

Data Traffic Dynamics and Saturation on a Single Link

Reginald D. Smith

Abstract—The dynamics of User Datagram Protocol (UDP) traffic over Ethernet between two computers are analyzed using nonlinear dynamics which shows that there are two clear regimes in the data flow: free flow and saturated. The two most important variables affecting this are the packet size and packet flow rate. However, this transition is due to a transcritical bifurcation rather than phase transition in models such as in vehicle traffic or theorized large-scale computer network congestion. It is hoped this model will help lay the groundwork for further research on the dynamics of networks, especially computer networks.

Keywords—congestion, packet flow, Internet, traffic dynamics, transcritical bifurcation

I. INTRODUCTION

The 1969 the Internet (then ARPANET) was first established as a distributed packet communications network that would not only reliably operate if some of its nodes were destroyed in an enemy attack, but allow easier communications of computer research results by universities. Today the Internet has grown to become a sprawling network of every aspect of humanity dwarfing previous technological mediums in both complexity and behavior. It was therefore only a matter of time that advanced statistical techniques, such as those developed by physicists in statistical mechanics were applied to investigate it.

Since the late 1990s, the Internet has been of Interest to the physics community, becoming aware to most in the seminal Nature paper of Watts and Strogatz [1]. This was continued or paralleled by the work of countless others [2], [3], [4], [5], [6], [7]. However, until recently this research has focused mostly on the topological aspects of networks and much less on dynamics. A particularly fertile area on network dynamics, and one related to this paper is the study of phase transitions from free flow to congestion in computer networks [8], [9], [10], [11]. Most results give a critical packet flow on networks which separates free flow from congested traffic. There have been some investigations of dynamics aspects such as synchronization of coupled oscillators [12], [13], [14], [15] and some metabolic dynamics [16], [17] as well thus dynamics is rapidly moving from being a peripheral to a primary discussion about networks.

A very interesting reverse situation is visible in the studies of the statistical mechanics of vehicular traffic. Vehicle traffic on roads has been investigated, also using statistical mechanics, but focusing on the dynamics and flow of traffic versus the

topology of the road network [18], [19], [20]. Though there is some overlap between these two topics, traffic flow has been described in many ways with the most common description using the fundamental diagram of vehicle flow vs. vehicle density. Most models propose a two or three phase model of traffic. The first phase is free flow, where cars drive near the speed limit with little congestion and influence on each other's velocity. The final phase is congested traffic where traffic flow becomes spontaneously congested after reaching a critical density. In the three phase model, there is an intermediate phase called synchronized flow where traffic is not congested but cars match their speed at a reduced speed level effectively increasing the correlation length of the system as a prelude to congestion. In [21] Gábor and Csabai conduct a similar study comparing vehicle traffic flow to data packet flow using the number of TCP connections as the variable for flow density. They find that a fundamental diagram like pattern appears in data traffic when the flow density is modeled as the number of TCP connections between two different endpoints.

II. INTERNET TRAFFIC RESEARCH ISSUES

In investigating Internet dynamics, one is struck by how similar the dynamics of the network can be superficially similar to vehicle traffic. Internet networks are made of the flow of countless packets across network links and can be prone to the same free flow or congestion that vehicle traffic can be. Like traffic data, data on Internet traffic is easily obtainable either by setting up packet sniffers and traffic analyzers between certain nodes or using public data sets such as the WIDE Project's MAWI data traces from trans-Pacific US-Japan T1 lines [22] or the Stanford Linear Accelerator Center (SLAC)'s Pinger project which has monitored ICMP ping response against different nodes across the Internet on a continuous basis for years [23].

However, Internet traffic data analysis is complicated by many factors that are not easily accountable for in theoretical models. First, since the Internet is a decentralized network based on dynamic routing, the route of the traffic is not completely transparent. The path from one point to another can be fairly fluid and changing even within a continuous flow of transmitted data. Many of the intermediate routers or autonomous systems (AS), which are basically Internet service providers, have private and constantly changing routing rules and configurations that alter traffic in unpredictable ways to obtain certain quality of service metrics or traffic shaping priorities [24]. This makes analysis of data collected from Internet traffic fraught with questions and confusion as far as disambiguating the effect of network topology on dynamics.

