The Nature of Network Traffic

Kate Lance
Department of Computer Science
The University of Newcastle
Callaghan NSW 2308
clance@cs.newcastle.edu.au

Abstract

Analysis of the composition and behaviour of real large-scale network traffic is a new field of research. Most people have a feeling for how much of their own usage goes into email, telnet, ftp, netnews or www, but the overall composition of traffic is not well known: how do the relative proportions of the major protocols change over time, and how do traffic measurements from Australia compare to data from sites in the USA?

As the extraordinary growth in network usage continues, unusual phenomena are being observed for the first time, such as periodic synchronisation and self-similarity in traffic flows. Previous models appear to be inadequate to describe the behaviour of aggregated network traffic, but more realistic models based on fractal systems are providing new understanding.

1. Composition of Network Traffic

The behaviour of large-scale aggregated traffic may be illuminated by study of its component protocols, such as **ftp**, **telnet**, **snmp** (email), **nntp** (Netnews), **www** (World Wide Web), and **domain** (name services). We all know that traffic levels are constantly increasing, but it is the relative *proportions* of traffic for the major protocols that are changing most dramatically over time.

NSF Backbone traffic statistics [1] were collected until April 1995. Figure 1 is derived from that data and shows the percentages of the total number of bytes per month for the top eight protocols. The very large increase in **www** usage is often discussed, but this plot shows how truly explosive that growth has been. (Since this is a graph of relative percentages of a continually rising total, the protocols like **domain** and **irc** that remain at the same level are actually increasing in usage, and the ones that are apparently falling are probably maintaining their userbase, but not increasing.)

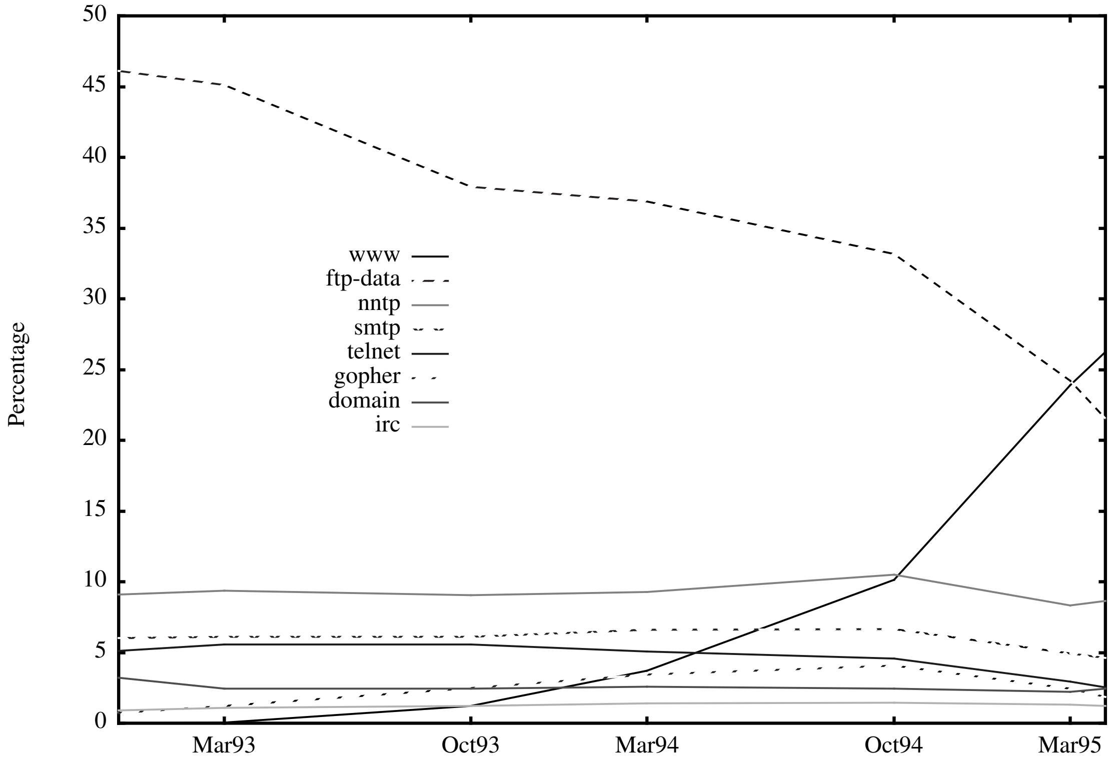

Figure 1: The relative percentages of total bytes per protocol on the NSF backbone from December 1992 to April 1995.

