Internet Failures – An Emergent Sea of Complex Systems and Critical Design Errors?

Jon Crowcroft
The Computer Laboratory,
University of Cambridge
Cambridge, UK
jon.crowcroft@cl.cam.ac.uk

December 7, 2009

Abstract

Complex systems researchers have looked to the Internet as a possible source of interesting emergent behaviour. Indeed, some high profile failures, and some low level phenomena, might easily be construed as evidence of a complex system. In this paper, I look at the local and global consequences of the Internet design, and show that few, if any, of these problems are actually consequences of emergent properties in the pure technical sense. However, there are lessons for network architecture from these problems. The influence of local decisions on global behaviour of the network is a source of some of the difficulties that protocol designers must cope with, but it is also a source of great wealth and innovation, and as such should be regarded in a positive light.

1 Introduction

A number of high profile failures in the Internet have come to be regarded by outsiders as possible indications that the system is suffering from complex, emergent behaviour.

In this lecture, I will look at the design of the Internet from the perspective of these systemic failures, and determine the root cause. Specifically: firstly, we have seen local decisions on policies to do with access to certain sites (in particular, certain instances of user contributed youtube video content from a particular country the middle east) impact the ability of most of the world to reach anything at all on youtube; secondly, a small misconfiguration of a database led to the "zero-ing" of the root of the Domain Name System, which is used to provide human-friendly ways of naming resources (crucially, web sites) so that users don't have to type complicated numerical addresses, frequently erroneously. This led to wide-spread un-reachability; thirdly, small errors in configuration of the underlying access control list management system by operations staff at Google led to most sites being marked by their search engine as potentially "unsafe" when many of the sites were probably no less safe than average.

There are lessons to be drawn from these failures. Most important is that the Internet has evolved as a set of cooperating systems organised vertically over many layers, and horizontally over many organisations. Incremental growth depended on local arrangements, but the growth in value has been super-linear because these systems are globally reachable. This means that local decisions impact global behaviour. While this leads to critical dependence on some configurations being correct, it is not clear it could correctly be termed "complex" rather than complicated. A consequence of this is that we may need techniques to reduce this complicatedness, to reduce the chances of these errors. Combating the impact of these errors would lead to a very different Internet. Indeed, a new Internet Architecture that was less prone to these sorts of mis-configurations might enjoy certain other advantages (less prone to unsolicited communication such as Spam or Denial-of-Service attacks), but might never see any deployment. I will end the paper by speculating on the lesson for Future Internet Architectures and the basic rules we observe here that determine whether a new networked system can succeed or fail.

2 Complex Systems and Emergent Behaviour

Why should we care if the Internet is a complex system or not? Principally, we need to know what skill set is best brought to bear on solving design problems. The Internet is now a critical resource, like clean water, sewers, electricity, and roads, and as such, we need to be able to depend on it. It may never be possible to build a machine that is 100% reliable, but we need to know what to expect of a system so that we can plan contingencies. Complex systems have a tendency to surprise, and that is not pleasant in this context.

Complex systems are built out of a myriad of simple components which interact, and exhibit behaviour that is not a simple consequence of pairwise interactions, but rather, *emerges* from the combination of interactions at some scale. A famous example of emergent behaviour in nature is that of fireflies flashing in synchronisation. It arises because the fireflies see each other, and stop and think before flashing again – after a small number of pauses, they all end up flashing at the same frequency and phase.

Flocking is an even clearer example of emergent behaviour, where a pattern of apparently coordinated action by a set of birds is in fact a product of lots of continual small local adjustments in flight plans. Looking at a pair of fireflies, or a pair of birds, you would not see either of these phenomena.

Some emergent behaviour occurs because of non-linearity in the system. This makes the results even more unexpected and spectacular.

In engineering, a famous example is that of resonance. A 42 MPH wind set the Takoma Narrows road bridge oscillating violently at around 5Hz[1].

3 From Local to Global

One of the big claims about the Internet is that local changes can have a global impact and that this is evidence of emergent behaviour of a complex system.

Let's look at three famous recent examples of this, to see if there is any truth in this claim:

Routing Firstly, there was the incident when a small Internet Service Provider (ISP) in the middle east decided to configure its policy routes to disallow traffic to servers run by the Youtube service, since there was content there deemed offensive (and possibly illegal) in the ISP's home country[2].

