Published in IET Communications
Received on 22nd February 2008
Revised on 31st July 2008
doi: 10.1049/iet-com.2008.0111

In Special Issue on Simulation, Analysis and
Measurement of Broadband Network Traffic

ISSN 1751-8628

Analysis, simulation and measurement in large-scale packet networks

R. Mondragon[1] *A. Moore*[2] *J. Pitts*[1] *J. Schormans*[1]

[1]*Network Research Group, School of Electronic Engineering and Computer Science, Queen Mary University of London, E1 4NS, UK*
[2]*Computer Laboratory, University of Cambridge, UK*
E-mail: john.schormans@elec.qmul.ac.uk

Abstract: The authors review analysis, simulation and measurement techniques, the three fundamental methods for performance evaluation in packet networks, looking at what's known, what's new and some outstanding issues. In trying to avoid re-reviewing material which has already been well summarised elsewhere, the authors concentrate on areas that are relatively new or possibly less generally well appreciated. So, under analysis, the focus is on models for network topologies and connectivity, and on wireless access. In the simulation section the focus is on techniques for scalable simulation for large-scale packet networks. Compared to the other two areas measurement is relatively new anyway, and more time is spent on motivation, techniques and some recently discovered limitations.

1 Introduction

This paper reviews the three techniques that are available for packet-level performance evaluation in packet networks: mathematical analysis, computer simulation and direct measurement. Any such paper cannot be exhaustive, as this would require three (long) books. So instead our intention is to point to some key problems, results and limitations.

Analysis can be the most convenient method to use. Where analytical formulas exist, numerical results can be generated in minimal time and with a relatively small requirement (usually) for computational power and storage space. There is a very long tradition of analysis in the evaluation of communications networks, going all the way back to the work of Erlang. However, reviews of teletraffic engineering, in general, and queue modelling, in particular, have been provided before, [1–4], so we will ignore many traditional queuing theoretical results. Instead we consider (briefly) queue models for statistical multiplexing in packet buffers, and from there we will spend a little time on the problems introduced by wireless access. Then, we move on to the area where new analytical work has recently blossomed: techniques for the analysis of network topology and connectivity, and the intersection of this with packet traffic flows.

Discrete-event simulation is widely used in the performance evaluation of complex networks and communication systems. Typically, the challenges of using simulation include modelling the target system, running the simulation and analysing the output. Associated with running any simulation is a number of challenging questions, including: how long should the simulation runs be, and how many are then required? It is necessary to balance the contradictory requirements of simulation cost and accuracy; for these reasons, a lot of recent attention has been focused on methods of accelerating simulations, and we will review some of these.

Given the limitations on the applicability of analytical techniques, and the resource requirement and the associated statistical problems that complicate the use of simulation, direct measurement of network performance may appeal. However, measurement methods require a real network to be available for experimentation. The advantage of direct measurement of network performance is that no detail of network operation is excluded: the actual operation of the real network can be monitored and measured. The limitation is that a revenue earning network cannot be pushed to its limits because customers may then complain. Alternatively, an experimental network may be limited in the number and type of traffic sources available. In this paper, we are concerned to explore two main facets of network measurement: 1) the motivational, with some pointers to the technological – how monitoring works – and 2) the limitations on accuracy when actually performing measurements of network performance.

It is important to see how the three methods fit together to provide network performance evaluation. An analytic solution can be easy to use, and provide a very general understanding of performance in a simplified scenario. This would, however, generally be backed up by a more detailed simulation of the scenario of interest, and finally by detailed network measurements. A good example might be the use of a simple queue model to provide an insight into how large is the packet buffer required in a router to multiplex VoIP packets. As reviewed in Section 2.1, only such a model operating through the load may well predict a requirement for some 10s of packet spaces in the buffer. A more detailed simulation of a network, however, may reveal that the packet arrival patterns are not simply Poisson, and so the prediction would be that the required performance is only achieved with a larger buffer (perhaps holding many 100s of packets). The performance engineer would usually then trust the simulation results, as these are based on a more detailed model. However, the subsequent measurement of a real network should help determine which (if either) of the performance predictions are sufficiently close to reality.

