

Self-Organization of Internet Paths

T. Kleiberg and P. Van Mieghem

Faculty of Electrical Engineering Mathematics and Computer Science
Delft University of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands
{T.J.Kleiberg, P.F.A.VanMieghem}@tudelft.nl *

Abstract The Internet consists of a constantly evolving complex hierarchical architecture where routers are grouped into autonomous systems (ASes) that interconnect to provide global connectivity. Routing is generally performed in a decentralized fashion, where each router determines the route to the destination based on the information gathered from neighboring routers. Consequently, the impact of a route update broadcasted by one router may affect many other routers, causing an avalanche of update messages broadcasted throughout the network. In this paper we analyze an extensive dataset with measurements on Internet routes between a set of highly stable testboxes for a period of five years. The measurements provide insight into the coherence between routing events in the Internet and we argue that the routing dynamics exhibit self-organized criticality (SOC). The SOC property provides an explanation for the power-law behavior that we observe in the operational times of routes.

1 Introduction

Interactive services in the Internet place strict bounds on the performance of end-to-end paths. Packet delay, delay variations and packet-loss have a severe impact on the quality of the Internet service and therefore it is important that end-to-end communication is reliable and predictable. The ability to control the end-to-end performance is seriously complicated by the connectionless nature of the Internet Protocol and the lack of any widely deployed Quality-of-Service implementation. As a consequence, the packets in the Internet are exposed to erratic network performance due to traffic fluctuations and routing dynamics. Traffic fluctuations can lead to temporary congestion of the router buffers, causing delay variations between the packets and packet-loss. Although congestion occurs very frequent in the Internet, measurements indicate that the traffic fluctuations are highly transient and the impact on the service performance often remains within bounds [25]. Routing dynamics correspond to the process where routing messages are propagated between sets of routers to advertise a route change. When a network event, for example a link or node failure, causes a route change, the network temporarily resides in a transient state while the routing tables of other

* The work is funded by the Next Generation Infrastructures foundation as part of the “Robustness and optimization of complex networks” project. The authors wish to acknowledge RIPE for providing the data and Fernando Kuipers, Steve Uhlig and Henk Uijterwaal for the valuable discussions.

routers are being updated. Routing dynamics contribute to most prolonged path disruptions and can last as long as 10 minutes, leading to serious degradation of Internet services [18, 20, 25, 26, 29].

In this work we study the dynamics of Internet paths with the use of an extensive dataset of traceroute measurements. In particular, we analyze how many routes are used between two end-hosts and how long a route remains operational. We regard the Internet as a “black box” and consider the path dynamics as the result of a collective behavior that organizes the thousands of autonomous nodes into a single complex system. Routing dynamics correspond to perturbations in the Internet and we argue that the statistical properties of the measured perturbations hint towards self-organized critical (SOC) behavior in the Internet. Self-organized criticality is often found in nature and other complex systems and arises as a collective result of the interaction between many autonomous sub-systems. When routers are considered as autonomous (sub-)systems that communicate via route-updates, the routing plane in the Internet can be considered as a SOC system. By comparing the characteristic features of SOC systems we argue that routing in the Internet also exhibits SOC. The presence of the SOC mechanism in the Internet implies that routing in the Internet is unpredictable in the sense that routes can change unexpectedly. The inter-event time between two routing events has no typical value and is widely varying. Packets associated with the same stream may follow entirely different routes, introducing a wide spread between the delivered packets and large delay variations. For the increasing number of real-time applications, such as interactive gaming, IP telephony, video and others, these variations can lead to dramatic degradation of the experienced quality. Furthermore, we find that permanent changes in the Internet cause the breakdown of existing routes and the discovery of new ones. On average, the number of discovered routes increases linearly in time at a fixed rate.

This paper is organized as follows: in Section 2 we briefly introduce self-organized criticality and present some features that are typical for SOC systems. Next, the measurements are described in Section 3. Section 4 contains the observations, followed by a discussion in Section 5. Section 6 presents an overview of related work. Finally, Section 7 presents the conclusions.