The Border Gateway Protocol(BGP) is a sophisticated system used to coordinate paths in the Internet, across sequences of Autonomous Systems(AS), each of which is typically a separate ISP. The system employs a variant of the now more than 50 year old distributed Bellman-Ford distance vector algorithm, but expresses path preferences based on policies, rather than simple performance metrics. This has led to a number of subtle business relationships emerging in terms of how ISPs relate to each other (e.g. customer-provider, or peers, or siblings).

Essentially, the small provider configured a special entry in their border routers to claim that they were, as an ISP, the best route to the Youtube service, but then, discarded all traffic to that set of sites rather than actually then routing it on to the correct destination. For a while, the effect was to cause all ISPs in the world to treat this ISP as the best route, thus "black-holing" all clients of Youtube – i.e. no-one could get at their favourite User Contributed Video site any more. However, note that there is an authoritative database kept separately which maps from site address (technically, IP prefix), to AS, and ISPs are advised to check advertisements for reachability to these prefixes from neighbour ISPs to see if they are valid against this authoritative database. If you like, this is a "sanity test".

Thus, in this case there were two errors necessary[3]. The small middle eastern ISP did not need to advertise to the rest of the world about the fact it was deliberately, and quite reasonably, black-holing a site its customers did not like. On the other hand, the rest of the world should

[1] This can be seen online at `http://www.youtube.com/watch?v=3mclp9QmCGs`

[2] See `http://government.zdnet.com/?p=3673`. for a full diagnosis of the problem and some of its political consequences

[3] In fact, as pointed out by Richard Clayton, a third error was perpetrated by Youtube themselves since they chose to announce a /23 prefix but had two /24s in the routing registry. Thus a /24 looked valid, but a /25 would not have looked valid - the devil is often in the details!

not have believed this advertisement in any case, and rapidly, moved to a configuration that correctly filters the advertisement.

Despite claims that these sort of problems can be caused because of the complexity of BGP, leading to errors in configuration[10], the reality is that this sort of event was a rarity, and was rapidly corrected by collective cooperative behaviour by network operators. It is really just something that is complicated, rather than complex.

Naming A second incident that had a high impact on a large number of users was when the organisation that maintains the "root" Domain Name System (DNS) servers for the `.com` domain accidentally removed *all* the entries.[4] Other similar unintended effects have been caused by poorly managed secure DNS deployment.

While the Internet routing system does not depend on names, the World Wide Web does, since websites are identified typically by URLs embedded in documents, and these contain DNS names as a component. If you cannot look up a name, you cannot get the network (IP) address of the site, and so you are lost. This is not complex system behaviour, but is the result of over-centralisation and insufficient sanity checking – one would expect the database to change somewhat from day to day (indeed, there are surprisingly large numbers of new or altered entries) but one would not expect the database to shrink much, if at all.

Searching The third example of global impact of a local configuration change is the incident when a very small change to rules used by the search engine giant, Google, that describe sites that may contain "malware", led to all sites that google returned in any and all search results, being marked as potentially infected. For (overly?) cautious users, this would lead to them not following the results of any search at all, which would lead to severe lack of information, since many people use search rather than bookmarks, to find even sites they visit frequently[5]. Again, this is not complex, it is merely that a small modification to a database was not flagged as potentially an error.

All three examples above were largely reported in the technical press, and made headlines, but have not led to new research initiatives in complex systems, simply because they are just examples of bad practice in sanity checking. Indeed, there is a principle, expounded by the late Jon Postel, known as the robustness

[4] Verisign use a commercial database to store the boot files for the root server – apparently, it was only a matter of a few key strokes by an unskilled operator to zero the database (i.e. delete all the entries). If the DNS servers are then rebooted from this database, all lookups for *anything*.com will return an error "non existent domain". Recently, some of the servers at the root have taken to returning a page full of advertisements instead of returning the error – this has other unintended consequences, as discussed by DNS expert Paul Vixie in `http://queue.acm.org/detail.cfm?id=1647302`.

[5] This is described in some detail in this technical news article `http://news.cnet.com/8301-1001_3-10153942-92.html`.

principle, explicitly stated in the standards document on one of the core Internet Protocols (see RFC 761 on Transmission Control Protocol, in `http://tools.ietf.org/html/rfc761`), wherein it is writ "Be conservative in what you do; be liberal in what you accept from others.": this would suffice to solve these problems if applied appropriately.