Reginald Smith is a researcher with the Bouchet Franklin Institute in Rochester, NY. Address at PO Box 10051, Rochester, NY, 14610, rsmith@sloan.mit.edu

In addition to topological constraints, the dynamics of Internet traffic itself can affect dynamics studies. Under the Internet Protocol (IP) suite, there are many transport level protocols such as TCP and UDP and countless application level protocols such as HTTP among others. Protocols at all levels, including transport and application levels, can influence data traffic. For example, TCP has a congestion control algorithm which will actually throttle network speed given the feedback it receives from packet loss data on the network, requires periodic acknowledgements from the destination before sending more data, and will buffer data to send depending on the round trip time (RTT) of the connection in order to guarantee delivery [25], [26]. Therefore, measurements of TCP/IP network speeds, even in relatively "clean" networks, can actually be extremely complicated and dependent on much more than the topology of the network or volume of traffic.

Finally, the statistical nature of Internet traffic is still poorly understood. Internet traffic volumes and packet interarrival times to not follow typical distributions in other communications networks such as Poisson or Erlang distributions but rather exhibit bursty, self-similar traffic patterns which have been very difficult to model and predict [27], [28], [29], [30]. Like in measurements of Internet topology where a relatively small amount of nodes have many edges, almost everything in the Internet that can be measured seems to have a long-tailed distribution as a matter of course. TCP and UDP traffic flows follow this trend where a relatively small number of flows carry to bulk of data transferred (so-called "elephant flows") [31], [32]. The origins of these patterns of traffic are still a matter of research and debate. Add to this other inherent uncertainties and patterns in Internet traffic flow such as trimodal distributions of packet sizes [33], traffic spikes to due to malicious code such as viruses or Trojan directed distributed denial-of-service attacks [34], [35], and periodicities in the volume of traffic caused by 12 hour, 24 hour, and 7 day cycles (with a 3.5 day harmonic) [36], [37]. There is also an issue of long-range correlations of Internet traffic between different routers which can be relatively uncorrelated or very correlated depending on the nature of traffic and the level of congestion [38]. All of these factors are mentioned to demonstrate that modeling and understanding Internet traffic dynamics is a problem likely of greater magnitude than topological analysis. Given these difficulties and more, understanding the basics of traffic dynamics and the interactions between topology and dynamics in computer networks is essential to understanding theoretical aspects, creating accurate simulations, and conducting useful experiments.

III. EXPERIMENTAL SETUP

The purpose of the experiment described by this paper is to ask a basic question about traffic dynamics: how is the throughput (speed) of a link affected by three fundamental variables that determine the nature of network traffic: the average packet size, the average packet flow rate, and the bandwidth of the link. The bandwidth of the network, in this case 100 Mbps (megabits per second), is the theoretically maximum possible throughput. The throughput itself can be

| A | B | C | D | E |
|----|----|---|--------|---|
| 14 | 20 | 8 | 1-1472 | 4 |

Fig. 1. Structure of a packet in this paper. Proportions based on a 50 byte payload. Numbers are size of headers or payload in bytes. A is the Ethernet header which contains MAC address source and destination and payload type, B is the Internet Protocol (IP) header, C is the User Datagram Protocol (UDP) header, D is the data payload and E is the Ethernet CRC checksum hash to prevent accidental corruption of the frame.

described in terms of the average packet size and average flow rate by

$$\langle T \rangle = \langle p \rangle \langle \lambda \rangle \quad (1)$$

Where $\langle T \rangle$ is the average throughput of the link, p is the average size of packets in a transmission over a given period of time and $\langle \lambda \rangle$ is the average flow rate in packets per second. In this experiment since all packets will be off the same size in each sample we can say

$$\langle T \rangle = p \langle \lambda \rangle \quad (2)$$

Two computers, both running Windows XP, were connected using an Cat 5e cable that connected to the Ethernet network adapter (NIC) of both computers. The Ethernet flow control was disabled to ensure that the flow of data, and not signaling between the computers to prevent dropped packets that flow control entails, determines the throughput. Traffic was generated using the program Iperf which generates a stream of identically sized packets at a throughput inputted by the user. The throughput in Iperf was designated at 100 Mbps in order to test what the maximum throughput the network would actually demonstrate given an attempt at maximum throughput. The transport protocol UDP was chosen over TCP for several reasons. UDP is a connectionless protocol, meaning it does not guarantee delivery, and will solely submit a string of packets. TCP in a connection based protocol whose delivery guarantee requires frequent "handshaking" between the source and destination and whose congestion control algorithm can affect throughput in a non-trivial fashion delivering a lower throughput due to protocol software, not the maximum throughput of the link. Finally, TCP can give different performance over links with different latencies, as measured by packet RTT, so the results may not be subject to larger generalization [25].