But are the relative proportions of protocols on the NSF backbone, or the new architecture that replaced it, the same as the proportions seen at endnodes like universities or research centres? In late February 1995 I was able to collect nine 1-hour traces on the University of Newcastle campus backbone (i.e. *before* the undergraduates gained computer access for the year), and compare it to two hours of traffic data from the University of California, San Diego campus backbone, from March 1993 (Claffy[2]); and a month-long dataset from the Lawrence Berkeley Laboratory, a basic research institute operated by the University of California, collected in March 1994 (Paxson[3]).

Figure 2 shows the relative percentages of the four currently largest-usage protocols at the three sites (1993 to 1995) compared to the NSF backbone (March 1995). The university/research sites have similar profiles, surprising given that they were collected in three separate years. They all differ from the NSF backbone data for **www** and **nntp** usage. The difference for the **www** data suggests that the enormous relative increase in **www** usage is being fueled by sites *other than* research or universities (at least during non-teaching periods): further traces collected during semester times would show whether students significantly affect university levels of **www** access.

The difference for the **nntp** data probably stems from the method of propagation: as Netnews floods from site to site, often via dial-up lines or regional networks, much of the traffic would not need to pass across the NSF backbone. It is interesting to note just how much of the traffic at an individual site is Netnews: at Newcastle it is now the largest usage protocol for both incoming and outgoing traffic.

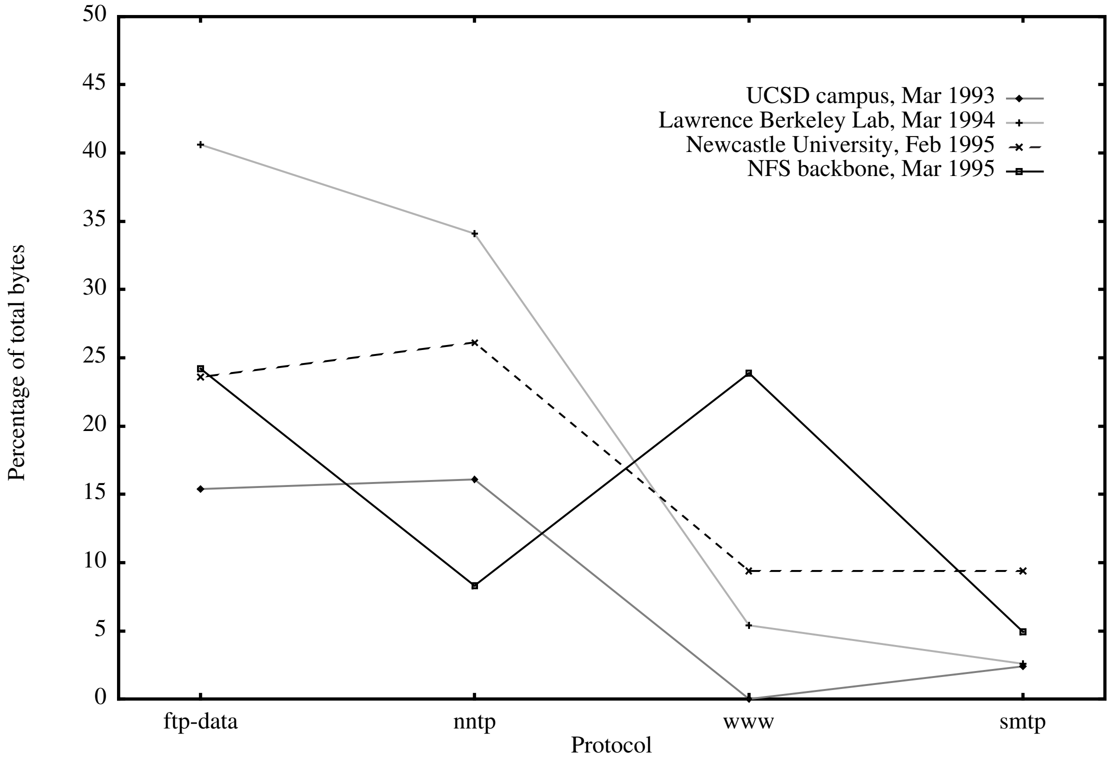

Figure 2: The relative percentages of bytes for the four major protocols at university and research sites, compared to the NSF backbone.