Four other instances of alleged complexity relate to traffic patterns, and have led to some research, although most have been transient or rare occurrences which the Internet now avoids.

Gridlock In the late 1980s, there was a massive congestive collapse of the Internet due to the lack of any resource management system in the network, unlike previous networks such as the ARPANET. What emerged was borrowed from other networks (notably from DECNET, see `http://en.wikipedia.org/wiki/DECNET`), but the ideas were present in the original article by Donald Davies (see `http://en.wikipedia.org/wiki/Donald_Davies` which identified the idea of packet switched networks), and from control theory, and has led to an entire research industry in what is now known as "end-to-end congestion control". This eschews any resource management within the network itself, rather relying on end system cooperation. Some more recent work has suggested that one might improve the efficiency and fairness of the scheme with a modest amount of network help, e.g. [22], but the core idea remains the same – self-admission control on a packet by packet basis (each acknowledgement lets you send another packet and potentially increase your rate; each lack of acknowledgement gives you a hint to slow down).[6].

While the congestion collapse was an emergent property, its solution was the application of good practice in control theory, and made use of known techniques. Since then, researchers have tried to identify if there are other *phases* of the network, as well as "collapsed", or "flow controlled".

Phase Change While many people have tried to find evidence of *phase changes* in the overall Internet traffic time series, as far as I can find, no-one has actually succeeded.[7] There has also been work on traffic engineering[15], which shows the system is not as complex as some would like to believe.

Synchronisation The one case of an interesting and elegant exposition of an emergent behaviour is that of Routing Update Message Synchronisation[11]. This is undesirable as it leads to spikes of work for routers processing received messages, and potentially leads to loss of forwarding capability, or

[6] There are over 5000 papers on Google Scholar, and other bibliographic databases, on the topic of congestion control. Very little has changed in 20 years, however, since the basic idea of packet conservation and stable feedback control

[7] The thesis work of Dina Papagiannaki looked at traffic on the Sprint backbone in great detail `http://en.scientificcommons.org/konstantina_papagiannaki`, and that ISP allowed the author access to detailed traffic traces, some of which have now been published. While there are patterns, these are pretty much what you would expect from the demographics of the user base.

even loss of control messages which would lead to faulty results for the route computation. The solution (again, a well understood technique in distributed computing systems work) is to randomise the timers for the route update process – this is something systems like Ethernet protocols (small CSMA/CD) have done since the mid 1970s.

Fractal Internet traffic is “self-similar”[17]. However, this is not an emergent property, but simply reflects the nature of the users and the use. The superposition of a set of on-off processes accessing files, where the “on” part of the process is distributed sub-exponentially, is sufficient to explain the behaviour. This can in turn be explained by users downloading web pages, then reading/thinking about the web page rendered in their browser window. The typical file distribution and popularity distribution of files in the Internet is well known to be Zipf[9] – this is no different than library books for the last 1000 years. That’s the necessary and sufficient reason for the self-similarity

There have been attempts to explain the self-similarity as an emergent property, arising from the fractal nature of TCP, but while this may be somewhat true, and the jury is still out[21]: it is not necessary to invoke such a complex explanation for the heavy tailed distribution, when a simpler one is available and suffices.

Failures in the the Internet may have quite an impact, but they don’t seem to indicate complex system characteristics rather than good old fashioned configuration mistakes which can be controlled by relatively simple sanity checks. Where there is some evidence of complexity impacting the global system on the short time frame, it is quite rare.

In the next section, I will look at some emergent properties that seem to impact the system on the long time frame, leading to economic and social structures that are novel and interesting.

4 From Global to Local

In the previous section, I looked at the short time scale, and tried to see, without success, if there were local incidents that could lead to macroscopic behaviour in the Internet which might be considered emergent. In this section, I look at the longer timescale, and show that there are properties of the network that seem to derive from the underlying structures, but are surprising. In each case, I try to identify a principal lesson that can be extracted from these emergent social and business practises.