To test the performance of the network under different sized packets, the UDP packet payload was varied from 25 bytes up to 1450 bytes in 25 byte increments. In order to ensure that equation 2 holds, you must ensure that all packets are the same size. The structure of packets under Ethernet/IP is shown in figure 1. The payload, a variable selected in the Iperf software, is encapsulated by a header for UDP (8 bytes), IP (20 bytes), and a header and footer in the Ethernet frame (total 18 bytes). Frame is just a general term for an Ethernet packet. These headers mostly provide routing data, priorities, checksums, and other information important to packet logistics. In addition, in standard Ethernet the maximum frame size, minus Ethernet headers and footers, is 1500 bytes. With the Ethernet overhead the total maximum size for an Ethernet frame is 1518 bytes.

Therefore, at 1475 bytes payload, the total would be 1503 bytes and you would have packet fragmentation - instead of one frame you would have two, one with a frame payload of 1500 bytes and a second with a frame payload of 3 bytes. This would affect throughput by showing a sudden change since average frame payload size would drop to about 750. Given the phenomenon of fragmentation, an actual analysis of the effect of packet size on throughput is only useful up the fragmentation size limit.

In the experiment, Iperf delivered a packed stream of UDP traffic from the client to the server computer, trying to send as close to bandwidth as possible, and outputted the average throughput in Kbps (kilobits per second). It also gave the packet loss, and a measure called jitter which is not used but measures the deviation in packet interarrival times versus interdeparture times. For throughput, Iperf measures a related measured called goodput which measures data speed in terms of the payload size, not including any packet overhead in the calculation of bytes transferred. However, the packet flow rate is accurate and is calculated from equation 2. Therefore the actual throughput was recalculated using packet sizes that include both payload and packet overhead.

IV. RESULTS

At all packet sizes, packet loss was small, much less than 1%. The results of the experiment are shown in figure 2. First, it is clear that throughput decreases with decreasing packet size. This first fact is well-known in the network engineering community [39]. This is an inherent property of all network adapters, Ethernet or otherwise. In fact, one of the key requirements in next generation networks is the network capability to send "jumbo frames" where fragmentation limits at the network layer are much larger than 1500 bytes, up to 9000 bytes in some cases. These larger packet sizes cause higher throughput and more efficient networking because within the network adapter and computer hardware, there is a per-packet processing overhead.

Despite the bandwidth rating of network adapters, be it 10, 100, or 1000 Mbps, there is a maximum packet flow that they can effectively handle. Given equations 1 and 2 it is clear that to maintain any given throughput, by lowering the packet size you are increasing the packet flow. Because of the packet flow processing bottleneck in the hardware, however, this can make high throughput impossible at low packet sizes.

As seen in figure 2, for large packet sizes, the throughput is very close to bandwidth and can be roughly equivalent to free flow traffic. At a specific critical packet size, however, p_c , the throughput begins to rapidly degrade to the point it is only about 6% of bandwidth at 25 bytes. This is the saturated state. Here saturation is used instead of congestion since congestion is usually a network wide phenomenon while this packet slowdown in throughput is due to overwhelming the processing power at the NIC. In a superficial way this behavior is similar to the fundamental diagram flow-density curve in vehicle traffic. Packet flow increases with smaller packet size until saturation forces packet flow to begin to slow its increase and approach a maximum value. Comparisons

