Hawaii International Conference on Systems Sciences, January 2002, Hawaii.

Growth and Propagation of Disturbances in a Communication Network Model

D. E. Newman
Physics Dept., University of Alaska
Fairbanks, AK 99775 USA
ffden@uaf.edu

Nathaniel D. Sizemore
Oak Ridge National Laboratory
Oak Ridge, TN 37831

B. A. Carreras
Oak Ridge National Laboratory
Oak Ridge, TN 37831
carrerasba@ornl.gov

V. E. Lynch
Oak Ridge National Laboratory
Oak Ridge, TN 37831
lynchve@ornl.gov

Abstract

In a society in which information is one of the highest-valued commodities, information networks are the economic backbone. Therefore, network vulnerability is a major hazard. Analysis of communications system traffic suggests that there are a variety of similar dynamical characteristics in widely varying systems. There are some universal features of communication networks that may determine the main dynamics of communications. These are: 1) hardware limitations on network capacity, 2) forcing from the increasing number of users towards the network capacity, and 3) a high level of complexity in those systems. These common features may drive the system dynamics and determine its main intrinsic vulnerabilities. The nontrivial character of network dynamics has already been found in traffic analysis of local area networks that display self-similarities and long-range correlations. This suggests a complex dynamics in the traffic of packets over the network. In this work simple distribution network models with different levels of complexity are examined and compared to each other and to characteristics of real Internet data.

1. Introduction

Complex systems such as, communication networks, distribution systems, and electrical power transmission grids which run near their operational limits, can develop non-periodic major cascading disruptions which can have serious consequences. Large-scale disruptions of these systems have an obvious impact on national security and may point out a vulnerability in the Nation s infrastructure. Individually these disruptions can be attributed to a specific cause, or causes, such as hardware failure, unexpected traffic, lightning strikes or shorts through trees. However, finding these individual causes overlooks the global dynamics of the system in which repeated major disruptions, from a wide variety of sources, are a virtual certainty.

Increased system security might decrease the risk of an individual event, but does not eliminate the possibility of global disruptions of the system in question. The reason for this is that these types of disruptions are intrinsic to the dynamics of many complex systems. This type of global dynamical system behavior has been connected to the concept of Self-Organized Criticality (SOC)[1]. A SOC system is one in which the nonlinear dynamics in the presence of perturbations organizes the overall average system state near to, but not at, the state that is marginal to major disruptions. These systems are characterized by a spectrum of spatial and temporal scales of the disruption events that exists in remarkably similar form in a wide variety of different physical systems.

Analysis of Ethernet traffic measurements has already shown the self-similar character of the data and the existence of long time correlations that can not be explained by simple statistical models[2-5]. The algebraic tails of the probability distribution function of this data is suggestive of SOC type dynamics. Although sophisticated statistical models can be (and are being) used in interpreting the data, there is the need for development of dynamical models if we want to understand the temporal operation of the systems as well as having some ability to control the flows (information, power etc.) and to understand as well as avoid vulnerabilities in the networks. The SOC character of communication network has been put forward for traffic flow by Nagel and collaborators[6]. We utilize the universality of the SOC concept to investigate the dynamics of communication and distribution systems, with the emphasis on catastrophic failure, from a very general system point of view. The applications can be

directed to general distribution systems as well as to information networks and power transmission systems, discussed elsewhere, in particular.

If communication networks are in a SOC state then some direct results of understanding this could come in the following forms: First and most straight forward is a more realistic failure statistic which would allow for more realistic risk analysis. Next, understanding the general system dynamics could allow changes in design, making the system inherently less susceptible to major cascading disruptions. Finally, the possibility of semi-quantitative predictions of the onset of major system collapses, through real time monitoring, could allow preventative steps to be taken to mitigate or even prevent the disruption.