ISP Business Relationships One of the emergent properties of the net that I don’t believe was fully intended as a consequence of its original design, but has turned out to be incredibly rich, is the range of business relationships between ISPs[19] induced by the properties of the policy routing system, as embodied in BGP. This has become an entire industry of research in and of itself, leading to many papers about the AS topology of the Internet. The graph that is formed by the ISPs in customer-provider and peering relationships, and its evolution[14], is very interesting – it is not *optimal* in the sense of being designed for some single traffic provisioning goal, but it has led to a very competetive and innovative world in terms of different styles and types of Internet Service, some offering "walled gardens", others being mere ultra-fast bit-pipes, and others being a mix.

Application Layer Networking Most of the recent innovation in the Internet has been in the application layer, rather than in the network layer or transport protocols. Web 2.0[8] has led to many new systems being devised such as User Generated Content sites, tagging and reputation/recommendation systems, social networking technologies and so forth. One of the results of the way that Web 2.0 uses the Hypertext Transfer Protocol (HTTP) is that mobile applications on smart phones (iPhones, Android Phones, Blackberries, Palm Pres, and pretty much any other system now) make mobility almost seamless, even though attempts to make *network layer mobility* work seamlessly with various mobile IP routing schemes and handovers, have all dismally failed. The Internet technologists have *routed around* the barriers (in this case, the barrier being the inability to get the core standard Internet Protocol to change at the router technology level).

But before this happened, the pressure on the network to provide other services, such as *multicast*, or group distribution of content had failed for a variety of reasons. Instead, application layer Content Distribution Networking flourished, both as a full-on commercial service CDN as exemplified by Akamai[9][16] and Youtube[7], and in the trenches, as exemplified by BitTorrent[6]. The interesting thing about these technologies is that they both capture business relationships in subtle ways. Akamai distribute content for very large content creation service providers, including Microsoft Windows Update, and CNN and the BBC. BitTorrent is widely used for legal open source software distribution (as well as for breaking copyright) and incorporates a neat tit-for-tat incentive matching scheme to reduce the amount of free riding that plagues some file sharing systems. This latter scheme is a classic result from game theory, as any computer scientist will recognize.

Social Networking As mentioned above, Social Networks have emerged from Web 2.0 technology, and represent a mix of new business models[7] [8] [4]. User Contributed Content rises up from the human need to give, and the success of Gift Economy has been studied by Anthropologists for over 100 years. For content which actually has any production cost, the gift system originally arose in peer-to-peer file sharing and often from copyright theft, but has now successfully evolved into systems that either thrive on free

[8] Web 2.0 `http://en.wikipedia.org/wiki/Web_2.0`

[9] Akamai, Digg

content with advertising revenue, or segue smoothly into new subscription systems (with free samples).

Market Model The Internet business community has been quick to understand the value of combining the freely indexable content, and the interest in it. Such implicit or explicit viral marketing[18] is well understood in the commercial world, and has other interesting applications, such as the optimisation of CDNs[20].

New Privacy While the new content networks have capitalised on visibility of content and users' interest, this has led to concerns about privacy. A new research area has arisen that tries to apply information theory and security systems architecture thinking to the problem of how to retain the market, but improve privacy[2][13][1]. Some go as far as not trusting the content owner or CDN, and propose decentralising the entire system [5].

In the set of success stories above, we can extract a number of principles being applied (either by conscious design, or through natural evolution). I would claim that Postel's Robustness principle is present in the Web design. The way that application layer networking has routed around (worked around, in Internet community language) the ossified Internet infrastructure protocols is simply an example of the power of indirection. Tunnelling protocols over other protocols is of course just the recursive application of layering.

5 Global to Global

I'm indebted to Mark Handley for pointing out that there has been one near catastrophe in the Internet caused by a complex system interaction, and that was the incident of the slammer worm, which is well dissected at `http://www.spamlaws.com/slammer-worm.html`[10].

The worm attacks end systems and is an example of a scanning system that propagates very rapidly exploiting code weaknesses in end-systems, together with lists of addresses to visit next.

The problem that showed up in this particular case was that the end systems were often routers (normally one considers a router as an intermediate system, but when traffic is directed to the router itself, it should be thought of as an end system). Some routers were so inundated with traffic that they were unable to compute and send routing updates to their neighbours. The net effect of this is that they are "dropped" from routing tables, and so connectivity starts to fail. This interaction between end system role and intermediate system role (control plane and data plane) shows up as an emergent property, which is that the network starts to partition.