An interesting point for universities and government research organisations who previously paid by the megabyte for incoming overseas traffic, but not for Australian traffic: apart from the points at which they enter the country, **nntp** packets carry Australian source addresses, so they would not have previously been charged for. So any site working out their potential bills on the basis of the accounted usage of their previous ones may be surprised by how much additional data is charged for once the apparently "Australian" **nntp** component enters the picture under the new charging schemes.

2. The Behaviour of Network Traffic

Observation of real large-scale network traffic is a new area of research. Until recently models of small-scale networks appeared to be fairly well understood, but as network usage grows without apparent bounds, unusual phenomena are being observed for the first time and previous models are not adequate to describe them.

It's a common and reasonable assumption that, since the sources of network traffic are independent of each other, and arise from deterministic systems with precise rules of activity and interaction, then the aggregated traffic flows remain uncorrelated and independent, and affect each other only when conditions such as congestion occur.

Yet evidence is accumulating that previously unthought-of physical phenomena are occurring in all kinds of network environments, from TCP/IP LANs to major backbones, from telephone services to ATM networks. Although the components of these systems may be initially simple and deterministic, when they accumulate to even a minor degree curious behavior begins to manifest. This paper will discuss two such manifestations: one known to physicists for centuries, and the other, recognised for the first time because of a recent revolution in understanding of dynamical systems. In both cases, a strange kind of order is seen to emerge from extreme complexity.

2.1 Periodic Synchronisation

Synchronisation has been known since at least the 17th century, when Huygens observed that two unsynchronised pendulum clocks on a wall would start to beat in time because of the minute vibrations each set up in the wall. It is observed in all dynamic systems when there is a weak coupling between the elements of the system: in mechanical objects, in electronic circuits, and even in biological populations. Synchronisation occurs independently of the initial conditions and physical constants of the system.

Floyd and Jacobson[4] point out that computer networks are complex coupled systems that may easily evolve to high states of order and synchronisation. They describe several known examples of unexpected synchronised behaviour:

- TCP window increase/decrease cycles for separate TCP connections sharing a common bottleneck gateway.
- Synchronisation to an external clock: for example, peaks in ftp traffic for the most recent weather maps from servers every hour on the hour.
- Users' periodic background scripts and cron jobs.
- Client-server models: multiple clients may become synchronised as they check for response from a busy or recovering server (Sprite OS).

Floyd and Jacobson found a number of patterns of packet drops in multiple ping experiments across the Internet in 1992. One at 90-second intervals was traced to certain routers, but the causes of others at 15, 45 and 318 seconds were not identifiable, but it was clear that a significant amount of packet loss was occurring because of periodic processes.

They analysed the possibility of synchronisation of periodic update messages between routers, and proposed that the way the timing of a router's outgoing message is affected by an incoming update message provides a mechanism of weak coupling between the systems. They simulated a number of systems and found that synchronised behaviour occurred for a wide range of inputs, and the transition to it was very abrupt. It could be prevented by adding a random amount of time to the update intervals of each router, but the amount of time necessary was unexpectedly large, at least one full second.

They point out that, while this provides a solution to that particular problem, the phenomenom of synchronisation is inescapable in complex systems: it is possible to engineer it out in one application, but that may simply invoke it somewhere else. A further problem is that increasingly large components of Internet traffic are periodic, like user-generated scripts and variable bit rate video. Synchronisation is a wide-scale problem, and its occurrence must be a matter of serious consideration for future protocol and network design.

2.2. Self-Similar Traffic

Mathematical models are used widely in network engineering to provide algorithms for network devices like routers to most efficiently handle real packet traffic. They have to mimic the queueing behaviour of real traffic and, naturally, be computationally feasible. Similar models are the basis of software for testing network applications, like new implemetations of congestion or admission strategies, or simulations of realistic background traffic, for current and future high-speed networks.