Simple analysis of communications system traffic dynamics, as well as that of most other distribution systems, suggests that though the details differ there are in each similar dynamical characteristics. In a relatively brief time, networks have developed from relatively small LANs to the Internet. While each type of network has its own distinctive characteristics and communication technology both of which are rapidly evolving in time, there are some universal features of network communication that may determine the main dynamics of communication. These are: 1) hardware limitations on the network capacity, 2) the forcing from the constantly increasing number of users towards the limiting capacity of the network, and 3) a high level of complexity in those systems. These common characteristic features may drive the dynamics of communication and determine its main intrinsic vulnerabilities.

The nontrivial character of communication network dynamics has already been shown[7,2] in the analysis of traffic in local area networks. Although the arrival of communication packets in a local network maybe thought to be random, the arrival process is neither Poisson nor compound Poisson. The analysis of local area network traffic has shown that it is self-similar in nature with long-range correlations. This suggests a complex dynamics in the traffic of packets over the network. This complex behavior has been characterized by the existence of a dynamical phase transition in the Internet traffic[3].

Simple models of the communication traffics have been build[7,8] that illustrate the existence of the phase transition. These models are characterized by a fixed rate of packet creation and do not have steady state traffic in the congested regime. They are inspired by traffic models that have such a phase transition between flowing and congested traffic[9]. A cellular automata model[6] has been used to investigate some aspects of the collective behavior of the computer network. This model is 1-D and is basically concerned with the queuing problem.

Here, we investigate the dynamics of communication models incorporating some possible self-organization dynamics that can allow the development of steady state traffic even in the congested regime. We base some of these mechanisms on existing congestion control methods being applied to communication networks.

2. Basic Communication Model

We have developed a communications model based on the model of Ohira and Sawatari[8]. The model is applicable to arbitrary network configurations, but here we consider only a two-dimensional lattice network configuration (Fig.˚1). These lattices are square and identified by N, the number of nodes on a side. Thus an N˚=˚20 lattice contains 20^2 nodes. Nodes on the edges of the lattice are denoted as hosts that could both create and receive packets. (Note that hosts are not connected to each other directly and that nodes at the corners of the lattice were not used at all, as they can be connected only to hosts.) Interior nodes are denoted as routers, responsible for moving packets in the direction of their destination. Each router can move one packet per time step. If the router receives more packets than can be processed, they are buffered by the router into a queue and processed in the order received. This buffer can have a fixed size; if that limit is exceeded, the router becomes incapable of receiving packets until the queue size is again below the limit.

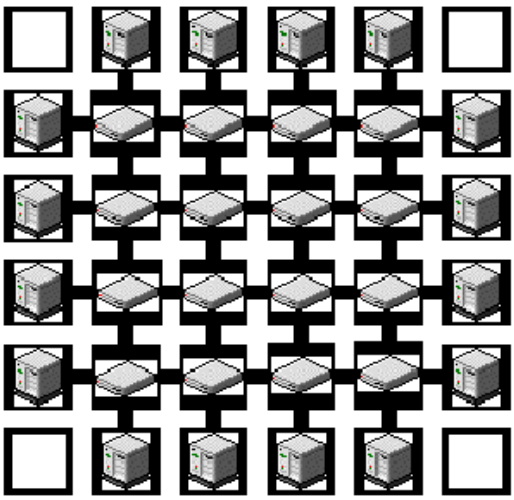

Fig.˚1.˚˚An N˚=˚6 lattice network. Nodes along the edges of the lattice are designated as hosts; interior nodes are routers. Nodes at the corners can be connected only to hosts, and therefore are not used.

A shortest path algorithm determines distance between nodes in the lattice. This is implemented as a breadth-first search through the network where each connection between nodes is assumed to have a weight of one. These distances are calculated at the beginning of the calculation and accessed using a lookup table thereafter.

The model is based on the following rules:

1. Hosts create packets with a probability λ, and insert them into the system. Generated packets are given a random destination chosen from the hosts in the lattice.
2. Routers forward the top packet in the queue to the neighboring node closest to the destination of the packet. If the closest distance is shared by more than one node, selection is made randomly between them. If a neighbor s queue is full, the router chooses the next closest neighboring node.
3. Hosts that receive packets addressed to them delete the packets —this removes the packets from the system.

