

Scale-free behavior of the Internet global performance

R. Percacci¹ and A. Vespignani^{2,a}

¹ International School for Advanced Studies SISSA/ISAS, via Beirut 4, 34014 Trieste, Italy

² Laboratoire de Physique Théorique^b, bâtiment 210 Université de Paris-Sud 91405 Orsay Cedex, France

Received 25 January 2003

Published online 7 May 2003 – © EDP Sciences, Società Italiana di Fisica, Springer-Verlag 2003

Abstract. Measurements and data analysis have proved very effective in the study of the Internet's physical fabric and have shown heterogeneities and statistical fluctuations extending over several orders of magnitude. Here we focus on the relationship between the Round-Trip-Time (RTT) and the geographical distance. We define dimensionless variables that contain information on the quality of Internet connections finding that their probability distributions are characterized by a slow power-law decay signalling the presence of scale-free features. These results point out the extreme heterogeneity of Internet delay since the transmission speed between different points of the network exhibits very large fluctuations. The associated scaling exponents appear to have fairly stable values in different data sets and thus define an invariant characteristic of the Internet that might be used in the future as a benchmark of the overall state of "health" of the Internet.

PACS. 89.75.-k Complex systems – 05.70.Ln Nonequilibrium and irreversible thermodynamics – 87.23.Ge Dynamics of social systems

The Internet is a distributed network whose size has already scaled five orders of magnitude since its inception. Given the extremely complex and interwoven structure of the Internet, several research groups started to deploy technologies and infrastructures aiming to obtain a more global picture of the Internet. This has led to very interesting findings concerning the Internet maps topology. Connectivity and other metrics are characterized by algebraic statistical distributions that signal fluctuations extending over many length scales [1–5]. These scale-free properties and the associated heterogeneity of the Internet fabric define a large scale object whose properties cannot be inferred from local ones, and are in sharp contrast with standard static graph models with Poissonian degree distributions. The importance of a correct topological characterization of the Internet in routing protocols and the parallel advancement in the understanding of scale-free networks [6] have triggered a renewed interest in Internet measurements and modeling. Considerable efforts have been devoted also to the collection of end-to-end performance data by means of active measurements techniques. This activity has stimulated several studies that, however, focus mainly on individual properties of hosts, routers or routes [7,8,5]. Only recently, an increasing body of work focuses on the performance of the Internet as a whole, especially to forecast future performance trends [9,10].

These measurements pointed out the presence of highly heterogeneous performance and it is our interest to inspect the possibility of an associated scale-free behavior.

The basic testing package for Internet performance is the original PING (Packet InterNet Groper) program. Based on the Internet Control Message Protocol (ICMP), Ping works much like a sonar echo-location, sending packets that elicit a reply from the targeted host. The program then measures the round-trip-time (RTT), *i.e.* how long it takes each packet to make the round trip. Organizations such as the National Laboratory for Applied Network Research (<http://moat.nlanr.net/>) and the Cooperative Association for Internet Data Analysis (<http://www.caida.org/>) use PING-like probes from geographically diverse monitors to collect RTT data to hundreds or thousands of Internet destinations. Our Internetwork Performance Measurement (IPM) project currently participates in the PingER monitoring infrastructure (<http://www-iepm.slac.stanford.edu/>). PingER was developed by the Internet End-to-end Performance Measurement (IEPM) group to monitor the end-to-end performance of Internet links. It consists of a number of beacon sites sending regularly ICMP probes to hundreds of targets and storing all data centrally. Most beacons and targets are hosts belonging to universities or research centers; they are connected to many different networks and backbones and have a very wide geographical distribution, so they likely represent a statistically significant sample of the Internet as a whole.

^a e-mail: alexv@th.u-psud.fr

^b UMR du CNRS 8627

We have analyzed two years worth of PingER data, going from April 2000 to March 2002. We have selected 3353 different beacon-target pairs, taken out of 36 beacons and 196 targets. For each pair we have considered the following metrics: the geographic distance of the hosts d (measured on a great circle), the monthly average packet loss rate r (the percentage of ICMP packet that does not reach the target point), the monthly minimum and average round-trip-times RTT_{\min} , and RTT_{av} , respectively. These data offer the opportunity to test various hypotheses on the statistical behavior of Internet delay. Each data point is the monthly summary of approximately 1450 single measurements. The geographic position of hosts is known with great accuracy for some sites, but in most cases it may be wrong by 10–20 km. Consequently, we have discarded pairs of sites that are less than this distance apart.