The combination – analysis, simulation and measurement – does not have to operate in the sequence as just described. Another possibility is for network measurements to inform analytical models and then simulation modelling too. An example of this is discovery of long range dependent (LRD) traffic patterns in measurement from real networks [5]. Subsequently, LRD became a hot topic for analysis, and was then adopted in many simulation studies.

The subject of packet networks could be impractically large, so we limit the focus of this review paper to large-scale packet networks. Inevitably, this means that some subjects will be missed, and one of the most notable left out of this paper is MANETs. For example, there is considerable recent work in the area of MANET mobility modelling and topology generation. We do not include both here, mainly because of space constraint and internal coherence.

2 Analysis

2.1 Models for packet buffers

Packet buffering is fundamental; much of the delay and loss that contributes to the overall packet-level performance happens as a result of buffering. These buffers are located in various places within (and without) the switching fabrics that are fundamental to all generations of routers, for example, see Fig. 1. Many router designs feature ubiquitous buffers: at inputs, at outputs and in many other places too, see Fig. 1.

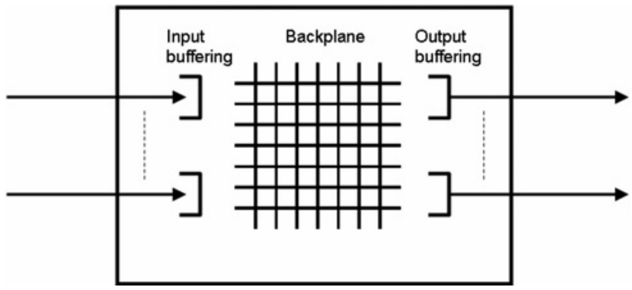

Figure 1 *Schematic representation of buffers within a generic router*

Queuing theory has long been used to evaluate delays and losses in packet switches and routers. The most basic queue models operate through the load, ρ, only. Load is the ratio of the packet service rate (reciprocal of the mean packet transmission time) and the mean packet arrival rate, that is $\rho = \lambda/\mu$.

The simplest model for variable packet lengths is the `M/M/1` in which the mean queue length is $q = \rho/(1-\rho)$. The mean packet delay time, d, through the buffer is available through Little's law: $d = q/\lambda$. These results are extremely well known, and fuller reviews can be found in [1–4]. They are recapitulated here as they are used again in Section 2.3.

It is not always sufficiently accurate to operate though the load only. Probably, the next basic model on from `M/M/1` (and `M/D/1` for fixed packet lengths) is the queue model that features a multiplex of homogeneous on–off sources, feeding a packet buffer. Such a model could be denoted `N*On-Off/D/1` if the packets are all of fixed length, and hence of fixed (deterministic) service time. This queue would arise as a natural model for, for example VoIP traffic multiplexing. Inactive (off) states would be caused when silence suppression in the Codecs results in periods in which zero packets are generated. Note that this relatively straightforward queue model could be compromised by Codecs that, for example, react mid-call to increased networks congestion (perhaps by reducing their transmission rate). One of the key researchers in modelling packet voice, in general, and VoIP, in particular, is Henning Schulzrinne. He has published a considerable number of papers on many aspects, including early work for the IETF [6] and review material [7], work specifically on signaling [8], experimental work on packet voice over WiFi access [9] and recently investigations into SKYPE [10].

In an `N*On-Off/D/1` queue, it actually rarely matters if the packet service times do vary (provided they are not power-law distributed). It is the arrival process – overlapping bursts of activity from the multiplexed on–off packet sources – that is more critical to generating large queues, and associated delays and losses.

doi: 10.1049/iet-com.2008.0111

Overlapping on–off source activity periods generate a different form of queuing to that which arises in simple 'classical' queue models like the `M/D/1`, `M/M/1`. The classical models produce (approximately) a single decay rate on a log/lin plot; however, as shown in Fig. 2, multiplexing on–off sources produces both this classical straight line, and another that 'emerges' when enough sources are active at the same time to overload the output capacity (bandwidth) of the buffer for significant periods of time. Because this effect often results in very large queue sizes (often with non-trivial probabilities), it is an important driver for the build-up of large delays (and packet loss probabilities) in large buffers, and large packet buffers is the default model of choice in packet networking. Analytic results for the `N*On-Off/D/1` are reviewed in [2–4].