2 Self-Organized Criticality

Self-organized criticality was introduced by Bak *et al.* [3] as a property of a system that, through a self-organized process, always evolves to a “critical state”, regardless of the initial state of the system. A common feature observed in SOC systems is the power-law temporal or spatial correlations that can extend over several decades. SOC systems organize into clusters, with a scale-free spatial distribution, with minimally stable, critical states. A perturbation in that system in a critical state can propagate through the system at any length scale from a local change to an avalanche by upsetting the minimally stable clusters. The magnitude of the perturbations is only limited by the size of the system. The lack of a characteristic length leads directly to the lack of a characteristic time for the resulting fluctuations. Hence, a power-law distribution arises for the lifetime distribution of the fluctuations in the system. From the inter-event time of the fluctuations a time signal can be constructed, where the perturbations are modeled

as a series of Dirac pulses [12],

$$I(t) = \sum_k \delta(t - t_k) \quad (1)$$

where t_k corresponds to the time of the k -th event. The time signal can now be transformed into the frequency domain and the power-spectrum of (1) is found as,

$$S(f) = \lim_{T \rightarrow \infty} \left\langle \frac{2}{T} \left| \sum_{k=k_{min}}^{k_{max}} e^{-j2\pi f t_k} \right|^2 \right\rangle \quad (2)$$

where T denotes the whole observation time, k_{min} and k_{max} are the minimal and maximal values of the index k in the interval of observation and the brackets $\langle \dots \rangle$ denote the averaging over realizations of the process. The averaging is necessary since we are only interested in the process that leads to the power-spectrum and not a realization of the process. From (2) it follows that the estimation of the spectrum improves as the observation time increases. It can be shown that a power-law distribution in the inter-event time leads to a power-law spectral density [11],

$$S(f) \propto 1/f^\beta \quad (3)$$

where β is typically close to 1. The power-law spectrum in (3) is referred to as $1/f$ noise and is widely found in nature. The $1/f$ noise phenomenon is often observed in large systems that act together in some connected way and can be seen as a measure of the complexity of the system. The noise can arise as the result of the coherence between events in the system, the so-called long-range-dependence. It is also seen as a naturally emergent phenomenon of the SOC mechanism [3, 4, 28].

3 Measurements

3.1 Methodology

The measurement apparatus consists of a set of testboxes that are deployed by RIPE as part of the Test Traffic Measurements service (TTM) project¹. A testbox measures the Internet paths towards a set of pre-determined destinations by repeatedly probing the router-level path from source to destination. The path is measured by the traceroute tool, which sends probes to the destination host and infers the forwarding path by analyzing the response from the intermediate hosts. The traceroute messages are transmitted at exponentially distributed random intervals, with an average of 10 messages per hour from one source to one destination. Besides the IP path, the AS path is obtained by inspection of BGP data and matching each address in the IP path with an AS prefix². In the translation from the IP route to the AS path, the duplicate AS entries are removed,

¹ A detailed technical description of the design and features of the TTM testboxes can be found in [8] and on the TTM website, <http://www.ripe.net/projects/ttm/>.

² The AS path information has not been recorded in the initial phase of the project and is available only from the beginning of 2003.

which result from the multiple IP hops in one AS. Hence, the AS path consists of a list of unique AS numbers.

Each source maintains its own list of destinations, which is a subset of the other testboxes deployed by RIPE. The testboxes are placed at customers' sites, typically ISPs residing in various countries, just behind their border routers. Since the enrollment of the TTM project in 1999, the number of testboxes has increased from approximately 30 up to around 150 active boxes, today. Between the roughly 160 testboxes that exist, or have existed, around 10,000 source-destination pairs have been registered, where pair (A,B) is different from (B,A). Hence, the data that is available from 1999 does not include all the testboxes available today. In addition, testboxes can be temporarily offline for managerial or other purposes and several testboxes have disappeared completely, indicating the termination of the TTM service at the customer's site. The configuration and (geo)location of the testboxes is very stable. The IP address of a testbox seldom changes and only few testboxes are discontinued. The stability of the testboxes facilitates the measurement trustworthiness in the sense that the observations indeed reflect the network state and not so much the measurement setup. The high fidelity of the measurements and the extent of the observation time make the TTM measurements an excellent set to study long-term Internet path dynamics. In fact, it is the only publicly available dataset containing router-level information for this time span with such high accuracy.