Vital engineering decisions about factors such as buffer sizes, resource allocation and control strategies are often based upon the kind of queueing delays predicted by a specific model. Queueing theory uses the notation **A/B/m** to describe a queueing system, where **A** is the probability distribution that describes the *time intervals between packet arrivals*; **B** is the probability distribution of the time to handle a service request; and **m** is the number of servers. The standard functions used for **A** are Poisson or Poisson-like distributions, which provide *exponential* probability densities.

Recent studies (Leland, Taqqu, Willinger, and Wilson [5], Paxson and Floyd[6]) have started to question the previously widely-held assumption that packet interarrivals are well modelled by exponential functions, which imply that packet arrivals are independent over long timescales, and as traffic is aggregated into longer time-interval bins, the obvious variations and bursts of activity level out, so mean traffic levels become increasingly smooth.

Leland *et al.* show that a *theoretical* plot of arrival intervals begins to approximate white noise (i.e. a straight line) after being aggregated over just a few seconds, but that *real* network traffic shows a "bursty", irregular structure with roughly the same unpredictable profile, no matter what the scale of observation, from milliseconds up to hours. This characteristic is called **self-similarity**, and it is this property that several researchers have now shown is a characteristic of packet traffic on all sorts of networks, ranging from Ethernet LANs to CCSN telephone networks, ISDN D-channels, and variable bit-rate video over ATM networks.

Statistically this means that, rather than large-scale traffic being essentially random with some short-scale correlations between packets (the standard model), there are some correlations that persist over very long time-scales, which have a powerful effect on the overall structure, and that these correlations decay like power laws rather than like exponentials.

This has striking consequences for a number of packet traffic engineering issues, such as queueing behaviour, packet losses, buffer sizes, usable capacity, admission and rate control, performance predictions and network resource allocations. In particular, queueing performance can be shown to be fundamentally degraded by long-range correlations. Figure 3 shows that in the case of real data, delays on a queueing system start to significantly increase at around 55% of utilisation, yet standard queueing theory predicts this should not happen until utilisation reaches about 80%.

Erramilli, Naryan, and Willinger [7] divided a sequence of packet arrival data into a number of blocks. The values inside *each* of the blocks were randomly shuffled among themselves, but the overall sequence of the blocks remained the same: i.e. the short-range correlations were destroyed and the long-range ones preserved. This new arrangement of data also showed the sharp increase in delay at 55% utilisation, identical to the original, indicating that the short-range correlations could not be the cause of the delays.

Then the reverse experiment took place: the data inside each block was left untouched, but the overall sequence of the blocks was randomly shuffled, which destroyed the long-range relationships. This data then performed nearly as well as the model, with delay increasing at around 75% utilisation, which showed that it was the **long-range dependence** that was the previously unrecognised reason for delays to start at much lower utilisation levels than predicted.

