the shortest path. *Physical Review E*, 71, to appear 2005.
[15] D. J. Watts. *Small Worlds, The Dynamics of Networks between Order and Randomness*. Princeton University Press, Princeton, New Jersey, 1999.
[16] D. J. Watts and S. H. Strogatz. Collective dynamics of "small-worlds" networks. *Nature*, 393:440–442, June 1998.