2.1.1 More complex situations:

Very soon these queue models run into more realistic scenarios that are much harder to analyse. In no particular order, these include the following.

Source heterogeneity: The nice approximations for multiplexed on–off traffic sources assume that all the sources have the same parameters: mean on and off times, and traffic generation rate in the active (on) mode. Very rarely will this be exactly so in practice. Consequently, it may be necessary soon to resort to simulation, or make a further assumption that an average of the (heterogeneous) parameters can be used in the `N*On-Off/D/1`, again reducing the situation to a homogeneous scenario while providing acceptable accuracy.

Non-exponentially distributed on/off times: The main motivation here has been the discovery of power-law distributions, particularly in the active periods of real packet traffic, going back to the discovery of self-similarity in aggregate traffic patterns [5]. Explanations for this include Zipfian distributions of the sizes of downloaded objects, and the effect of aggregated TCP-controlled data traffic sources. Both of these are now recognised as major factors that result in aggregate traffic processes demonstrating a form of self-similarity. However, these are not the only mechanisms that give rise to the generation of power laws; others have also been studied for some time [11]. Origins discussed include some associated explicitly with video sequences – it has been shown that the autocorrelation of some VBR video sequences decay hyperbolically, and can therefore be well modelled using self-similar processes; also the tail behaviour of certain bandwidth distributions (in packet video) can be accurately described using the Pareto distribution. (However, there is still a lively debate regarding the effect of statistical multiplexing on buffer packet streams, even when power laws are clearly present. On the one side, it is pointed out that '...*traffic exhibits long-range dependence across many time scales so that aggregating traffic streams does not provide the statistical smoothing that would be expected from random traffic* ...' [12], whereas other authors suggest that although the long-range dependence does not disappear, its effect on queuing performance depends on the number of multiplexed active connections, and may eventually converge to a Poisson limit [13]. We do not really have room to explore this debate here.)

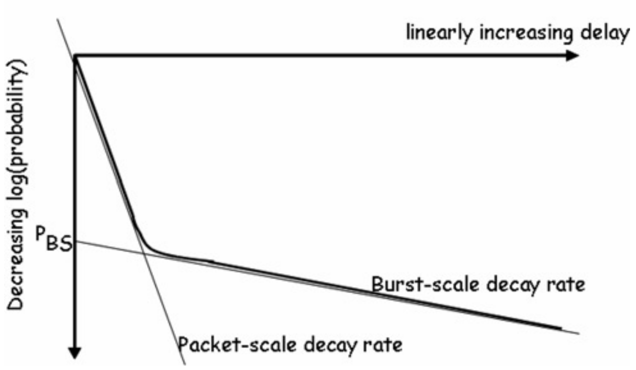

Figure 2 *Emergence of a burst scale when on–off sources are multiplexed*

A range of possible aggregate models that create self-similar processes now offer relatively parsimonious ways of reproducing 'reality' [14]. An alternative, widely used approach is to replace the whole arrival process (from *N* sources) with an aggregate model like that for Fraction Gaussian Noise (fGN) or fractional Brownian Motion (fBM). The reader interested to find out more is referred to the book self similar processes in telecommunications [15].

Traffic control mechanisms: Probably, the most widely used traffic control mechanism is random early detection (RED), [16, 17]; the operation is also explained in [4]. Because (usually) a weighted moving average of the queue length is used to affect the queue length (by early packet dropping), it is a system that is extremely hard to analyse. However, some results have been published [18, 19].

Non-FIFO scheduling: The emerging commercial model for all-packet networking relies on differential service level guarantees [20, 21]. To support this, there needs to be (at least) scheduling mechanisms that react to congestion by preferring higher priority traffic over lower. This is easier to design (at least in principle) than it is to analyse [22]. There has been extensive work in this area; the interested reader is well served by starting at Kleinrock's Queueing Systems Vol. 2 [23] and come up-to-date with the work of Bruneel *et al.* [24]. There is also some useful material in [25].

Elastic traffic: With elastic traffic, queue modelling is considerably harder. The fundamental new problem is that the traffic sources are reacting to the state of the network – speeding up, and then slowing down when they detect congestion (through loss of packets). Of course, this is the result of having TCP at layer 4.