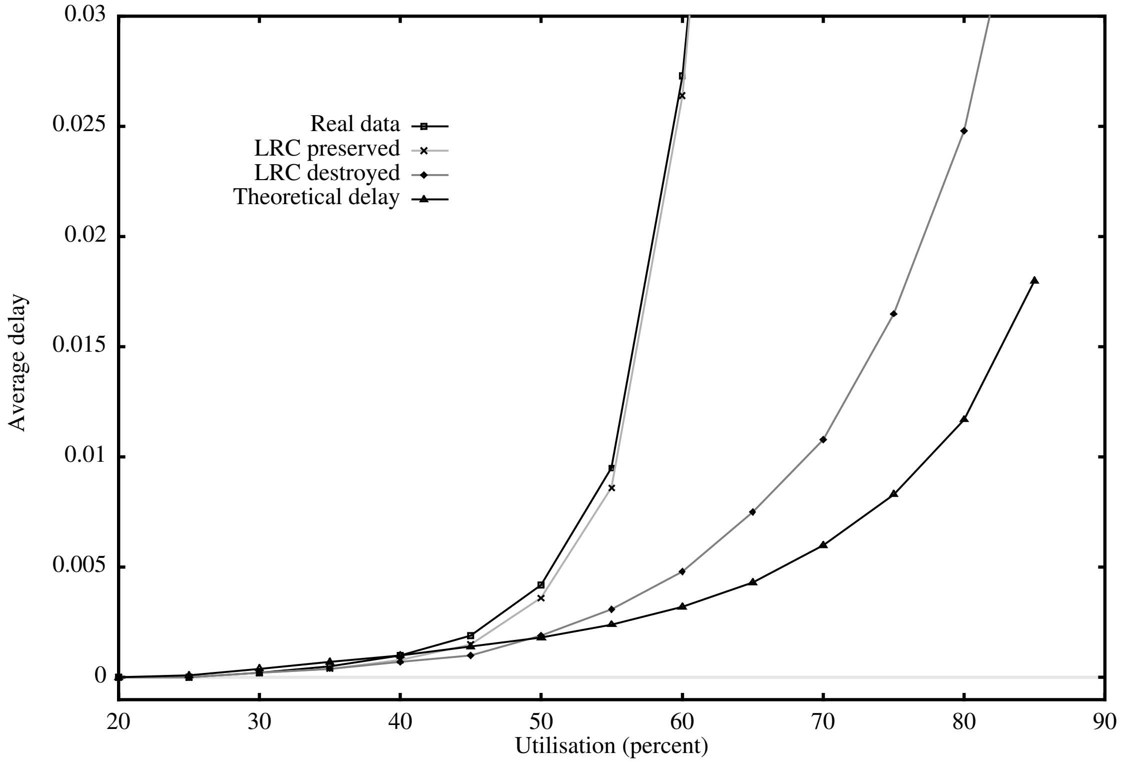

Figure 3. From Erramilli et al. [7]. Queueing delay vs. utilisation level when long-range correlations (LRC) are preserved or destroyed in Ethernet data.

So where do such strange long-term correlations between streams of traffic come from? How can randomly-generated packets, from vast numbers of different users working under wildly varying conditions, come to create patterns that can be seen at timescales as short as thousandths of a second, yet still persist at timescales of hours? Leland *et al.* picture it as "traffic *spikes* riding on longer-term *ripples* that in turn ride on still longer term *swells*".

Erramilli *et al*. suggest that one reason that aggregated traffic shows long-range dependence is that at the level of individual sources (i.e. people), on/off activity is characterised by power-law (heavy-tailed) distributions that arise from occasional periods of extended activity or inactivity. When accumulated into multiplexed traffic, the sources with the heavier-tailed behaviour have a statistical influence that predominates despite their relative scarcity.

While this describes one particular means by which self-similar characteristics might appear in random network traffic, it seems that within *any* complex dynamical system can be found simple patterns and relationships that are the same between systems of the same kind, no matter what the initial conditions might be, or how widely the behaviour of the system might diverge from case to case. Examples of this are the the way biological populations grow and shrink, how fluids behave as they become turbulent, how stock markets vary, and how the weather behaves.

Over the last 15 years or so it has been recognised that within all systems that have wildly unpredictable outcomes from minute variations in initial conditions, the **chaotic dynamical systems**, there exist simple constants and descriptors with **fractal** properties, that provide a means of identifying common characteristics across divergent behaviour. Self-similarity is one of those characteristics. Figures 4 and 5 show examples of the self-similar fractal images that can be generated by simple equations undergoing large numbers of computer iterations [8].

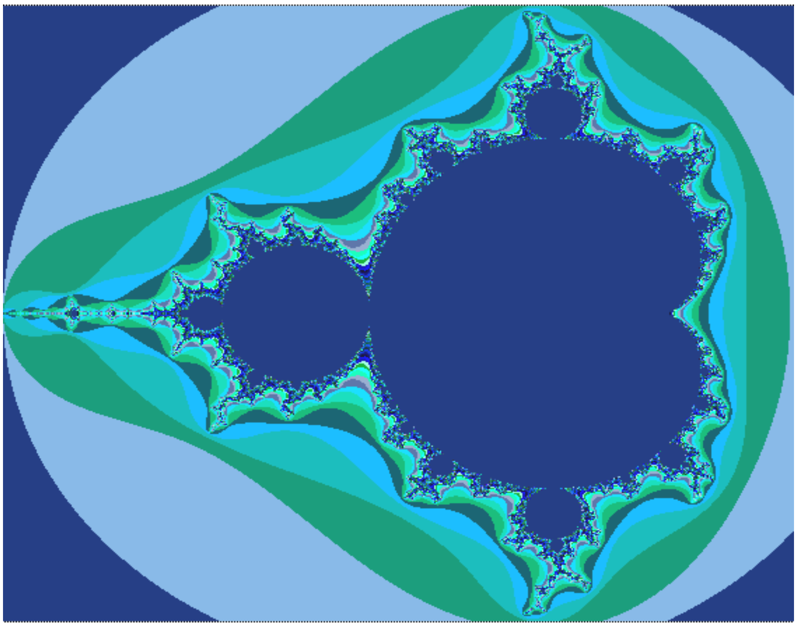

Figure 4: The Mandelbrot set.