The end-to-end delay is governed by several factors. First, digital information travels along fiber optic cables at almost exactly $2/3$ the speed of light in vacuum. This gives the mnemonically very convenient value of 1ms RTT per 100 km of cable. Using this speed one can express the geographic distance d in light-milliseconds, obtaining an absolute physical lower bound on the RTT between sites. The actual measured RTT is (usually) larger than this value because of several factors. First, data packets often follow rather circuitous paths leading them through a number of nodes that are far from the geodesic line between the endpoints. Furthermore, each link in a given path is itself far from being straight, often following highways, railways or power lines [13]. The combination of these factors produces a purely geometrical enhancement factor of the RTT. In addition, there is a minimum processing delay δ introduced by each router or switch along the way, of the order of 50–250 μs per hop on average, summing up to a few ms for a typical path [13]. This can be significant for very close site pairs, but is negligible for most of the paths in the PingER sample. On top of this, the presence of cross traffic along the route can cause data packets to be queued in the routers/switches. Let t_R be the sum of all processing and queuing delays on a path. When the traffic reaches congestion, t_R becomes a very significant part of the RTT and packet loss also sets in.

We have considered minimum and average values of the RTT over one month periods. It is plausible that even on rather congested links there will be a moment in the course of a month when t_R is negligible, so RTT_{\min} can be taken as an estimate of the best possible communication performance on the given data path, subject only to the intrinsic geometrical enhancement factor and the minimum processing delay. For instance, for a US-Europe path, the night between Saturday and Sunday (US time, corresponding to Sunday morning in Europe) is a time when the links and nodes along the path can generally be expected to carry little traffic. On the other hand, RTT_{av} for a given site pair is obtained by considering the average RTT over one month periods. This takes into account also the average queuing delay and gives an estimate of the overall communication performance on the given data path.




Fig. 1. RTT_{\min} between 2114 host pairs (PingER data set of February 2002) as a function of their distance d . Each point correspond to a different host pair. The line indicates the physical lower bound provided by the speed of light in transmission cables. It is possible to observe the very large fluctuations in the RTT_{\min} of different host pairs separated by the same distance. For graphical reasons the picture frame is limited to 400 ms, however, several outliers up to 900 ms are present in the data set.

We studied the level of correlation between geographic distance and the RTT_{\min} and RTT_{av} of source-destination pair. In Figure 1 we report the obtained relationship for RTT_{\min} compared with the solid line representing the speed of light in optic fibers at each distance. While it is possible to observe a linear correlation of the RTT_{\min} with the physical distance of hosts, yet the data are extremely scattered. The RTT_{av} present a qualitatively very similar behavior, and it is worth remarking that both plots are in good agreement with similar analyses obtained from different data sets [10–12]. While several qualitative features of this plot provide insight into the geographical distribution of hosts and their connectivity, it misses a quantitative characterization of the intrinsic fluctuations of transmission delays and their statistical properties.

A more significant characterization of the end-to-end performance is obtained by normalizing the latency time by the geographical distance between hosts. This defines the absolute delay metrics $\tau_{\min} = \text{RTT}_{\min}/d$ and $\tau_{\text{av}} = \text{RTT}_{\text{av}}/d$ which represent the minimum and average delay time for unit distance, *i.e.* the inverse of the overall communication velocity (note that if we measure d in light-milliseconds τ_{\min} and τ_{av} are actually dimensionless). These metrics allow us to meaningfully compare the delay between pairs of hosts with different geographical distances. The highly scattered plot of Figure 1, indicates that end-to-end delay fluctuates conspicuously in the whole range of geographic distances. In particular,

looking at collections of host pairs at approximately the same geographical distance, we find delay times varying up to two orders of magnitude. The best way to characterize the level of fluctuations in delay times is represented by the probability $P(\tau_{min})$ and $P(\tau_{av})$ that a pair of hosts present a given τ_{min} and τ_{av} , respectively. In contrast with usual exponential or Gaussian distributions, for which there is a well defined scale, we find that data closely follow a straight line in a double logarithmic plot for at least one or two orders of magnitude, defining a power-law behavior $P(\tau_{min}) \sim \tau_{min}^{-\alpha_{min}}$ and $P(\tau_{av}) \sim \tau_{av}^{-\alpha_{av}}$. In Figure 2 we show the cumulative distributions $P_{cum}(\tau) = \int_{\tau}^{\infty} P(\tau')d\tau'$ obtained from the PingER data. If the probability density distribution is a power law $P(\tau) \sim \tau^{-\alpha}$, the cumulate distribution preserves the algebraic behavior and scales as $P_{cum}(\tau) = \tau^{-(\alpha-1)}$. In addition, it has the advantage of being considerably less noisy than the original distribution. From the behavior of Figure 2, a best fit of the linear region in the double logarithmic representation yields the scaling exponents $\alpha_{min} \simeq 3.0$ and $\alpha_{av} \simeq 2.5$. It is worth remarking that the presence of a truncation of the power law behavior for large values is a natural effect implicitly present in every real world data set. This is especially true in the present case where packet loss occurs after a maximum timeout. Power-law distributions are characterized by scale-free properties, *i.e.* unbounded fluctuations and the absence of a meaningful characteristic length usually associated with the probability distribution peak. In such a case, the mean distribution value and the corresponding averages are poorly significant, since fluctuations are gigantic and there are non negligible probabilities to have very large τ_{min} and τ_{av} compared to the average values in the whole system. In other words, Internet performances are extremely heterogeneous and it is impossible to infer local properties from average quantities.