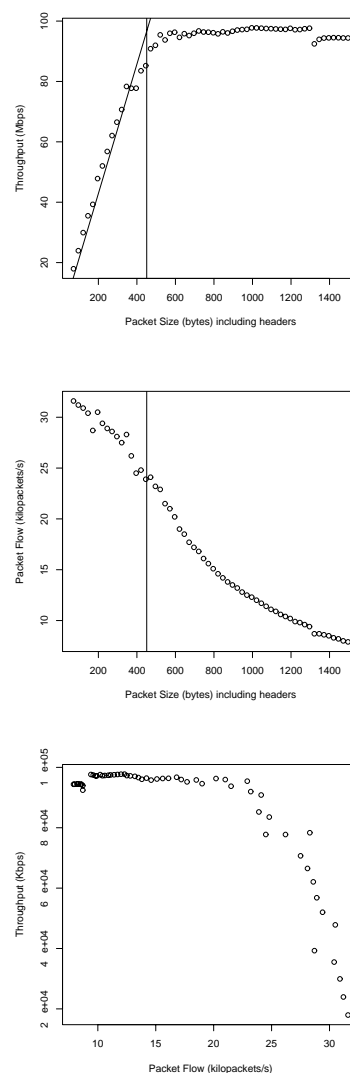


Fig. 2. Graphs of the average throughput vs. packet size, average packet flow vs. packet size, and average throughput vs. average packet flow respectively. Vertical lines on the first two graphs represent the calculated critical packet size of 451 bytes. The fitted line on the first graph is the line predicted by maximum packet flow. The slight decrease at high packet size in the first graph is due to unknown system effects and not inconsistent enough to reject experimental results.

between vehicle traffic should be qualified though. In vehicle traffic, there is interaction between cars on the road giving rise to the collective dynamics which justify a statistical mechanics interpretation. On data networks, packets do not interact with each other and packet collisions are errors, not intrinsic aspects of the packet flow. Therefore, as will be shown below, the transition from free flow to saturation should be viewed as a bifurcation in the system dynamics, not as a phase change. On the network level where there are many interacting nodes, perhaps congestion can be seen as a phase change but this perspective is not appropriate at the single link level.

As stated earlier, in free flow the throughput is nearly bandwidth and comparatively, though not completely, steady state. In this region, there is a mutual relationship between

packet size and packet flow. Differentiating equation 2 we have

$$d\langle T \rangle = p d\langle \lambda \rangle + \lambda dp \quad (3)$$

assuming $d\langle T \rangle$ is 0 in free flow regardless of the packet size or flow we can conclude

$$\frac{dp}{d\langle \lambda \rangle} = -\frac{p}{\langle \lambda \rangle} \quad (4)$$

So there is a tradeoff curve, like the production possibility frontiers in economics, between packet size and flow in free flow traffic. Also,

$$\frac{dp}{p} = -\frac{d\langle \lambda \rangle}{\langle \lambda \rangle} \quad (5)$$

demonstrating that every increase in packet flow is matched by a corresponding decrease in packet size and vice versa. Because the networking and computer equipment have various processes and imperfections our free flow region never reaches bandwidth and steadily erodes with smaller packet size, however, the high throughput feature is relatively constant compared to the saturated state.

At a packet size p_c , we have an increasingly rapid breakdown in throughput. This corresponds approximately to the maximum flow to the network adapter and the breakdown into saturation. We calculate the theoretical maximum flow as

$$\langle \lambda \rangle_c = \frac{\langle T \rangle_{max}}{p_c} \quad (6)$$

where in the theoretically ideal situation $\langle T \rangle_{max} = B$ where B is the bandwidth. Therefore in the saturated region, the throughput $\langle T \rangle$ is given by

$$\langle T \rangle = \frac{p}{p_c} \langle T \rangle_{max} \quad (7)$$

so the ratio of packet size to the critical packet size determines the throughput under saturation compared to the maximum possible throughput. In the first graph of figure 2, is a comparison of this prediction with the observed data in the saturated region where $\langle T \rangle_{max}$ is about 96 Mbps. Given data from the saturated region for $\langle T \rangle$, $\langle T \rangle_{max}$, and p we can estimate the critical packet size p_c . In this case, the critical packet size is approximately 451 bytes at around 25k packets/s. Though this model seems to accurately predict the throughput values for almost the entire region of saturation, there is a seemingly glaring contradiction in the second graph of figure 2 where the packet flow keeps increasing with decreasing packet size. Though the packet flow does not stagnate, as it should according to ideal theory, the third graph shows that the decline in throughput over a fairly short range of packet flow demonstrates packet flow is the limiting factor in the saturated region. In the congested region the packet flow increases from about 25 to 30 thousand flows per second which is a definite increase but relatively small compared to the region of free flow when it increased from 7 to about 25 without affecting throughput. Note, p_c is specific to the equipment and configuration used and does not have a general value of 451 bytes.