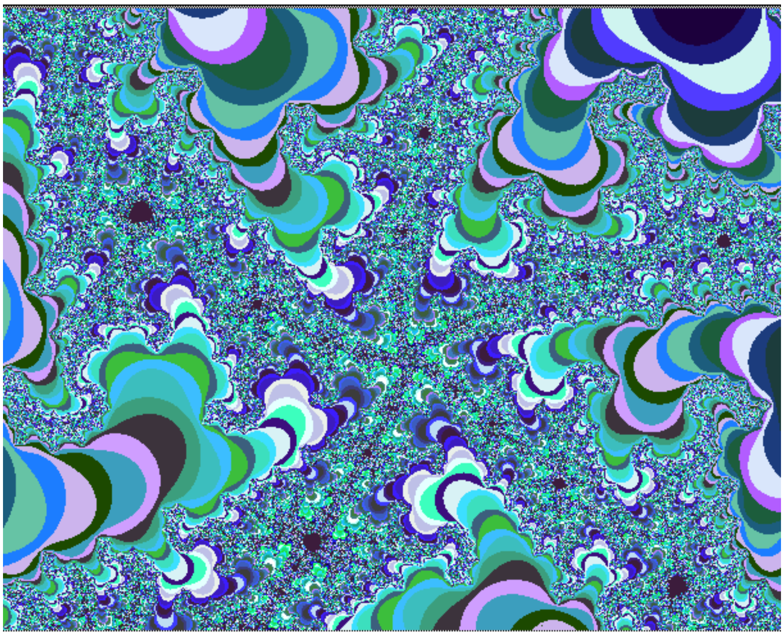

Figure 5: Detail from the Mandelbrot set.

Apart from their striking forms, the fascinating thing about them is how the patterns at one scale reccur over and over with tiny differences, no matter what the scale of observation - extraordinary illustrations of self-similarity in action. These are examples of generating extremely complex forms from simple descriptors, whereas the study of chaotic systems seeks to find the simple descriptors hidden among the highly complex manifestations.

Interestingly, the very first application of fractals in problem-solving was in the field of networks. Mandelbrot, the IBM researcher who has done so much to encourage investigation of fractals in many disciplines, was asked to work on the problem of spontaneous noise on communication lines in the early 1960s. He found a means of analysing it in terms of clusters of noise and noise-free periods that, within each cluster, had smaller self-similar clusters of noise and noise-free periods, and so on. He could not find from the data that this process reached a limit or could be prevented in any way, because increasing the signal strength just increased the noise - so he suggested instead that engineers use a strategy of redundancy and error correction to handle the problem. (A similar design approach, of course, was fundamental to the success of the flexible and robust TCP protocols.)

After that ground-breaking work, and until the recent studies discussed here, fractal insights in networking have lagged far behind that of other disciplines. Whole new fields of understanding have arisen from fractal applications in areas as diverse as fine-particle physics (surfaces of complex chemical and geological particals), ecology (population dynamics), epidemiology (disease propagation), artificial life, geophysics, astronomy (galactic clustering), economics, meteorology, fluid dynamics, music of all cultures, and even the flow of cars on expressways.

The common characteristic is that fractal behaviour is an element of all complex systems, whether they arise *from nature or from human manufacture*. Mandelbrot [9] said "...fractal methods can serve to analyse any "system", whether natural or artificial, that decomposes into "parts" articulated in a self-similar fashion, and such that the properties of the parts are less important than the rules of articulation." This suggests that a fruitful approach to network analysis is to study the forces that give rise to self-similarity among packet flows rather than try to study the behaviour of packet flows themselves.