Without some method of congestion control, the network undergoes a phase transition at a value of λ, denoted by λ_{crit}. When this critical value is exceeded, the network becomes locked in a traffic jam where routers are receiving packets more quickly than they can process them. This is illustrated in figure 2 where we have used a square network with 400 nodes and carried out the calculations for 10^5 steps. A transition is observed for a value of the probability of creating packages $\lambda = 0.12$. When the probability of package creation is increased above this value, we observe a sharp increase in the average time for delivering a package, while the averaged distance travel per package remains the same.

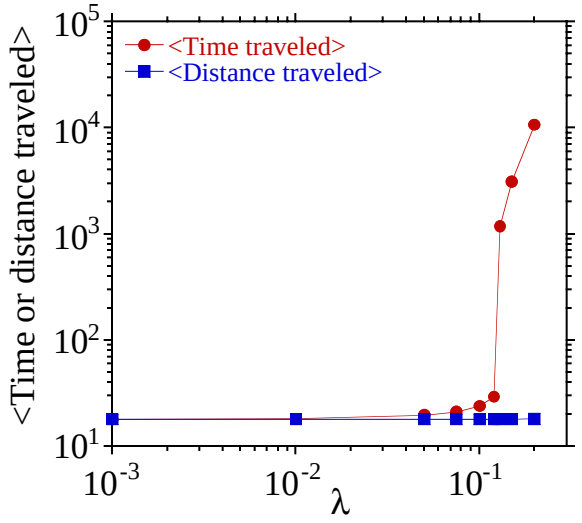

Fig. 2. Averaged time and distance traveled by each packet as a function of the packet creation rate

After the transition, the averaged traveled time simply depends on how long the calculation runs. At this point, there is a global traffic jam and the system is not in steady state. Packets accumulate at all routers and the delay in delivering the packages increases with time. There are additional ways of detecting the transition, one being through calculating the averaged number of packages waiting in line to be transmitted by the routers. There is a continuous increase in the list length with an increase of the probability λ. However, at the transition point, there is a clear jump. The size of the jump depends on the time that the code has been running for the reasons discussed above.

3. Communication models with self-organization

In order to deal with these traffic jams and maintain the circulation of information in the system, it is necessary to implement methods of congestion control. These methods can induce a self-organization of the communication network and can help in understanding the operation dynamics of the real system. Four methods of congestion control were extracted from actual networking techniques and implemented in the simulation[11]. We assume that each router node has a buffer of limited capacity able to hold a fixed number of packets to be transferred. The congestion control methods that we have considered are:

- *Dropping packets* — The simplest technique for dealing with congestion is to delete those packets that arrive at a node with a full buffer.
- *Congested signaling* — This scheme involves introducing a time delay, T_d, to deal with congestion. If a buffer overrun occurs, the host that generated the packet that caused the overrun is put into a suspended state where it will not generate packets for a set number of time steps.
- *Simplified choke packet* — This technique also introduces a time delay; however, instead of suspending the generating host, the node that delivered the packet is suspended.
- *Backpressure* — Analogous to a backpressure-plumbing valve, this congestion method deals with congestion by changing the value of λ at the generating host by an amount Δλ.

In real-world networks, communicating commands to delay hosts, change the rate of packet generation, etc. would have to be done with packets as well, thus adding more traffic to the network. This extra traffic was not deemed statistically significant and for simplicity, the model does not include them.

Several diagnostics are used to measure the performance of the calculation and to compare network congestion schemes. They include: throughput, given by the number of packets delivered divided by the total time in the simulation; the mean, variance, and probability distribution function of the time traveled by the packets; and the mean, variance, and probability distribution function of the distance traveled. Throughput as defined and measured in this study is analogous to throughput in a real-world network —that is, how much information can be moved in a given

time. The effective packet creation rate, λ_{eff}, of the system is calculated as the throughput divided by the number of hosts and used for comparison with the value for λ, the desired packet creation rate. The average time traveled by the packets is analogous to real-world network speed — how quickly a given piece of information in averaged can be moved from one point to another.