3.2 Dataset

Section 2 emphasizes the importance that the observation time is long with respect to the interval times between subsequent events. On the other hand, increasing the observation time reduces the number of usable source-destination pairs, because less testboxes were available in the early stage of the project. Furthermore, measurements between source-destination pairs can fail: the list of destinations of each testbox can change over time and testboxes can be temporarily offline. To decrease the influence of these dynamics in the analysis, the set of usable source-destination pairs is restricted by the maximum time a source-destination pair was inactive. Increasing the stringency on the outage restrictions will reduce the number of usable source-destination pairs. Hence, a trade-off is made between the number of usable source-destination pairs versus the observation time and outage restrictions. The resulting dataset consists of the traceroutes between all the source-destination pairs that were active the entire period from January 1, 2003 until January 1, 2008, where any outage between a source-destination pair is restricted to maximally 28 days. Source-destinations pairs of which both source and destination belong to the same AS are excluded from the dataset. The dataset, that we will denote by \mathcal{D} , contains 64 source-destination pairs out of a set of 10 testboxes located within 8 different European countries.

3.3 Measurement artifacts

The anonymous and dynamic nature of the Internet inherently adds noise to the measurements which consequently incurs errors in the analysis. These measurement artifacts include persistent forwarding loops, transient routing loops, infrastructure fail-

Route result	probes	% probes	routes	% routes
Successful delivery	30986955	98.90	13496	43.19
Persistent forwarding loop	24732	0.07	1422	4.55
Transient routing loop	411	0.00	231	0.73
Infrastructure failure, destination not reached	36240	0.12	2561	8.20
Packet delivered at non-listed IP address	3995	0.01	619	1.98
Anonymous reply, destination reached	278066	0.88	12920	41.34

Table 1. Statistical overview of the pathologic routes and probes in \mathcal{D} .

ures and anonymous replies by the intermediate routers. Persistent forwarding loops are generally related to mis-configured routers, while transient forwarding loops are often a manifestation of routing dynamics. Infrastructure failures lead to premature termination of the traceroute probe. Anonymous replies are due to unresponsive routers or rejected probes. For a detailed discussion on these pathologies we refer to [18, 19]. We have also found several cases where the packet was not delivered at the correct destination address. This may be the result of erroneous routing or a configuration problem in the measurement setup such that the packet is delivered at an unknown IP address. Finally, we would like to mention the presence of “third-party” addresses as a source of noise in the traceroute measurements [9]. But since such occurrences are rare [9] we disregard them in our analysis. The classic traceroute tool developed by Van Jacobson and used in the TTM project is unable to detect such pathologies and different modifications have been developed to address short-comings of the classic traceroute tool [1, 7, 24]. These recent changes were not available in 1999 and are therefore not included in the TTM project. Routes which exhibit the above mentioned artifacts have been filtered from the dataset before processing. Table I presents an overview of the frequency of the measurement artifacts. From Table I, we can deduce that the pathologies contribute to slightly more than 1 percent of the measured probes. Hence, we argue that the measurements are barely affected by the pathologies.

3.4 Route fluttering

Between the successful routes, there is also a significant fraction of aliasing routes. Route aliasing can be a manifestation of load-balancing, where a group of packets that are traveling between the same source and destination traverse different routes. The packets are separated based on their packet header or simply in a round-robin fashion. As a result, the samples of the IP path in the presence of load-balancing routers will consist of rapidly alternating routes, so called fluttering routes [18]. Load-balancing is the result of a decision process inside one router, it does not involve the advertisement of any routing updates to neighboring routers and does not affect the state of the routing tables³. Hence, load balancing does not contribute to the routing dynamics and we will handle fluttering routes as a single route, i.e. as if the packets were sent along one route.

³ Here we assume that the routing tables are not affected by the actual traffic due to some form of traffic engineering.