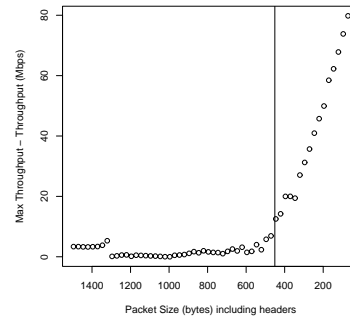


Fig. 3. Graph of the maximum throughput measured minus throughput vs packet size on a reverse axis. The estimated critical packet size is given by the vertical line.

V. BIFURCATION ANALYSIS

The fact that the stable throughput of the system changes with the parameter of packet size leads us to suspect a possible bifurcation. The exchange of stability between quasi-bandwidth throughput and packet flow limited throughput leads to the hypothesis of a transcritical bifurcation with stability changing at p_c . Assuming $\langle T \rangle_{max} = B$ the two stable throughput regimes are $\langle T \rangle = B$ and $\langle T \rangle = (p/p_c)B$, so

$$d\langle T \rangle / dt = (\langle T \rangle - B) \left(\frac{p}{p_c} B - \langle T \rangle \right) \quad (8)$$

and

$$d\langle T \rangle / dt = -\frac{p}{p_c} B^2 + \langle T \rangle B \left(1 + \frac{p}{p_c} \right) - \langle T \rangle^2 \quad (9)$$

This can be seen as the traditional form of a transcritical bifurcation where $dx/dt = px - x^2$, however, this is made more clear by making the independent variable $B - \langle T \rangle$:

$$d(B - \langle T \rangle) / dt = B \left(1 - \frac{p}{p_c} \right) (B - \langle T \rangle) - (B - \langle T \rangle)^2 \quad (10)$$

which fits the normal form for a transcritical bifurcation, $dx/dt = px - x^2$. The transcritical bifurcation can be more clearly seen in the graph of $(\langle T \rangle_{max} - \langle T \rangle)$ vs. the reverse order axis with p in figure 3.

VI. DISCUSSION

As mentioned before, the transition from free flow to saturated traffic here is a bifurcation not a phase change given the lack of interaction among the constituent particles in the system. This conclusion can also give pause to extrapolations of "classical" flow on network theory to complex networks such as computer networks. If looking at the weights and capacities of the network from the perspective of throughput, typical maximum flow algorithms such as max-flow min-cut may give incorrect answers if the limiting aspect of the flow is the packet flow rate, size of the packets, or another issue that is not plainly visible. Flows in computer networks do not behave like incompressible fluid or similar flows assumed in most flow models where any flow rate freely flows up to the capacity minus costs incurred by friction, etc.

Another point is other similarities to vehicle traffic dynamics. Interestingly, the first graph of Figure 2 is similar to a flow-density curve seen in traffic models. Even equation 1 has similarity to such traffic models where

$$Q = DV \quad (11)$$

Where Q is the traffic flow in cars/h and D is traffic density in cars/km and V is flow velocity in km/h. There is a critical traffic density that separates free flow from synchronized flow or congestion that plays a similar role to packet size in data networks.

This paper does not address the larger problem of dynamics on networks, however, it is the contention of this paper that understanding the simple dynamics at the network level is essential to understanding the wider implications of network sized dynamics. Given the problems with using Internet traffic data described earlier, more understanding in the basics of traffic dynamic may be obtained through experimental setups or computer network simulations. Future research should look at how topological invariants in networks affect dynamics and how dynamics may affect the evolution of networks and changes in their topological invariants.

As a final note regarding research on dynamics in networks, particularly data networks and the Internet, the author believes that more interaction and cross-referencing between the engineering and physics communities will help promote better understanding and advancement. Although there is research at an almost feverish pitch in both the physics and engineering community on networking, both sets of publications seem to be almost entirely ignorant of each other. In physics, one of the only such publications commonly cited is the Faloutsos team's work on the router topology of the Internet [6]. Engineering researchers, in journals such as those published by IEEE and ACM, also have rarely quoted physics literature besides the most well-known papers of Barabási, Watts, or Strogatz and seem largely unaware of the later results being reached by physics in the topology of networks. Without a detailed understanding of the network protocols and engineering literature regarding the function of data networks and the Internet, this paper would not have been possible. A similar situation is seen in sociology where a wealth of research has been done on aspects of social network dynamics such as the diffusion of trends along networks, in a quantitatively rigorous fashion, but again both communities do not usually inform each other of progress or motivate collaboration. As the study of networks expands and matures, it will become necessary to read and reach across interdisciplinary boundaries to share tools and knowledge that will allow the most profound and predictive insights to be reached.