Once the congestion control methods are in operation, the evolution of the system reaches a steady state independently of the creation rated desired. We can see this by just looking at the number of packets moved through the system every time step (Fig.˚3).

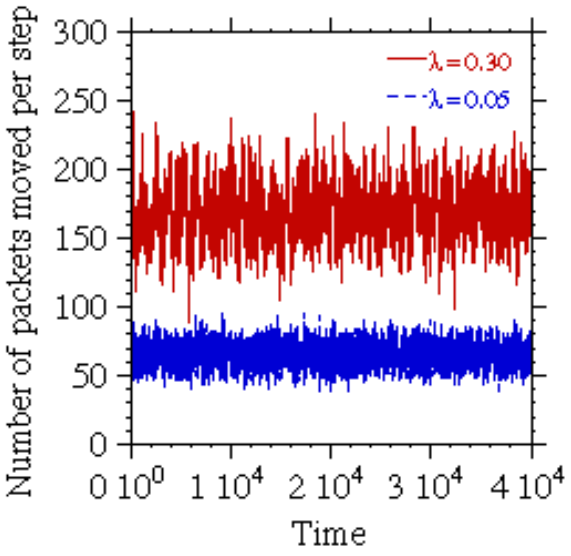

Fig.˚3.˚˚Time traces of the number of packets moved per step in a square network with 400 nodes and for two different values of the packet creation rate.

4. Dynamics of the packet traffic in a communications network model

Each of the congestion control techniques was tested with a network of N˚=˚20 for values of λ ranging from 0.02 to 0.8. The same random number seed was used in each case, and the number of time steps for the simulation was held constant. The buffer size for the routers for all cases was arbitrarily set to 60. A delay of 30 time steps was used for the congested signaling and simplified choke packet methods. This combination of values was found to be effective in giving congestion control for most of the methods and most values of l. A Δλ of 10% was chosen arbitrarily for the backpressure method. We compared the various control methods and examined the average time traveled for packets in the system as well as throughput. We found it useful to introduce a measure of the efficiency of the system as the averaged number of packets delivered per averaged packet transit time. This combines the previous measures in a single parameter.

A transition was again observed in the time traveled for each case as the network went from a non-congested state to one in which traffic jams are occurring (Fig.˚4). However, in the cases with the congestion control, the size of the change between presence and absence of traffic jams was both much lower and actually bounded relative to the calculations that did not include congestion control methods. Again, the efficiency is found to have a maximum close to the critical point.

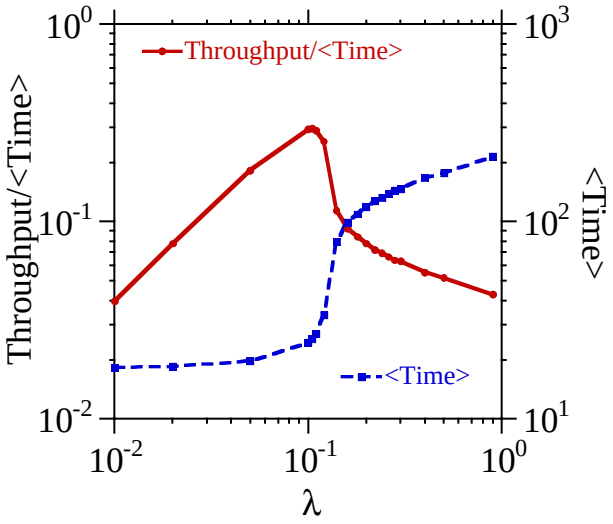

Fig. 4. Efficiency and averaged time traveled for an N˚=˚20 network with congested signaling control method.

Near the critical point, the first 3 congestion control methods have effectively the same behavior, with a variation in the averaged traveled time of at the most a factor of 2. The exception is the backpressure scheme that will be discussed later on. The congestion control methods effectively maintain the creation rate of packets close to the critical value (Fig.˚5).