To identify fluttering routes, we will adopt the heuristics presented by Paxson [18]: two routes are considered the same when the paths have an equal length and differ at maximally one consecutive hop. When routes are considered the same route, the samples of all the routes are aggregated as if they were samples from one route. E.g., if route R_1 is observed 1000 times and route R_2 is observed 500 times, then the aggregated route has 1500 observations. Prior to filtering the fluttering routes, the dataset \mathcal{D} contained 13496 routes. Afterwards 8612 routes remained.

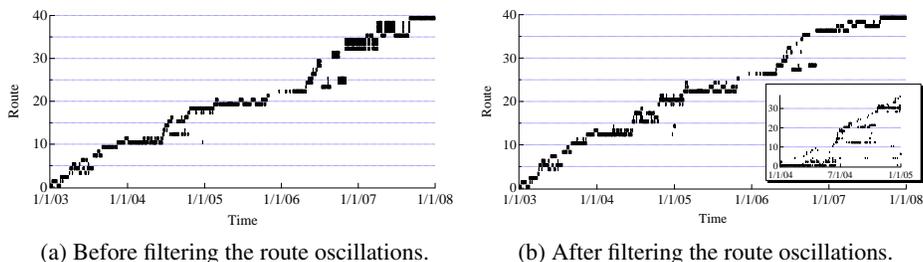
3.5 Metrics

The routing dynamics in the Internet can be measured by means of the time intervals between subsequent route events and their coherence. A route event can affect one or more routers on the route between a source–destination pair, such that the route is changed. Hence, the time-interval between two route events corresponds to the time that a route is operational without being interrupted, which we will denote by the *route duration*. In the RIPE measurement setup, the path between a source–destination pair is sampled at independent, exponentially distributed random intervals with an average of 360 seconds. The exact time of the traceroute call is rounded to seconds and recorded in the database. When the same path is sampled multiple, consecutive, times, only the time of the first and last call are recorded. Hence, there is no accurate information of the exact time of each call, only the number of calls and the start and end time of the sequence of calls. The number of calls qualifies as an alternative measure for the duration of the path, hence the route duration is defined as the number of successive occurrences once it is selected. Due to the Poisson measurement times, the PASTA property applies and the sampled time averages indeed reflect the real time averages [18]. The sampling rate prevents us from detecting the typically highly transient failures at the data plane due to congestion, which are typically in the order of seconds. Yet, the average inter-arrival time is sufficiently small to detect the slow route dynamics, which can last many minutes.

4 Observations

Figure 1 shows the IP routes between a typical source–destination pair, considered over a long period of five years. Only the 40 most dominant routes are displayed. Figure 1 depicts when routes are operational and demonstrates the presence of several phenomena in the Internet, which we will discuss here. First, Figure 1 exemplifies that route fluttering is a common artifact in the Internet and it is important to consider these oscillations in the analysis. Figure 1a shows several cases of route fluttering, e.g. at the end of 2006 equivalent routes appear that overlap in time. After resolving the equivalent IP paths, the fluttering routes have been merged, as illustrated in Figure 1b. The routes presented in Figure 1b are considered unique routes and a route change indeed corresponds to a change in the routing table and not to load-balancing.

Second, Figure 1 illustrates that route events occur frequently and at all time scales. Most of the time a single route prevails between a source–destination pair. The dominant route is sometimes interrupted, where the interruption can be very brief or some-



(a) Before filtering the route oscillations. (b) After filtering the route oscillations.

Figure 1. Example of a set of observed routes between one source–destination pair from the dataset. The horizontal axis represents the time axis, ranging from January 1, 2003 to January 1, 2008. The vertical axis indicates the different routes between the source–destination pair. The marked areas indicate when a particular route is operational. Only the 40 most prevalent routes have been shown out of a total of 213 routes for (a) and 158 for (b). The routes are sorted from bottom to top in order of their first appearance. The inset in (b) shows *all* the observed routes for the year 2004.

times last for days. The inset in Figure 1b shows all the routes between this source–destination pair for the whole 2004. The inset reflects that many routes are observed only occasionally and briefly. These routes can be the result of temporal route failures or routing dynamics.