REFERENCES

- [1] D.J. Watts and S.H. Strogatz, "Collective dynamics of 'small-world' networks", *Nature*, vol. 393, pp.440-442, 1998
- [2] R. Albert, H. Jeong, and A.L. Barabási, "Diameter of the World-Wide Web", *Nature*, vol. 401, pp.130-131, 1999
- [3] M.E.J. Newman, "Scientific collaboration networks. I. Network construction and fundamental results", *Phys. Rev. E*, vol. 64, 016131, 2001
- [4] M.E.J. Newman, "Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality", *Phys. Rev. E*, vol. 64, 016132, 2001
- [5] S.N. Dorogovtsev and J.F.F. Mendes, "Exactly solvable small-world network" *Europhys. Lett.* vol. 50, pp. 1-7, 2000
- [6] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology", *Comp. Comm. Rev.*, vol. 29, pp.251-262, 1999
- [7] R. Pastor-Satorras R and A. Vespignani, "Epidemic spreading in scale-free networks", *Phys. Rev. Lett.*, vol. 86, pp.3200-3203, 2001
- [8] T. Ohira and R. Sawatari, "Phase transition in a computer network traffic model", *Phys. Rev. E*, vol. 58, pp.193-195, 1998
- [9] M. Takayasu, H. Takayasu, and K. Fukuda, "Dynamic phase transition observed in the Internet traffic flow", *Physica A*, vol. 277, pp.248-255, 2000
- [10] R.V. Solé and S. Valverde, "Information transfer and phase transitions in a model of internet traffic ", *Physica A*, vol. 289, pp.595-605, 2001
- [11] L. Zhao, Y.C. Lai, K.H. Park, and N. Ye, "Onset of traffic congestion in complex networks", *Phys. Rev. E*, vol. 71, 026125, 2005
- [12] L. Lago-Fernández, R. Huerta, F. Corbacho, and J. Sig'uenza, "Fast Response and Temporal Coherent Oscillations in Small-World Networks", *Phys. Rev.Lett.*, vol. 84, pp.2758-2761, 2000
- [13] S.F. Wang and G.R. Chen, "Synchronization in small-world dynamical networks", *Intl J. of Bifurcation and Chaos*, vol. 12, pp.187-192, 2002
- [14] H. Hong, Y.M. Choi, and B.J. Kim, "Synchronization on small-world networks", *Phys. Rev. E*, vol. 65, 026139, 2002
- [15] M. Barahona and L.M. Pecora, "Synchronization in Small-World Systems", *Phys. Rev. Lett.*, vol. 89, 054101, 2002
- [16] E. Almaas, A.N. Oltvai, and A.L. Barabási, "The activity reaction core and plasticity of metabolic networks" *PLoS Comp. Bio.*, vol. 1, pp.557-563, 2005
- [17] R. Guimera and L. Nunes-Amaral, "Functional cartography of complex metabolic networks" *Nature*, vol. 433, pp.895-900, 2005
- [18] B.S. Kerner and H. Rehborn, "Experimental properties of phase transitions in traffic flow", *Phys. Rev. Lett.*, vol. 79, pp.4030-4033, 1997
- [19] C.F. Daganzo, M.J. Cassidy, and R.L. Bertini, "Causes And Effects Of Phase Transitions In Highway Traffic (Research Report)", University of California, Berkeley Institute of Transportation Studies Research Report, Report No. UCB-ITS-RR-97-8, 1997
- [20] D. Chowdhury, L. Santen, and A. Schadschneider, "Statistical Physics of Vehicular Traffic and Some Related Systems" *Phys. Rep.*, vol. 329, pp.199-329, 2000
- [21] S. Gábor and I. Csabai, "The analogies of highway and computer network traffic", *Physica A*, vol. 307, pp.516-526, 2002
- [22] Widely Integrated Distributed Environment (WIDE) Project, Kanagawa, Japan, MAWI Working Group Traffic Archive (WIDE Backbone traffic traces)[Online]. Available: <http://mawi.wide.ad.jp/mawi/>
- [23] Stanford Linear Accelerator Center (SLAC) PINGER End-to-End Performance Measuring Project [Online]. Available: <http://www-iepm.slac.stanford.edu/pinger/>
- [24] H. Chang, S. Jamin, Z.M. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices" in *Proc. of the 5th ACM SIGCOMM Conference on Internet Measurement*, Berkeley, CA 2005, pp.139-152
- [25] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, "The macroscopic behavior of the TCP congestion avoidance", *Comp. Comm. Rev.*, vol. 27, no. 3, pp.67-82, 1997
- [26] A. Veres and M. Boda, "The Chaotic Nature of TCP Congestion Control", in *Proc. of INFOCOM 2000*, vol. 3, Tel Aviv, Israel, 2000, pp. 1715-1723
- [27] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)", *IEEE/ACM Trans. on Networking*, vol. 2, pp.1-15, 1994
- [28] A. Feldmann, A. Gilbert, and W. Willinger, "Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic", *Proc. of ACM SIGCOMM '98*, Vancouver, Canada, 1998, pp.42-55
- [29] S. Uhlig, "Non-stationarity and high-order scaling in TCP flow arrivals: a methodological analysis", *Comp. Comm. Rev.*, vol. 34, no. 2, pp.9-24, 2004
- [30] K. Fukuda, H. Takayasu, and M. Takayasu, "Origin of Critical Behavior in Ethernet Traffic", *Physica A*, vol. 287, pp.289-301, 2000
- [31] W. Fang and L. Peterson, "Inter-AS traffic patterns and their implications" in *Proc. Of the Global Telecommunications Conference, GLOBECOM '99*, Rio de Janeiro, Brazil, vol. 3, 1999, pp.1859-1868
- [32] T. Mori, R. Kawahara, S. Naito, and S. Goto, "On the Characteristics of Internet Traffic Variability: Spikes and Elephants", in *Proc. of the 2004 International Symposium on Applications and the Internet*, Tokyo, Japan, 2004, pp.99-106