The two main parameters in the control methods, the time delay and the buffer size play an important role in the capability of the system to keep the packets moving and deliver them in the congested regime. If the value of one or both parameter is too small, the system collapses and the effective creation rate goes practically to zero. In the steady state and above the critical point seem to be only these two states. There is the normal operation state with effective packet creation rate close to the critical value and the collapsed operation state with the effective creation rate close to zero. This is illustrated in Fig.˚6 by changing the time delay for the congested signaling method. Similar results are obtained by

varying the buffer size and for the different congestion control methods. The exact value of these parameters for which the transition between states occurs is dependent on the congestion control method and the requested packet creation rate.

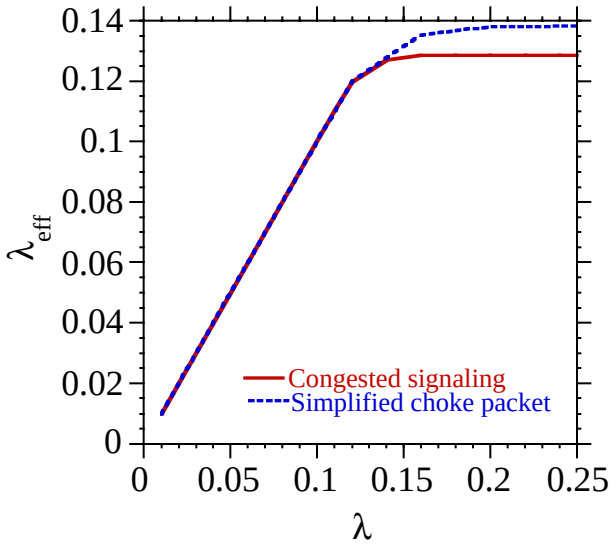

Fig. 5. Effective creation rate versus the requested creation rate of packets for to of the congestion control methods considered in this paper.

The backpressure scheme had a slight increase in time traveled around the critical point, but at higher values for λ returned to values seen before λ_{crit} was reached. This particular behavior in the backpressure method throughout the study prompted modified version of the technique to be created. In this altered backpressure technique, each host attempts to increase their value for λ by a given rate of change Δ′λ if they had been decreased by a router. This modified technique produced throughput and packet averaged time traveled curves that were consistent with the other techniques. Additionally, when the hosts were given an initial λ of one, this method produced an averaged λ that equaled λ_{crit}.

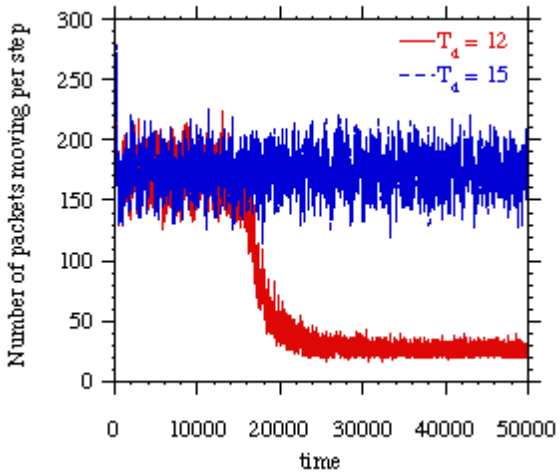

Fig. 6. Number of packets moved per time step for two values of the time delay using the congested signaling method.

An important consequence of operating near the critical transition point is that the probability distribution function of the travel time has an algebraic tail. In Fig. 7, we have plotted the probability distribution function of the time traveled by the packet for three of the methods considered. The plot is for λ =˚0.18, that is for a creation rate above the critical point. All PDFs show approximately a decade of algebraic decay with a decay index of about —1.3.