Finally, we can conclude from Figure 1 that the routes have a limited lifetime. Between a route’s first and last appearance many other routes can be operational, however, in all cases the route eventually disappears and is never seen again. The restriction on these lifetimes is a consequence of the evolution taking place in the Internet, that is stimulated by changes in peering relations, the birth and death of ASes and reconfigurations at the intranet level.

Figure 2 presents the measurements on the number of *unique* routes learned since January 1, 2003 till January 1 2008, on both the AS- and IP-level, averaged over the source–destination pairs. We consider an IP route *unique* when there exists no other route with the same sequence of IP addresses. Similarly, we consider an AS path *unique* when there exists no other AS path with the same sequence of AS numbers.

The measurements in Figure 2 exhibit a remarkable linear behavior on both the AS and IP level, which is in agreement with the findings in [17]. If $R(t)$ represents the number of routes that are learned as a function of the time t , then according to Figure 2 we can write $R(t) = \alpha t - g(t)$, where $g(t) = o(t)$ for large t . Hence, $\lim_{t \rightarrow \infty} R(t)/t = \alpha - \lim_{t \rightarrow \infty} g(t)/t = \alpha$, which corresponds to the “rate” at which new routes are discovered. Figure 2 shows that the discovery rate remains fairly constant for the entire observation period. A new IP route is learned approximately every 14 to 15 days, on average. Note that this is not the same as the average duration: the route that is actually used still frequently changes back to a previous route. The first few months of 2003 the discovery rate is slightly higher due to a learning phase. At the AS-level, a new route is discovered every 38 to 39 days, on average.

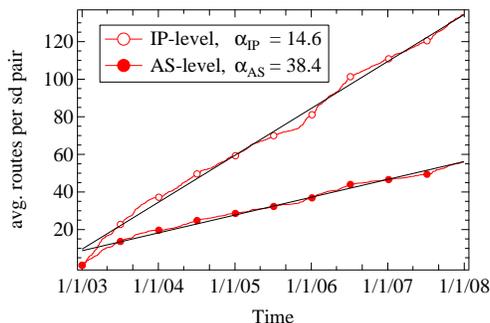


Figure 2. The average number of discovered routes between the source–destination pairs at both the AS- and IP-level, counting from January 1, 2003. The measurements have been fitted with a line. The fit at the IP-level resulted in a discovery rate of 14.6 days per route, while at the AS-level the fit provided 38.4 days per route.

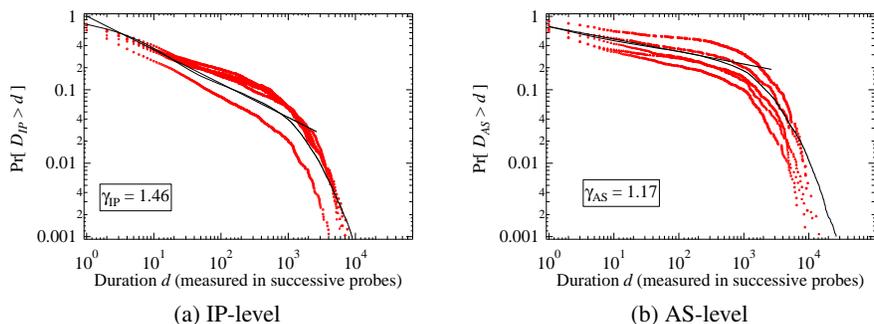


Figure 3. The complementary cumulative distribution function of the route duration on log-log scale. The markers represent distributions for five individual source–destination pairs with typical behavior. The measurements are aggregated into a single distribution (solid line) and fitted with $F(d) = Cd^{1-\gamma}$. The fit results are presented as the straight lines; the power-law coefficients are found as $\gamma_{IP} = 1.46$ and $\gamma_{AS} = 1.17$.