- [33] C. Dovrolis, P. Ramanathan, and D. Moore, "What Do Packet Dispersion Techniques Measure?" in *Proc. of INFOCOM 2001* Anchorage, AL, vol. 2, 2001, pp.905-914
- [34] J. Cowie and A. Ogielski, "Global Routing Instabilities during Code Red 2 and Nimda Worm Propagation (conference report)", CAIDA Internet Statistics and Metrics Analysis Workshop, San Diego, CA, December 2001
- [35] D. Moore, C. Shannon, D.J. Brown, G.M. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity" *ACM Trans. on Comp. Sys.*, vol.24, no. 2, pp.115-139, 2006
- [36] X.M. He, C. Papadopoulos, J. Heidemann, and A. Hussain, "Spectral Characteristics of Saturated Links (research report)", University of Southern California USC/CS Report, Los Angeles, CA, 2004
- [37] K. Papagiannaki, N. Taft, Z.L. Zhang, and C. Diot, "Long-term forecasting of Internet backbone traffic: observations and initial models" in *Proc. of INFOCOM 2003*, San Francisco, CA, vol. 2, 2003, pp.1178-1188
- [38] M. Barthélemy, B. Gondran, and E. Guichard, "Large scale cross-correlations in internet traffic", *Phys. Rev. E*, vol. 66, 056110, 2002
- [39] M. Mathis, "Pushing up the Internet MTU (conference report)", Joint-Techs Workshop, Miami, FL, 2003



Reginald Smith Reginald Smith, received his B.A. in physics and his B.S. in Commerce concentrating on management information systems and finance from the University of Virginia, Charlottesville, VA in 2002. He received his Masters in Business Administration from the MIT Sloan School of Management in 2007. He has published papers on complex networks, Internet traffic dynamics, plasma physics, and computational linguistics. He is a long-time member of the Society for Amateur Scientists, the premiere global organization for amateur scientists started by

McArthur fellow Shawn Carlson. He currently resides in Rochester, NY where he works in supply chain management for ITT.