Critical and scale-free behavior has been observed and characterized in queuing properties at router interfaces, probably affecting conspicuously the distribution of τ_{av} . It is, however, unclear why scale-free properties are observed also in the distribution of τ_{min} . In this case traffic effects should be negligible, and it is well known that the the distribution of hop counts between hosts has a well defined peak and no fat tails [12]. On the contrary, we find that minimum delays are distributed over more than two orders of magnitude. Potentially, cables wiggleness, Internet connectivity and hardware heterogeneities might be playing a role in the observed performance distribution.

A tendency to improved performance is observed over the two years period of data collections. Table 1 shows that the averages over all the site pairs of $\langle \tau_{min} \rangle$ and $\langle \tau_{av} \rangle$ decrease steadily, whereas the exponents α_{min} and α_{av} increase signalling a faster decay of the distribution tails. We can consider the improvement of the absolute delay as the byproduct of the technological drift to better lines and routers. On the other hand, the large fluctuations present in the Internet appear to be a stable and general feature of the statistical analysis. In order to have an independent check of the PingER results, we have considered also the

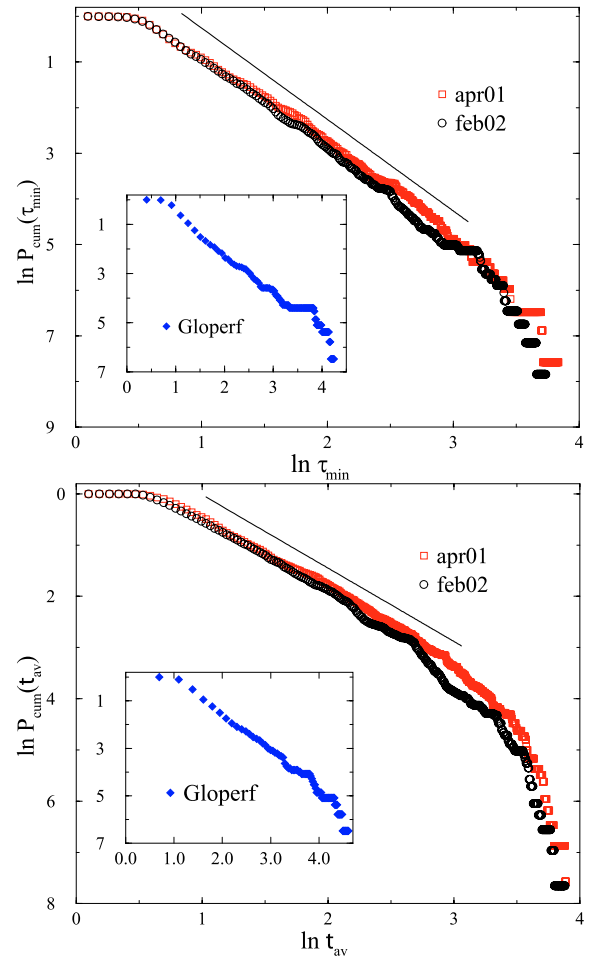


Fig. 2. Cumulative distributions, of the round-trip-times normalized with the actual distance d between host pairs. The linear behavior in the double logarithmic scale indicates a broad distribution with power-law behavior. (a) In the case of the normalized minimum round-trip-times τ_{min} , the slope of the reference line is -2.0 . (b) In the case of the normalized average round-trip-times τ_{av} , the reference line has a slope -1.5 . The insets of a) and b) report the distributions obtained for the Gloperf dataset. In both cases we obtain power-law behaviors in good agreement with those obtained for the PingER data sets (see Tab. 1).

Gloperf data set collected in the second semester of 1999 that was used in [10]. We have extracted a set of parameter values for each of 650 unique site pairs in the sample and analyzed the statistics. These results are also reported in Table 1. Although the averages depend on the specific characteristics of the sample (size, world region etc.) and differ significantly from the PingER case, the existence of power law tails and the values of the exponents seem to be confirmed. These exponents can thus be considered as one of the few and sought after reliable and invariant properties of the Internet [14].