Figure 3 shows the complementary cumulative distribution function (CCDF) of the duration of a route at the IP- and AS-level. The duration is measured as the number of successive occurrences of the same route until a new route becomes operational or the routing fails and an error message is received. We have selected five source–destination pairs that do not have any end-points in common and for which the routes show typical behavior. For each pair we have computed the route duration and plotted the CCDFs in Figure 3. Figure 3 demonstrates that the shape of the distributions in both the IP-level and AS-level cases, are very similar. The requirement that the source–destination pairs do not have any node in common argues that the behavior that we observe in Figure 3 is a true feature of the Internet and not an artifact of the dataset. The aggregate result of the entire dataset, which is obtained by accumulating the results for all the source–destination pairs, is presented by the solid line. In both cases, IP and AS, the aggregate distributions strongly resemble that of the individual pairs, which indicates

that the mean reflects a real property and can be considered as the average behavior of any source–destination pair in \mathcal{D} .

Figure 3 illustrates that the aggregate distributions of the duration relatively closely follow a power-law for the first three decades. At approximately $t = 10^3$ the distribution is cutoff followed by a steep decline. The CCDF of the aggregate result has been fitted with a power-law and is presented by the straight lines. The power-law exponent of the fits are $\gamma_{IP} \approx 1.46$ and $\gamma_{AS} \approx 1.17$. Power-law distributions with an exponent $\gamma < 2$ exhibit extreme behavior and have a divergent mean [15]. When sampling a power-law with extreme behavior, the mean is determined by the sample with the highest value, which will go to infinity when the number of samples becomes large. In real-world systems the mean is finite: the distribution is cut off in the tail because the system has a limited size. The measurements on the route durations are restricted by the limited sample space of the Internet (we cannot sample the entire Internet) and possibly the limitations of the measurement architecture (e.g. the limited uptime of testboxes due to managerial purposes, etc.) Compared to the IP routes, the power-law exponent at the AS-level is smaller, yielding a fatter tail. At the AS-level there are more long-lasting routes, which can be explained by considering that a route change at the IP level does not necessarily lead to a different AS path. This observation agrees with our finding from Figure 2. The discovery rate of AS paths is smaller than that of the IP routes, implying that multiple IP routes exists per AS path. The average duration of an AS path must therefore be greater than that of the IP routes.

Finally we will compute the spectral density of the routing dynamics and examine the existence of $1/f$ noise. The time signal (1) requires the time of occurrence of the routing events, which can be constructed from the route duration measurements by,

$$t_k = \sum_{i=0}^{k-1} D_i \quad (4)$$

where D_i is the duration of the route after i route changes. We have computed $I(t)$ separately for all the source–destination pairs. The power-spectrum $S(f)$ is then computed by averaging the transformed signal between the source–destination pairs. The result is presented in Figure 4, which demonstrates the presence of $1/f$ noise.

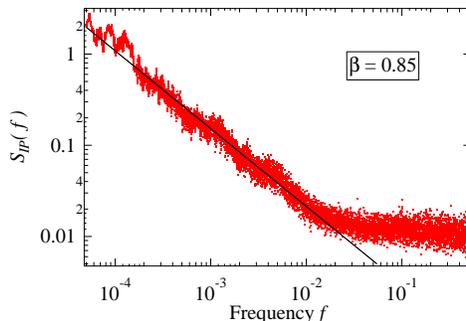


Figure 4. Power spectrum $S(f)$ of the inter-event time-signal $I(t)$ printed on log-log scale.

5 Discussion

The extreme power-law behavior observed in Figures 3a and 3b is often seen in real world systems [15]. When a distribution possesses such a heavy tail, the expected value and the variance tend to infinity. In practice this implies that the observed measure is highly unpredictable. In our case it means that the route is unpredictable and packets associated with the same stream may follow (many) different routes. At first glance, the unpredictable nature of the routes seems remarkable. The packets in the Internet are routed by the connectionless Internet Protocol, where each packet is routed individually without establishing a connection prior to the transmission. However, most of the lower layers, the datalink and physical layer, use connection-oriented technologies, such that the route towards the destination is known in advance.

The changes of the topology are relatively sparse and slow, because it is often manually managed. Hence, one would expect more stable paths. Yet, the measurements indicate that route changes occur at all time scales. Such frequent changes can have negative impact on streaming services that rely on packets arriving on time and in the correct order [23,26]. At the same time, the high variability demonstrates that the Internet is resilient and fastly adapts to changes.