Traffic elasticity arises because generic 'data' (as opposed to voice) can be transmitted very slowly. A useful way to think about this is through the notion of utility functions, see Fig. 3. Voice (usually) has to have a constant bitrate – 64 kbps, assuming no low rate or sophisticated Codecs that for example, can adapt the transmission rate mid-call. In contrast, non-real-time data can be transmitted almost as slowly as you wish. This means that the transmission rate can vary in response to the onset of perceived congestion; indeed, this is a fundamental aspect of the best-effort TCP/IP internet model.

Really useful queuing models are few, and do not necessarily reliably reflect network performance. One of the few simple models is the well-known approximate relationship between TCP throughput, B, and the packet loss rate, p

$$B(p) = k/(\mathrm{RTT}\sqrt{p}) \tag{1}$$

In (1), RTT is the flow round-trip time and 'k' is a constant. A more accurate formula is available in [26]. Of course, the throughput depends on the set of flows in progress; performance can (and usually does) rapidly deteriorate as the number of parallel flows increases.

Traffic elasticity causes problems for analysis; traffic sources that respond to the state of the network undercut many of the analytical techniques, for example, as based on queuing theory. It is quite likely that TCP traffic will be a very important component of packet network traffic for some considerable time (e.g. see http://ipmon.sprint.com/). Additionally, elasticity (TCP-friendliness) is being built into the application level, for example RTCP/RTP.

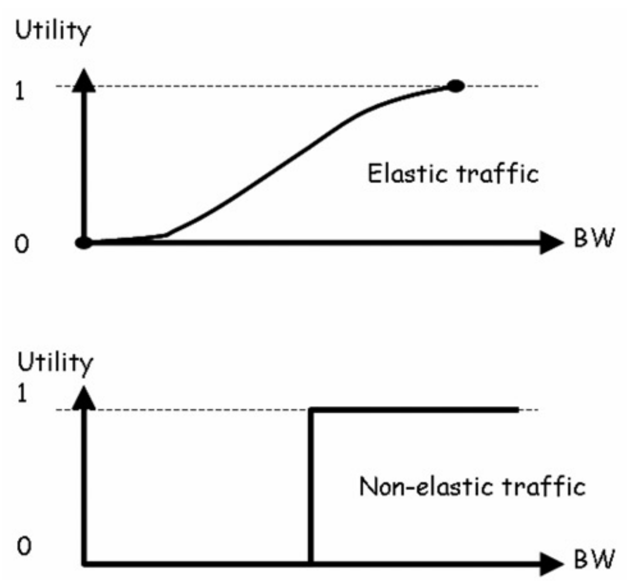

Figure 3 *Schematic representation of utility in the context of traffic sources*

2.2 Wireless access

In recent years, 802.11 WiFi has become one of the dominant Internet access technologies, both as a means of interconnecting devices in the home to a broadband ISP connection, and as a public Internet access technology in WiFi hot spots or blanket 'wireless cities' coverage [27]. From an analytical perspective, wireless access results in a very different packet multiplexing model to that traditionally used to model fixed-line access (reviewed in Section 2.1). This is because 'service' out of a packet buffer is dependent on avoiding collision with other packets also seeking to use the shared medium to access the wider (large-scale) network. A packet collision results in a back-off and a re-try after a random period. Because of this, analytical (e.g. queue) models for wireless access are much harder to develop than has been the case for many fixed-line scenarios. Additionally, any packet level analytical model has to account for physical layer effects, for example causing (non-congestive) packet loss.

When using simulation, a key problem is the increased size of the state space that occurs when the focus of the performance evaluation study is no longer just the packet level, but includes the effects of lower level (i.e. the physical level) on the packet layer performance. Measurement is also affected by this and by the much increased variability of the effects under study – we return to this subject in Section 4.2.