[10] Similar problems arose with the Welchia worm, which deployed a lot of ICMP traffic - the result was for ISPs to start blocking such packets

As in other plagues such as Ebola, which kill their victims too fast, the epidemic may be self limiting, but this very much depends on the relative timing of the attack, failure and route update periodicity.

The Internet's vulnerability to such Distributed Denial of Service (DDoS) attacks is well known and a matter of concern and cost for defence of end system services, but this interaction between such attacks and the routing and forwarding substrate is even more worrying, and certainly is a good reason to re-think some aspects of the Internet Architecture in response.

6 Conclusions

The Internet has been quite a success story. It is complicated, but I do not believe that its failures are as a result of being a Complex System, e.g. at the level of the AS topology or TCP behaviour. This is not to say that studying complex components in the net is not worth doing. However, none of the grand failures I have discussed are a result of complex system behaviour.

On the other hand, there are emergent properties in the higher levels, especially in business relationships and the way that innovation is forced to occur in those layers rather than below. Often, we can see computer science principles at play in the way this innovation takes place.

Are there other areas that could learn from the successes of the simplicity of the Internet? Some believe that treating the problem of sustainable energy in the same way that the Internet evolved might be one answer to that question[11]. The expansion of the stakeholder space (every home could be a generator as well as a consumer, and the network could switch small "packets" of energy) seems an attractive model that might free up a lot of creative ideas. The network itself is a starting point for creative thinking about saving energy, and much work has already taken place[3][12].

However, in answer to the question posed in the title of this paper, is the Internet "An Emergent Sea of Complex Systems *and* Critical Design Errors?", I would say that the answer to the first part is "It may be a complex system", and the answer to the second part is, "Yes, there are some critical design errors", but the errors aren't because of the complexity. The interesting conclusion, then, is that we need to redesign the Internet to reduce the likelihood of these simple operational mistakes in the future, but we need to retain the simplicity of the core Internet so that we don't introduce unnecessary, and unpredictable complex system behaviour with unexpected emergent consequences: i.e. don't throw out the baby with the bath water.

[11] See Creating the Grid OS: A Computing Systems Approach to Energy Problems `http://bnrg.eecs.berkeley.edu/~randy/Courses/CS294.F09/`, for example.

7 Acknowledgements

Thanks to Johna Till Johnson of Nemertes Research, Richard Clayton, David Evans, Tim Griffin, Steven Hand John Miller, Fernando Ramos and Vaughan Wittorff, from Cambridge University, Pan Hui from T-Labs, and to Professor Saleem Bhatti of St Andrews, and Professor Mark Handley and Dr Richard Clegg and Professor George Pavlou of UCL, for feedback on early versions of this paper.

8 Websites Mentioned

Google warns entire Internet is malware http://news.cnet.com/8301-1001_3-10153942-92.html

Pakistan on the YouTube black hole: Never mind http://government.zdnet.com/?p=3673

What DNS Is Not http://queue.acm.org/detail.cfm?id=1647302

How Akami Works http://research.microsoft.com/en-us/um/people/ratul/akamai.html

How Digg Works http://dig.com/faq

Web 2.0 http://en.wikipedia.org/wiki/Web_2.0

Internet Traffic http://en.scientificcommons.org/konstantina_papagiannaki

TCP http://tools.ietf.org/html/rfc761

Packet Switched Networks http://en.wikipedia.org/wiki/Donald_Davies

DECNET http://en.wikipedia.org/wiki/DECNET

Creating the Grid OS: A Computing Systems Approach to Energy Problems http://bnrg.eecs.berkeley.edu/~randy/Courses/CS294.F09/

Takoma Narrows http://www.youtube.com/watch?v=3mclp9QmCGs

Slammer http://www.spamlaws.com/slammer-worm.html

References

[1] Jonathan Anderson, Claudia Diaz, Joseph Bonneau, and Frank Stajano. Privacy-enabling social networking over untrusted networks. In *WOSN '09: Proceedings of the 2nd ACM workshop on Online social networks*, pages 1–6, New York, NY, USA, 2009. ACM.

[2] Randy Baden, Adam Bender, Neil Spring, Bobby Bhattacharjee, and Daniel Starin. Persona: an online social network with user-defined privacy. In *SIGCOMM '09: Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 135–146, New York, NY, USA, 2009. ACM.