The power-law spectral density and power-law behavior of the inter-event time can be seen as a manifestation of self-organized criticality. Bak *et al.* [2, 3] argue that SOC naturally arises in interactive dynamical systems with many degrees of freedom. The Internet is clearly a dynamical system with self-organizing behavior. The notion of SOC is a plausible explanation of the observed dynamics.

6 Related work

There has been considerable work devoted to Internet path measurement in several projects, such as Skitter⁴ and Rocketfuel⁵. The Skitter project, which also measures IP path information, was initiated around the same time as the RIPE TTM project, however its goal is Internet topology research which it does by querying roughly 400,000 destinations several times per day using only 20 source nodes. The Rocketfuel project combines BGP data with traceroute measurements to infer the ISP topologies [22]. To our knowledge, the database from RIPE is the only database that actually has IP-level measurements for a prolonged period (more than 9 years) at a relatively high sample rate between fixed and stable testboxes.

Several works have been published routing dynamics in the Internet, but none of them have examined the long-term correlations between routing events. Early work by Paxson [18] studies the stability of Internet paths through traceroute measurements performed during several months. Labovitz *et al.* [13] studies the path stability by combining IP and AS information. Iannaccone *et al.* [10] and Markopoulou *et al.* [14] study the path properties by monitoring the route update-messages within an AS and conclude that a small fraction of the links contributes to a large fraction of the route updates. Pucha *et al.* [20] and Wang *et al.* [26] study the impact of route changes on

⁴ <http://www.caida.org/tools/measurement/skitter/>

⁵ <http://www.cs.washington.edu/research/networking/rocketfuel/>

the path performance w.r.t. packet delay and jitter. Our work differs from the previous works in that the dataset that we use extends over a period of five years, which exposes long-term effects. Furthermore, we associate the routing dynamics with self-organized criticality and argue that routing is unpredictable leading to large variations in the path performance. Several works have related the SOC mechanism and $1/f$ -noise to the Internet. The self-similarity and long-range dependence of traffic patterns often reported in the Internet are considered related to the $1/f$ noise phenomenon [6, 27]. Csabai [5] measured the round trip times of packets and showed that the correlation between the round-trip times produces $1/f$ noise in the power spectrum. Ohira *et al.* [16] and Solé *et al.* [21] demonstrated that computer networks with self-organizing behavior show the maximum information transfer and efficiency at the critical state. In addition, Solé *et al.* demonstrate that near criticality, the network performance shows the highest variability in terms of packet latency. In this work we study the inter-event times of route changes in the Internet and argue that unpredictability and instability of Internet routes may be related to SOC.

7 Conclusions

Through analysis of traceroute measurements we have studied the lifetime of routes and the route dynamics. Based on the observations from our dataset we find that routing in the Internet is highly dynamic and results in unpredictable route durations. The actual cause of the perturbations in the Internet is hard to retrieve, since it is often related to configuration changes within and between ASes. The extreme power-law behavior suggests that the Internet exhibits self-organized criticality. The power spectrum obtained from the inter-event time of route changes confirms our conjectures. The impact of SOC at performance of applications and services on the Internet remains difficult to evaluate and requires further investigation. The unpredictability introduced by SOC may lead to quality degradation of Internet services, in particular services that heavily depend on a timely delivery of the packets. At the same time, the adaptability of the Internet paths demonstrates the resilience against failures and attacks.

References

1. B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with Paris traceroute. In *Internet Measurement Conference (IMC)*, pages 153–158, Oct. 2006. ACM.
2. P. Bak. *How Nature Works*. Copernicus, Springer-Verlag New York Inc., New York, NY, USA, 1996.
3. P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Physical Review Letters*, 59(4):381–384, July 1987.
4. K. Christensen, Z. Olami, and P. Bak. Deterministic $1/f$ noise in nonconservative models of self-organized criticality. *Physical Review Letters*, 68(16):2417–2420, Apr. 1992.
5. I. Csabai. $1/f$ noise in computer network traffic. *Journal of Physics A: Mathematical and General*, 27(12):L417–L421, 1994.
6. A. J. Field, U. Harder, and P. G. Harrison. Measurement and modelling of self-similar traffic in computer networks. *IEE Proceedings Communications*, 151(4):355–363, Aug. 2004.