The effects of channel contention in the MAC layer have also been widely studied, and most published analyses refer to a paper by Bianchi [28] that accurately models the 802.11 distributed coordination function (DCF) under ideal channel conditions. This uses a Markov chain to model the back-off algorithm behaviour, and predicts the throughput achieved under saturation conditions, that is, when each station always has a packet ready to send. Two important characteristics of wireless access are apparent from this analysis: (1) the total throughput (not just individual throughput) decreases with the number of stations accessing the wireless medium, and (2) the total throughput increases with the size of packets being sent. Both characteristics are indicative of the effect of the back-off algorithm on the time it takes for a station to access the medium. In terms of a classical queue model for the packet behaviour, the time to 'serve' a packet comprises the sum of the time to acquire the channel and the time to transmit the bits over the channel. It is the former that has the dominant effect, and is very dependent on the numbers of other stations attempting to access the channel. Hence, it is useful to think of an 802.11 wireless access point as having a packet processing limit determined by the number of stations accessing it, rather than having a link-speed defined in terms of the raw data bit-rate achievable under ideal channel conditions, as in a 'simple' queue model.

The non-congestive nature of packet loss in wireless access has an adverse impact on TCP operation, which normally

doi: 10.1049/iet-com.2008.0111

interprets loss as an indication of congestion. This has prompted a substantial amount of research on modifying TCP for use in wireless communications, involving a variety of techniques aimed at detecting whether losses are caused by congestion or not (see [29] for a review of these issues, and also Chapter 8 of [25]). Chakravorty *et al.* [30] was probably the first paper to approach the problem of explicitly telling apart packet loss caused by noise from packet loss caused by buffer overflow. The authors suggest a 'proxy' system, avoiding any changes to the underlying nature of TCP. This acts to aggregate TCP flows to any particular mobile host, and has the advantage of being able to utilise the statistical dependence between them, significantly improve network performance in terms of latency and throughput. They test the performance of their idea by measurement of an implemented system.

Considerable work has also recently focused on possible extensions to the 802.11 standard to better support QoS. In [31], an enhanced proposal for QoS support does not change the fundamental behaviour of DCF, with its back-off algorithm, but rather allows different categories of traffic to use different sets of parameter values. This enables, for example, voice packets to have an increased probability of succeeding in channel contention, hence suffering less delay than data packets. However, the emphasis on a packet processing limit (rather than a raw bit-rate limit) is still valid.

2.3 Connectivity, topologies and traffic flows

From here until the end of Section 2, we concentrate on reviewing the analytical techniques developed recently for connectivity, topologies and traffic flows [32].

An important part of any network model is the topology that defines the interaction between the links and nodes of the packet network. The topology of a network is correlated to the traffic that it carries because of the routers, traffic shapers, etc. For a large packet network, measuring and analysing its connectivity is a challenging task. For example, the connectivity of the Internet can be interpreted at many different levels. The router level can be considered as interconnected physical entities. At this level, the nodes describe the routers and switches managing the passage of traffic through the network. The links represent the different physical connections between the nodes, for example, optical fibres, copper, wires etc., and have specific directions between endpoint nodes. At the router level, a basic description of the topology should include the bandwidth of the links, the direction that the Internet traffic follows.

At the managerial level, a network is divided into subnetworks, where each subnetwork adheres to common routing conventions, usually the Interior Gateway Protocol. The management of a subnetwork and its routers fall under one administrative entity called an Autonomous System (AS). The AS should exhibit to other ASs a coherent interior routing plan with the destinations reachable through the AS.

At the AS level, the network can be considered as an abstract space where the pertinent property is the connectivity between ASs. At this level, we may disregard many physical properties of the network like the geographical location of the ASs, which could be in different continents, or the directionality of the links and the link capacities.

A network can also be considered at the applications level, for example in multicasting, where it is possible to consider the multicast network (e.g. MBONE) as a virtual network built on top of the underlying network. In this multicast network, the connections between multicasting nodes are represented by physical links or by tunnelling connections. The tunnel is a virtual link consisting of possibly many physical nodes and links. The multicasting-router considers the tunnel-connection as a single link connecting two routers.

In general, when using a topological description of a network, care has to be taken when defining the meaning of a node and a link, for example in internet protocol television (IPTV), which is a multicast network, the distribution of content is done using a 'walled-garden' network; in this case the nodes and links represent physical entities [33]. So, as noted by Floyd and Kohler [34], the right topology scenario for a packet-level study will require a problem-specific model.