[3] Luiz André Barroso and Urs Hölzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, 2007.

[4] Robert M. Bell and Yehuda Koren. Lessons from the netflix prize challenge. *SIGKDD Explor. Newsl.*, 9(2):75–79, 2007.

[5] Sonja Buchegger, Doris Schiöberg, Le-Hung Vu, and Anwitaman Datta. Peerson: P2P social networking: early experiences and insights. In *SNS '09: Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*, pages 46–52, New York, NY, USA, 2009. ACM.

[6] Damiano Carra, Giovanni Neglia, and Pietro Michiardi. On the impact of greedy strategies in bittorrent networks: The case of bittyrant. In *P2P '08: Proceedings of the 2008 Eighth International Conference on Peer-to-Peer Computing*, pages 311–320, Washington, DC, USA, 2008. IEEE Computer Society.

[7] Meeyoung Cha, Haewoon Kwak, Pablo Rodriguez, Yong-Yeol Ahn, and Sue Moon. I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system. In *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 1–14, New York, NY, USA, 2007. ACM.

[8] Meeyoung Cha, Alan Mislove, Ben Adams, and Krishna P. Gummadi. Characterizing social cascades in flickr. In *WOSP 2008: Proceedings of the first workshop on Online social networks*, pages 13–18, New York, NY, USA, 2008. ACM.

[9] Mark Crovella and Azer Bestavros. Explaining world wide web traffic self-similarity. Technical report, Boston University, Boston, MA, USA, 1995.

[10] Nick Feamster and Hari Balakrishnan. Detecting BGP configuration faults with static analysis. In *NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation*, pages 43–56, Berkeley, CA, USA, 2005. USENIX Association.

[11] Sally Floyd and Van Jacobson. The synchronization of periodic routing messages. *IEEE/ACM Trans. Netw.*, 2(2):122–136, 1994.

[12] Lakshmi Ganesh, Hakim Weatherspoon, Mahesh Balakrishnan, and Ken Birman. Optimizing power consumption in large scale storage systems. In *HOTOS 2007: Proceedings of the 11th USENIX workshop on Hot topics in operating systems*, pages 1–6, Berkeley, CA, USA, 2007. USENIX Association.

[13] Saikat Guha, Kevin Tang, and Paul Francis. NOYB: privacy in online social networks. In *WOSP 2008: Proceedings of the first workshop on Online social networks*, pages 49–54, New York, NY, USA, 2008. ACM.

[14] Hamed Haddadi, Steve Uhlig, Andrew Moore, Richard Mortier, and Miguel Rio. Modeling internet topology dynamics. *SIGCOMM Comput. Commun. Rev.*, 38(2):65–68, 2008.

[15] Murali Kodialam, T. V. Lakshman, James B. Orlin, and Sudipta Sengupta. Oblivious routing of highly variable traffic in service overlays and ip backbones. *IEEE/ACM Trans. Netw.*, 17(2):459–472, 2009.

[16] Tom Leighton. The akamai approach to achieving performance and reliability on the internet. In *PODC '07: Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, pages 2–2, New York, NY, USA, 2007. ACM.

[17] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Trans. Netw.*, 2(1):1–15, 1994.

[18] Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman. The dynamics of viral marketing. In *EC 2006: Proceedings of the 7th ACM conference on Electronic commerce*, pages 228–237, New York, NY, USA, 2006. ACM.

[19] Ricardo V. Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. In search of the elusive ground truth: the internet's AS-level connectivity structure. In *SIGMETRICS '08: Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 217–228, New York, NY, USA, 2008. ACM.

[20] Nishanth Sastry, Eiko Yoneki, and Jon Crowcroft. Buzztraq: predicting geographical access patterns of social cascades using social networks. In *SNS '09: Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*, pages 39–45, New York, NY, USA, 2009. ACM.

[21] Darryl Veitch, Nicolas Hohn, and Patrice Abry. Multifractality in tcp/ip traffic: the case against. *Comput. Netw.*, 48(3):293–313, 2005.

[22] Yong Xia, Lakshminarayanan Subramanian, Ion Stoica, and Shivkumar Kalyanaraman. One more bit is enough. *IEEE/ACM Trans. Netw.*, 16(6):1281–1294, 2008.