Finally, further evidence of large fluctuations in Internet performance is provided by the analysis of the packet loss data. Also in this case we are interested in the probability $P(r)$ that a certain rate r of packet loss occur on

Table 1. The table shows the improving performance over the years of the PingER data sample. As an independent check, we report the values obtained from the analysis of the data sample of the Gloperf project (second semester 1999).

Data set	α_{min}	α_{av}	$\langle\tau_{min}\rangle$	$\langle\tau_{av}\rangle$
April '00	2.7 ± 0.2	2.2 ± 0.2	3.7	6.6
Feb. '01	2.9 ± 0.2	2.4 ± 0.2	3.6	6.6
Feb. '02	3.0 ± 0.2	2.5 ± 0.2	3.1	5.3
Gloperf	2.7 ± 0.2	2.4 ± 0.2	5.4	7.8




Fig. 3. Probability density $P(r)$ for the occurrence of packet loss rate r on beacon-target pairs transmissions. The zero on the x axis corresponds to a 1% rate in packet loss. Note that the distribution has a linear behavior in the double logarithmic scale, indicating a power law behavior. The reference line has a slope -1.2 .

any given pair. We have analyzed the monthly average packet loss rate between PingER beacon-target pairs. In Figure 3 we report the probability $P(r)$ as a function of r . The plot shows an algebraically decaying distribution that can be well approximated by a power-law behavior $P(r) \sim r^{-\gamma}$ with $\gamma = 1.2 \pm 0.2$. The slowly decaying probability of large packet loss rate is another signature of the very heterogeneous performance of the Internet.

The results presented here have implications for the evaluation of performance trends. Models for primary performance factors must include the high heterogeneities observed in real data. Time and scale extrapolation for Internet performances can be seriously flawed by considering just the average properties. It is likely that we will observe in the future an improvement of the average end-to-end delay due to increased bandwidth and router speed, but the real improvement of the Internet as a whole would correspond in reducing the huge statistical fluctuations

observed nowadays. On a more theoretical side, the explanation and formulation of microscopic models at the origin of the scale-free behavior of Internet performance appear challenging, to say the least.

We thank C. Lee for sending us the Gloperf raw data. We are grateful to L. Carbone, F. Coccetti, L. Cottrell, P. Dini, Y. Moreno, R. Pastor-Satorras and A. Vázquez for helpful comments and discussions. This work has been supported by the Internetwork Performance Measurement (IPM) project of Istituto Nazionale di Fisica Nucleare, and the European Commission - Fet Open Project COSIN IST-2001-33555.

References

1. M. Faloutsos, P. Faloutsos, C. Faloutsos, ACM SIGCOMM '99, Comput. Commun. Rev. **29**, 251 (1999)
2. R. Govindan, H. Tangmunarunkit, in *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel (Ieee, Piscataway, N.J. 2000)
3. A. Broido, K.C. Claffy, in *SPIE International symposium on Convergence of IT and Communication*, (Denver, CO, 2001)
4. R. Pastor-Satorras, A. Vázquez, A. Vespignani, Phys. Rev. Lett. **87**, 2587011 (2001); A. Vázquez, R. Pastor-Satorras, A. Vespignani, Phys. Rev. E **65**, 066130 (2002)
5. W. Willinger, R. Govindan, S. Jamin, V. Paxson, S. Shenker, Proc. Nat. Acad. Sci. **99**, 2573 (2002)
6. A.-L. Barabási, R. Albert, Rev. Mod. Phys. **74**, 47 (2002)
7. M. Crovella, C. Lindemann, M. Reiser, Performance Evaluation **42**, 91 (2000)
8. W. Willinger, M.S. Taqqu, A. Erramilli, in *Stochastic Networks: Theory and applications*, edited by F.P. Kelly, S. Zachary, I. Ziedins (Clarendon Press, Oxford, 1996), pp. 339–366
9. V. Paxson, IEEE ACM T. Network **5**, 601-615, (1997)
10. C. Lee, J. Stepanek, *On future global grid communication performance*, 10th IEEE Heterogeneous Computing Workshop, May 2001
11. B. Huffaker, M. Fomenkov, D. Moore, E. Nemeth, K. Claffy, http://www.caida.org/outreach/papers/2000/asia_paper/
12. B. Huffaker, M. Fomenkov, D. Moore, K. Claffy, *Proceedings of the PAM 2001 Conference, Amsterdam, 23-24 April 2001*
13. C. Bovy, H.T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, P. Van Mieghem, *Proceedings of the PAM 2002 Conference, Fort Collins, Colorado, 25-26 March 2002*
14. S. Floyd, V. Paxson, IEEE/ACM Transactions on Networking **9**, 392 (2001)