2.3.1 Fundamentals of network topologies and connectivity modelling:

The basic characteristics of a network are its number of nodes N, its number of links L and the connectivity between the nodes via the links. There is considerable data describing the Internet at the router and AS level [35]. Network connectivity is obtained by direct probing of the network or by using routing tables, and the connectivity of the nodes is described by an adjacency matrix whose a_{ij} entry is 1 if node n_i is attached to node n_j and 0 otherwise. For undirected networks, like the AS network, $a_{ij} = a_{ji}$. A principal parameter to characterise a network is the degree k of a node, which is the number of neighbours that a node has, $k_i = \Sigma_j a_{ij}$. Fig. 4 shows the description of a network by the adjacency matrix and its node's degree.

A first step to analyse a network is to measure the degree distribution $P(k)$; the fraction of nodes in the network with degree k. In 1999, Faloutsos *et al.* [36] analysed the connectivity of the Internet at the AS level. Critically, they discovered that the degree distribution follows a power-law decay $P(k) \sim k^{-\gamma}$, $\gamma = 2.2$. The implication of this discovery is that in the Internet, there are a large number of nodes of very low degree and relatively few nodes with a very high degree of connectivity. Previously, networks were

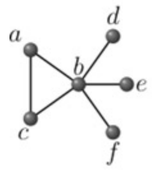

	a	b	c	d	e	f
a	0	1	1	0	0	0
b	1	0	1	1	1	1
c	1	1	0	0	0	0
d	0	1	0	0	0	0
e	0	1	0	0	0	0
f	0	1	0	0	0	0

degree	
a	2
b	5
c	2
d	1
e	1
f	1

Figure 4 *A simple undirected network, its adjacency matrix and the degree of its nodes*

described using simple random networks, where the degree distribution is well approximated by a binomial distribution. This would give a false representation of the Internet, see Fig. 5.

The observation that the degree-distribution decays as a power law has important repercussions for modelling packet networks. Take, for example, the length of the shortest path. The shortest path l_{ij} between node n_i and n_j is the path with the smallest number of hops. This quantity is relevant for routing processes and in the convergence of the traffic dynamics in a network simulator. The average of all possible shortest paths $\langle l \rangle = 1/(N(N-1)) \sum_{i=1}^{N} \sum_{j=1, j \neq 1}^{N} l_{ij}$ measures, on average, how far is a node from another. If the Internet was a random network, the average shortest path will scale as $\langle l_{\text{rand}} \rangle = \ln(N)/\ln(\langle k \rangle)$, implying that the distance between a pair of nodes is small when compared with N, the size of the network. However, for a power-law network, the average shortest path scales as $\langle l_p \rangle \sim \ln(\ln(N))$, meaning that in power-law networks, the distance between a pair of nodes is extremely small. In terms of the average shortest paths, a power-law network grows very slowly, so the growth of the number of hops is not an issue when evaluating its performance. This is particularly relevant for the analysis and performance of an evolving network.

Topology and traffic: When considering the shortest path routing mechanism, a topological quantity that is being used in Sociometrics can be applied. This is the concept of betweenness centrality [37]. The betweenness centrality for a node is defined as follows: given a source s and destination d node, the number of different shortest-paths between them is $g(s,d)$. The number of shortest-paths that contain the node w is $g(w;s,d)$. The proportion of shortest-paths, from s to d, which contain node w is $p_{sd}(w) = g(w;s,d)/g(s,d)$. The betweenness centrality of node w is defined as

$$C(w) = \sum_{s=1}^{N} \sum_{d=1, d \neq s}^{N} p_{sd}(w) \qquad (2)$$

In (2), the sum is over all possible pairs of nodes with $s \neq d$. In the context of packet networks, a node with a high betweenness centrality has a high 'status' because it stands between many other nodes on the paths of communications. It is also possible to define the centrality of a link, that is, the proportion of routes that use a given link. If the packets on the network are distributed evenly through all the shortest paths, then the normalised betweenness $\widehat{C}(w) = C(w)/\sum_{v=1}^{N} C(v)$ gives the proportion of usage of node w. Notice that the betweenness and the average shortest path are related by $\sum_{i=1}^{N} C(i) = N(N-1)l$.