And what about the recent traffic traces from the University of Newcastle - do they have a fractal structure? Figure 6 shows plots of packets per time interval over timescales aggregated from one-hundredth of a second to 10 seconds. It is clear that the arrivals are bursty in the same sort of way at each timescale, and there is no evidence that the profile is smoothing out to white noise. The strange manifestation of self-similarity, independent of the composition or amount of traffic, is as evident in an hour-long traffic trace from a small academic site in Australia in 1995, as it was in 27 consecutive hours in 1989 at a large telecommunications research organisation in the US [5].

Possibly the most attractive aspect of analysing network traffic in terms of its fractal properties is that, as well as being a more accurate description, it is a highly economical approach. In the past, realistic Poisson models demanded numerous parameters (which "tend to infinity as the sample size increases" [7]) whose only attraction was mathematical convenience, with no real physical interpretation. Fractal traffic models, on the other hand, are parsimonious - they need very few parameters, which are both easily estimated and have meaningful physical interpretations. They actually become more accurate the higher the levels of activity on a network.

So not only can fractal analysis offer us new methods of perceiving and interpreting the curious order hidden among the torrents of network traffic, but it can offer more efficient techniques for handling and providing appropriate resources for that traffic, as well as more realistic models to test future protocol and network strategies.

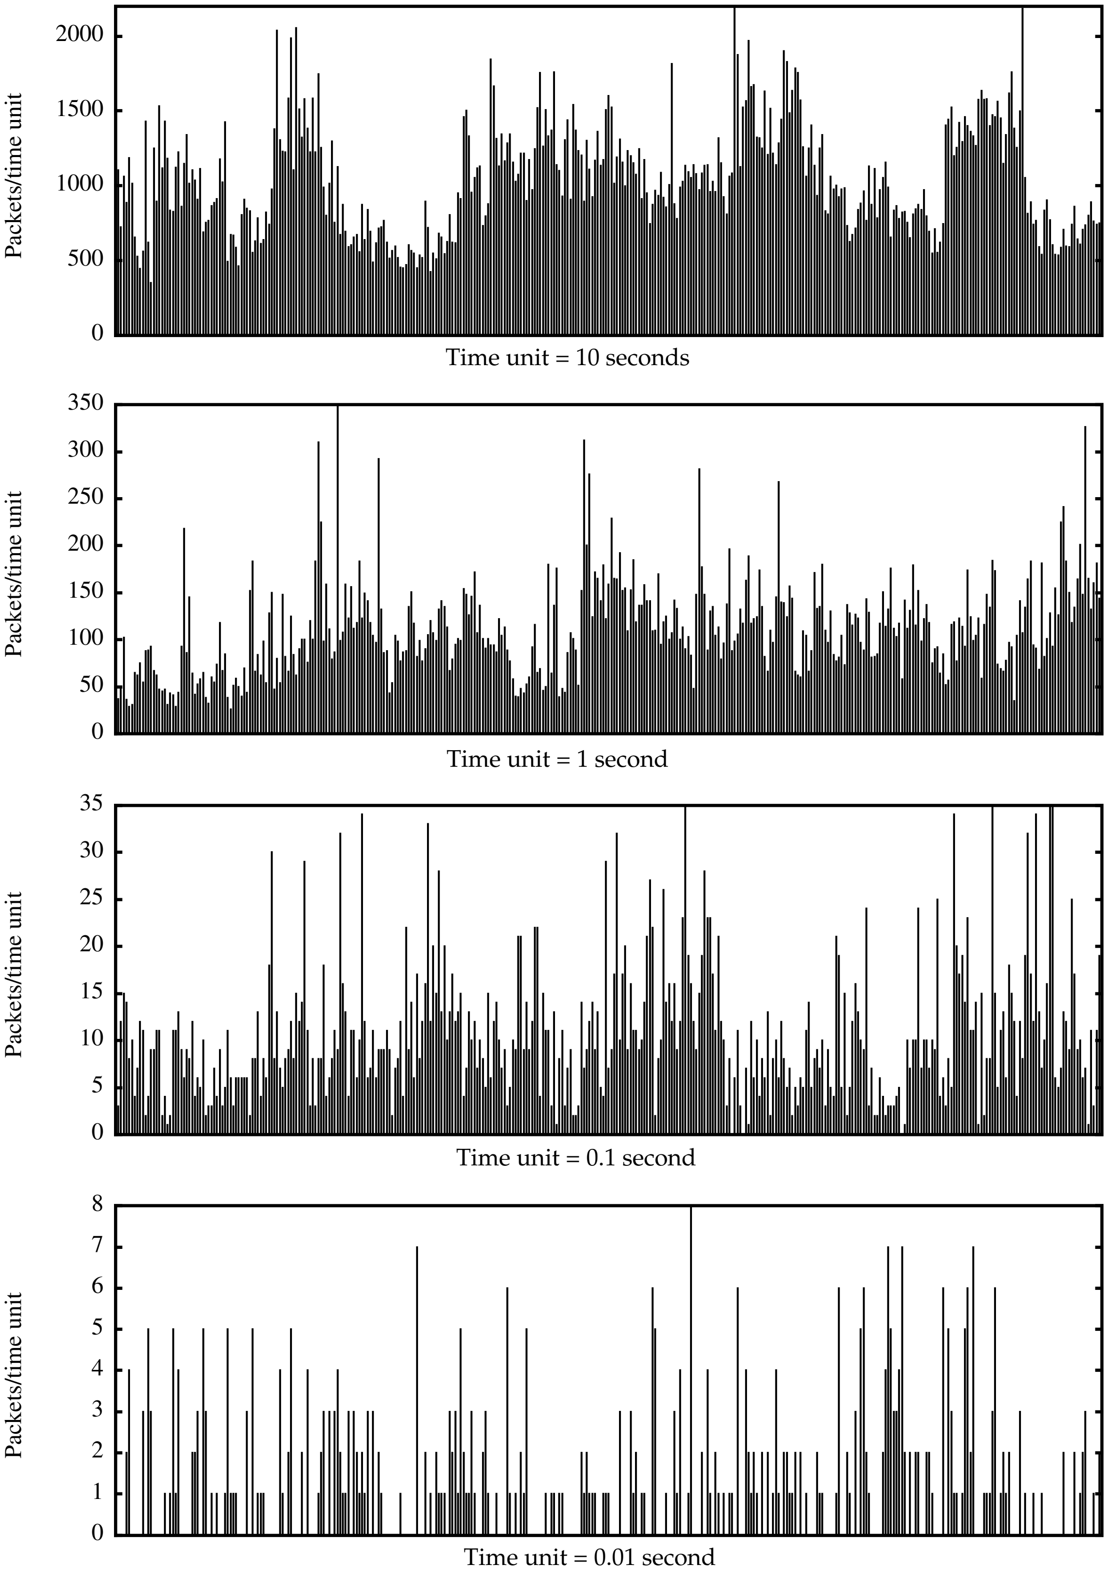

Figure 6: Packet arrivals over different timescales, University of Newcastle data.

Bibliography

[1] NFSNET Merit statistics from http://www.cc.gatech.edu/gvu/stats/NSF/merit.html

[2] K. Claffy, "Internet Traffic Characterization," PhD Dissertation, University of California, San Diego, June 1994. Available via anonymous ftp from ftp.sdsc.edu, in /pub/sdsc/anr/papers/kc-thesis94.ps.Z

[3] V. Paxson, "Growth Trends in Wide-Area TCP Connections," *IEEE Network*, 8(4), pp 316–336, August 1994.

[4] S. Floyd and V. Jacobson, "The Synchronization of Periodic Routing Messages," *IEEE/ACM Transactions on Networking*, 2(2), pp 122–136, April 1994.

[5] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Transactions on Networking*, 2(1), pp 1–15, February 1994.

[6] V. Paxson and S. Floyd, "Wide-Area Traffic: the Failure of Poisson Modelling," Submitted to *IEEE/ACM Transactions on Networking*, 1995. Preliminary version via ftp from ftp.ee.lbl.gov, in /papers/WAN-poisson.ps.Z

[7] A. Erramilli, O. Naryan, and W. Willinger, "Exprerimental Queueing Analysis with Long-Range Dependent Packet Traffic, submitted to *IEEE/ACM Transactions on Networking*, 1995.

[8] xfractint-2.01.tar.Z, fractal image generator. Anonymous ftp from life.anu.edu.au, in /pub/complex_systems/fractals/unix.

[9] B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman, New York, 1983.