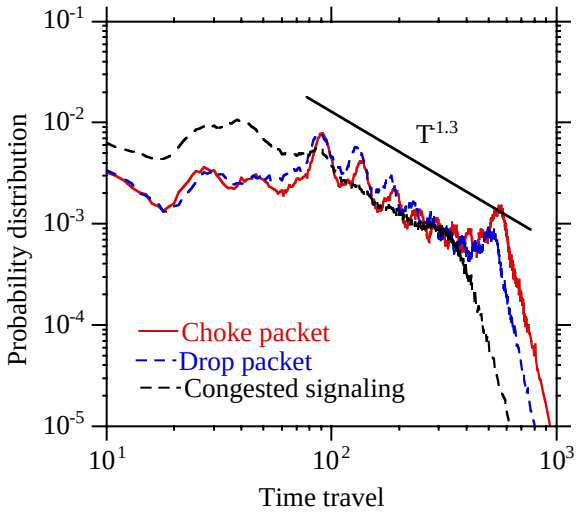

Fig. 7. Probability distribution function of the time traveled by the packet for three of the methods considered and for λ=0.18

For single packet dynamics as we have described, the distribution of time between arrival of packets in a given host node is Poisson like. This, results disagrees with many of the measurements of arrival times done in real networks. The main reason for that is that we have not considered trains of packets. As the host is ready to send some amount of information over the network, this information is broken in a number of packets that are being sent. These packets are initially correlated and this correlation is translated in the breaking of Poisson statistics. We have implemented a variation in the dynamics to take into account the size of the information to be sent. The method that we have implemented is as follows. If a host pulls decides to send a message by picking a random number, it also selects randomly the number of packets to be send. If there is more than one packet, they will be sent in successive time steps. In this situation the tail of the PDF of arrival times changes from exponential to an algebraic tail. An example is shown in Fig. 8, where the PDF of the time between arrivals of packets is plotted for different values of the maximum possible length of an individual message. As we allow for larger messages, a longer algebraic tail develops.

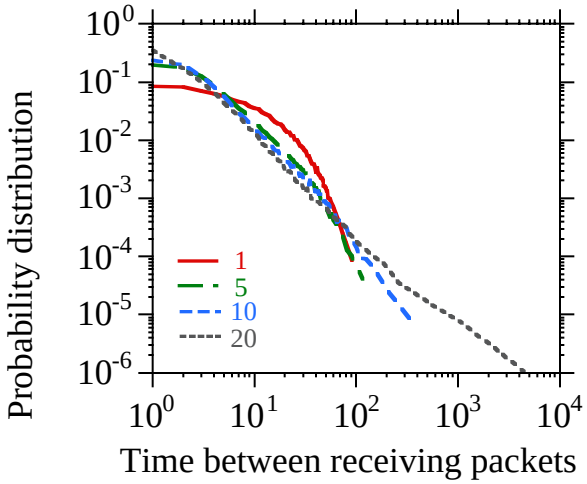

Fig. 8. Probability distribution function of the times between packets for different values of the maximum number of possible packets in each message sent.

5. Dynamics of the packet traffic in real communications networks

In order to make comparisons to the modeling results, we have performed a series of analyses to compare data traffic flow over two types of network connections. Others have performed similar analysis on a variety of network data as discussed earlier. Our analysis was performed on data collected over the period of two weeks by continuous pinging of two sites. One site had a path that remained in the ESnet network (swxaty-p2-9.ens.ornl to brxatyf1.ens.ornl.to orgwy.ctd.ornl to ornl-rt3f0.ctd.ornl to lbl-atms.es.to lbnl2-es-fddi.es.to sas.nersc.gov) while the other site had a path dominantly on the open Internet. Esnet was a closed network for energy systems research that had a large amount of traffic but was usually not near its limit. Both paths had multiple hops, with the total paths being from Oak Ridge to LBNL for one and the other from Oak Ridge to UW-Madison. The pinging was done with a one second interval while retaining the round trip time as well as a packet loss index. The raw data itself is instructive with clearly visible correlations much longer then a few ping times in the busy (stressed) network (Fig.˚9) and much less apparent correlations in the less stressed (ESnet) network data (Fig.˚10).