The load at a node in terms of its betweenness can be obtained using Little's law. Little's law can be thought of as a flow conservation law, which can be stated that, at steady state, the number of delivered packets is equal to the number of generated packets [38]. In a homogeneous

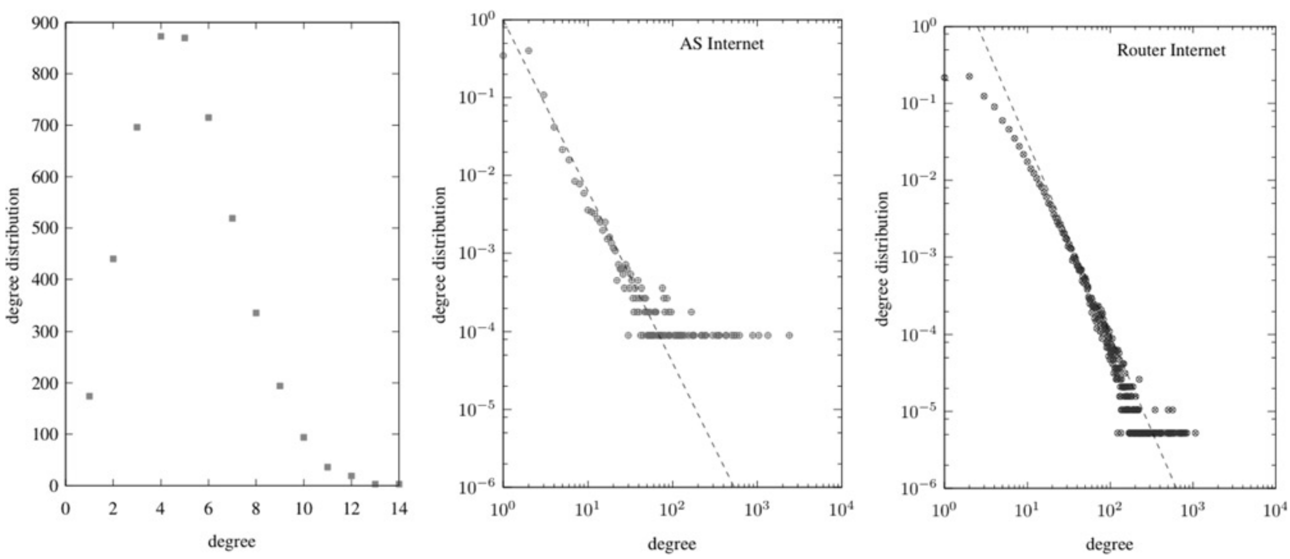

Figure 5 *Comparison of the degree distribution for (left) a random network (middle) the Internet at the AS level and (right) the Internet at the router level*

Degree distribution of the Internet at the AS and router level was obtained using the data published by CAIDA

IET Commun., 2009, Vol. 3, Iss. 6, pp. 887–905
doi: 10.1049/iet-com.2008.0111

network, where the density of source and sinks of traffic in a network is δ and each node produces a traffic load $\Lambda_i = \Lambda$, the change in the total number of packets $n(t)$ in the whole network is given by

$$\frac{\mathrm{d}n(t)}{\mathrm{d}t} = \delta\Lambda N - \frac{n(t)}{\tau(t)} \tag{3}$$

In (3), $\delta\Lambda N$ is the average rate of packets generation per unit of time, $\tau(t)$ is the average time spent in the system, and $n(t)/\tau(t)$ is the number of packets delivered per unit of time. Little's law does not depend on the arrival distribution of packets to the queue, the service time distribution of the queues, the number of queues in the system or the queuing discipline. If the load is low, the queues at the nodes tend to be empty and the average delay time is the average shortest path$\langle l\rangle$, that is $\tau(t) \sim \langle l\rangle$. For higher loads, the transit time can be approximated as the average shortest path plus the time in the queuing system $\tau(t) \sim \langle\tau\rangle = \sum_{i=1}^{N} T_i/N$ where T_i is the time spent in queue i plus the service time of the server. If the network is not congested, the steady-state solution $\mathrm{d}n(t)\mathrm{d}t = 0$ gives $\bar{n} = \delta\Lambda N\langle\tau\rangle = \Lambda\sum_{i=1}^{N} T_i$ and the general solution is $n(t) = \Lambda N\langle\tau\rangle + k_1\mathrm{e}^{-t/\langle\tau\rangle}$ where k_1 is an integration constant.