7. E. Gavron. NANOG traceroute, 1995. for latest version, see <ftp://ftp.login.com/pub/software/traceroute/>.
8. F. Georgatos, F. Gruber, D. Karrenberg, M. Santcroos, A. Susanj, H. Uijterwaal, and R. Wilhelm. Providing active measurements as a regular service for ISPs. In *Passive and Active Measurement Conference PAM*. Springer, 2001.
9. Y. Hyun, A. Broido, and kc Claffy. On third-party addresses in traceroute paths. In *Passive and Active Measurement Workshop*, Apr. 2003.
10. G. Iannaccone, C.-N. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot. Analysis of link failures in an IP backbone. In *Internet Measurement Conference (IMC)*, pages 237–242, Nov. 2002. ACM.
11. B. Kaulakys, V. Gontis, and M. Alaburda. Point process model of $1/f$ noise vs a sum of Lorentzians. *Physical Review E*, 71(5):051105, 2005.
12. B. Kaulakys and T. Meškauskas. Modeling $1/f$ noise. *Physical Review E*, 58(6):7013–7019, Dec. 1998.
13. C. Labovitz, A. Ahuja, and F. Jahanian. Experimental study of Internet stability and backbone failures. In *FCTS*, pages 278–285, June 1999. IEEE Computer Society.
14. A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot. Characterization of failures in an IP backbone network. In *INFOCOM*, Mar. 2004. IEEE.
15. M. E. J. Newman. Power laws, Pareto distributions and Zipf’s law. *Contemporary Physics*, 46(5):323–351, Sept. 2005.
16. T. Ohira and R. Sawatari. Phase transition in computer network traffic model. *Physical Review E*, 58(1):193–195, 1998.
17. R. V. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of Internet AS topology. In *SIGCOMM*, pages 313–324, Aug. 2007. ACM.
18. V. Paxson. End-to-end routing behavior in the Internet. *IEEE/ACM Transactions on Networking*, 5(5):601–615, Oct. 1997.
19. V. Paxson. *Measurements and Analysis of End-to-End Internet Dynamics*. PhD dissertation, University of California, Lawrence Berkeley National Laboratory, Apr. 1997.
20. H. Pucha, Y. Zhang, Z. M. Mao, and Y. C. Hu. Understanding network delay changes caused by routing events. In *SIGMETRICS*, June 2007. ACM.
21. R. V. Solé and S. Valverde. Information transfer and phase transitions in a model of Internet traffic. *Physica A*, 289(3):595–605, Jan. 2001.
22. N. T. Spring, R. Mahajan, D. Wetherall, and T. E. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2–16, Feb. 2004.
23. R. Teixeira and J. Rexford. Managing routing disruptions in Internet service provider networks. *IEEE Communications Magazine*, 44(3):160–165, Mar. 2006.
24. M. Toren. Tctraceroute, 2001. for latest version, see <http://michael.toren.net/code/tctraceroute/>.
25. F. Wang, N. Feamster, and L. Gao. Measuring the contributions of routing dynamics to prolonged end-to-end Internet path failures. In *GLOBECOM*. IEEE, Nov. 2007.
26. F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush. A measurement study on the impact of routing events on end-to-end Internet path performance. In *SIGCOMM*, pages 375–386, Sept. 2006. ACM.
27. W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, 5(1):71–86, Feb. 1997.
28. J. Yuan, Y. Ren, and X. Shan. Self-organized criticality in a computer network model. *Physical Review E*, 61(2):1067–1071, Feb. 2000.
29. Y. Zhang, V. Paxson, and S. Shenker. The stationarity of Internet path properties: Routing, loss, and throughput. ACIRI technical report, AT&T Centre for Internet Research at ICSI, 2000.