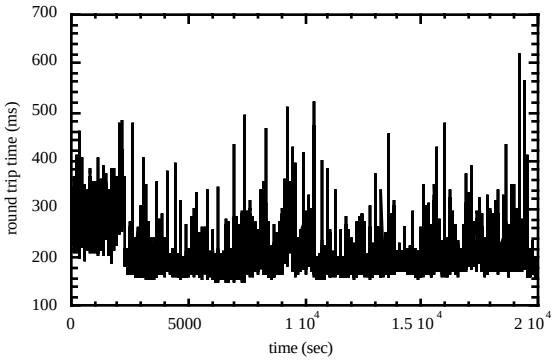

Fig˚9 Raw data from the busy (open internet) route

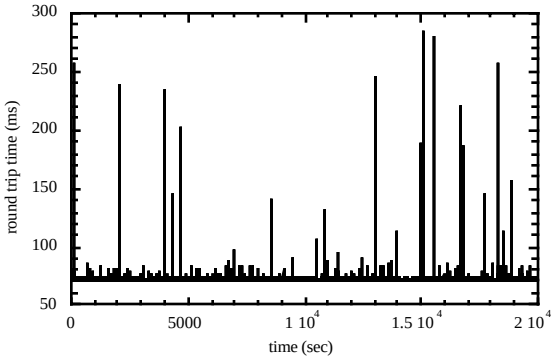

Fig 10. Raw data from the ESnet route.

More significantly, the power spectrum of the round trip ping time displays the 1/f behavior characteristic of

a SOC system both for the busy and the less stressed network (Fig. 11). In the highly active system (the open internet route) the 1/f region extends over nearly 2 decades reminiscent of a strongly driven SOC system. In the less stressed (ESnet) system the 1/f regime is much smaller showing much less power in the low frequency regime. This could suggest that even a network system that is not close to its operational limit still displays SOC behavior. This could have implications for both how to control and prevent network collapses as well as providing a possible tool for diagnosing the proximity of a network to trouble. Another test of long time correlations in systems is the R/S statistic[12-14]. A Hurst exponent of 0.5 signifies gaussian random numbers and therefore no long time correlations. In the busy (open internet) route there is a large regime in which the Hurst exponent (the slope on a log-log graph) is close to 1 signifying strong long time correlations. The less busy network shows much less long time correlated dynamical behavior which is consistent with the power spectra.

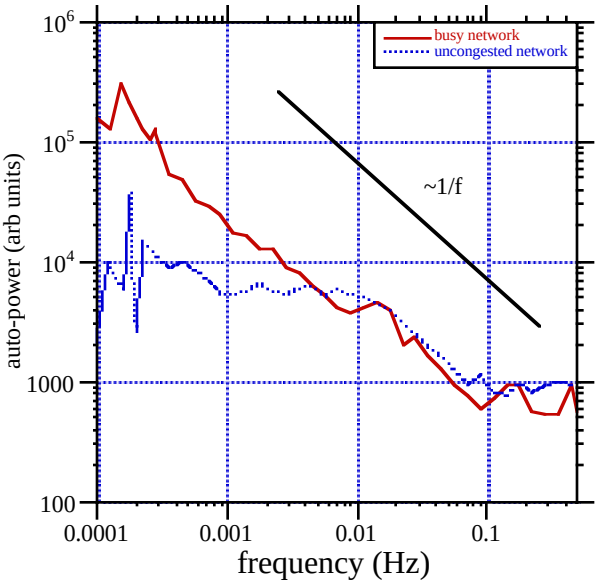

Fig˚11. Frequency spectra for the two routes showing a clear 1/f region in both data sets. The range of the 1/f region is much larger in the busy networks data set.

The PDFs of the travel time are another instructive measure. Figure 12 shows the travel time PDF for the busy network and for the model with a value for λ of 0.12. The model data is rescaled so as to have the same time scale as the real network data, a reasonable thing to do because the definition of time and the system size are clearly different. It should be noted that the meaning of this rescaling could be related to the intrinsic simplifications in the system that lead to the fixed time scale. The value of 0.12 is important because it is approximately the critical value of λ. The agreement between the model data and the real network data is striking and is made even more sticking by noting that for other values of λ the curves do not overlap as well or sometimes at all. This suggests that the real network maybe operating near its critical point. If so, this is likely because of the self-organization of the system through its self-regulating mechanisms.

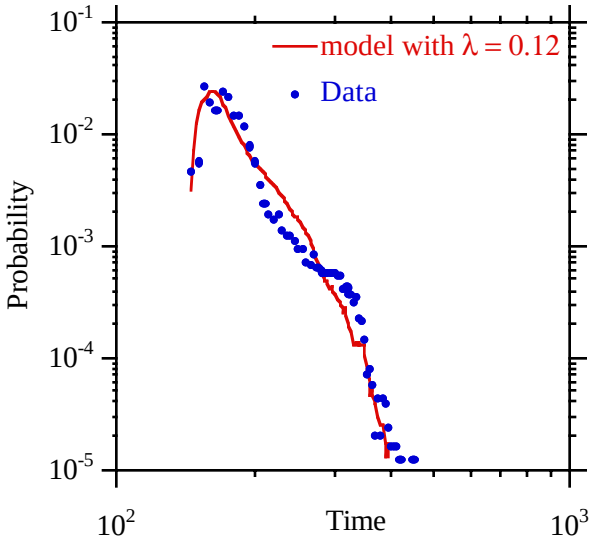

Fig. 12 Probability distribution function of packet travel time for real busy internet data and for rescaled model results with λ=0.12.

6. Conclusions

A simple dynamical model of communication network traffic displays a phase transition between a smoothly flowing regime and a clogged regime. When simple self-regulating congestion rules are added to the model, the system seems to self-organize at a point near to but not necessarily right at the critical phase transition point. The results of this could help explain how such a system finds its operating point. Analysis of real network data suggests that busy networks (such as the internet) have long time correlations and other characteristics of SOC systems. Perhaps even more importantly, the PDFs of the internet travel times closely resembles the model data for values of the control parameters very close to the critical point. This suggests that the real network might be operating near to its critical point which has important implications for the probability of disruptions and network risk analysis

and suggests that the real system could have intrinsic vulnerability to large-scale disruptions.

In these studies, we have limited the application of the model to simple networks systems with single self-regulation schemes. Further work to allow more realistic systems dynamics and regulation are being undertaken.

Acknowledgement

Part of this research has been carried out at Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract number DE-AC05-00OR22725. David Newman gratefully acknowledges support in part from NSF grant ECS–0085647.

References

[1] P. Bak, C. Tang and K Weisenfeld, *Self-Organized Criticality: An Explanation of 1/f Noise,* Phys. Rev Lett. 59, 381 (1987).

[2] R. Jain and S. A. Routhier IEEE j. Selected areas Communications, SAC.-4, 986 (1986).

[3] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, *On the self-similar nature of ethernet traffic*, *IEEE/ACM Transaction on Networking*, vol. 2, pp. 1-15, 1994.

[4] H. J. Fowler and W. E. Leland, Local area network traffic characteristics, with implications for broadband network congestion management, IEEE J. Selected Areas Communications, vol. SAC-9, pp. 1139-49, 1991.

[5] W. Willinger, M. S. Taqqu, W. E. Leland, and D. V. Wilson, Statistical Science 10, 67 (1995).

[6] K. Nagel and M. Paczuski, *Emergent traffic jams,* Phys. Rev. E 51, 2909 (1995).

[7] M. Takayasu, H. Takayasu, and K. Fukuda, Physica A 277, 248 (2000).

[8] T. Ohira and R. Sawatari, Phys. Rev. E 58, 193 (1998).

[9] R. V. Sole and S. Valverde, Physica A 289, 595 (2001).

[10] J. Yuan, Y. Ren, and X. Shan, Phys. Rev. E 61, 1067 (2000).

[11] W. Stallings, Data and Computer Communications (6^{th} ed.). Upper Saddle River, NJ: Prentice Hall, 2000.

[12] H. E. Hurst, *Long-term storage capacity of reservoirs*, *Trans. Am. Soc. Civil Eng.* 116, 770 (1951).

[13] B. B. Mandelbrot and J. R. Wallis, *Noah, Joseph, and operational hydrology*, *Water Resources Research* 4, 909-918 (1969).

[14] B. B. Mandelbrot and J. R. Wallis, *Some long-run properties of geophysical records, Water Resources Research* 5, 422-437 (1969).