To express $\bar{n}$ in terms of the betweenness, we can start by assuming that the queues are M/M/1 queues with average arrivals λ_i, service rate μ_i and traffic intensity (load) $\rho = \lambda_i/\mu_i$ then $T_i = 1/((1-\rho)\mu_i)$. The average number of packets that arrive at node 'i' is given by (4), see [39]

$$\lambda_i = \Lambda N\delta\left(l\widehat{C}(i)\right) = \Lambda\left(\frac{C(i)}{N-1}\right) \tag{4}$$

where N is the number of packets generating an average load Λ per unit of time, l is the average shortest path of the network and accounts for the average number of packets that were produced in the past and they are still in transit and $\widehat{C}(i)$ is the proportion of all the packets in transit that pass through the node i. Notice that l and $\widehat{C}(i)$ are dimensionless topological quantities. The total number of packets as a function of the betweenness is $\bar{n} = \sum_{i=1}^{N}(\Lambda(N-1))/(\mu_i(N-1) - \Lambda C(i))$. For low load $\Lambda \approx 0$, then $\bar{n} \approx \Lambda Nl$. For high loads, the majority of the packets in the network are in the busiest queue. If m labels the busiest queue, then $\bar{n} \approx \bar{Q}_m$, at the congestion point $\bar{n} \to \propto$ and the critical load for node m is $\Lambda_m^* = \mu_m(N-1)/C(m)$. Equation (4) is an example of how properties of the traffic can be estimated from topological considerations. The model presented here is a very simple approximation of a network; nevertheless, it gives an insight on the possible relationships between topology and traffic [38, 40–43].

Looking ahead to the next topic, simulation, it is clear that a key question is – can we exploit the topology of the network to speed up the simulation of a large network without losing accuracy? This is still an open question, and in part it is related to our understanding of the network connectivity and how to model it. It is known that describing a network like the Internet using only the degree distribution $P(k)$ does not give a full description of how the elements of the network are connected. A better description can be obtained from the correlations between the degrees of different nodes [44–46]. In a finite network, this correlation is defined by the degree–degree distribution

$$P(k, k') = \frac{1}{N}\left\langle \sum_{i,j=1}^{N} \delta(k_i, k)a_{ij}\delta(k_j, k')\right\rangle \tag{5}$$

Equation (5) gives the probability that an arbitrary link connects a node of degree k with a node of degree k' ($\delta(n,m)$ is the Kronecker delta).

In scale-free networks, because of the small number of nodes with high degree and the finite size of the network, it is not possible from the network's measurements (Section 4) to evaluate accurately the degree–degree distribution. Hence, the structure of the network is characterised by using different topological measures that are related to the problem under study. For example, the clustering coefficient is used to measure the number of short loops in the network, relevant for alternative routing, or the vulnerability of a node measures how the removal of that node affects the lengths of the shortest paths. There are many other measurements [47]; however, there is some redundancy in these measurements as they tend to be correlated, and there is no consensus in the research community about which measurements we should use to obtain a good description of the network [48–50].

Topological models: The networks research community has been producing topological models of large networks, in particular trying to generate networks with a power-law degree distribution. Until now, there is no consensus in the research community in respect of the mechanism responsible for the power-law distributions observed in packet networks. Power-laws appear in many different fields of research and can be generated in many different ways, and a good review of thinking across a range of disciplines is provided in [51]. More and more network topological models tend to reflect this diversity [52, 53]. A good topological model is a practical tool as it can be used to test 'what–if' scenarios and it can provide predictions of the network's evolution. There are two general approaches [54–56]: (1) static models [57–59] based on random networks, and (2) dynamical models based on network growth models [60, 61]. The latter are considered as being more promising as, if correct, they can describe the evolution of networks [56, 61]. These two approaches can also be divided into descriptive models, based on matching various topological properties of a network: these are used to study which topological properties give a good description of a network. Alternatively, explanatory models attempt to encapsulate the core principles and factors responsible for the network's structure and evolution, and in particular the router network [61].

A proper understanding of the connectivity of large networks and validation of all the network models is lacking because of the incomplete measurements of the Internet connectivity and the limited quality of the available measures and measurements [56]. We return to the whole issue of measurements in Section 4.

3 Scalable simulation for large-scale packet networks

Simulation investigations form a vital part of networking research, and have long been used by the research community in, for example, the design of protocols and evaluation of quality of service mechanisms. A thorough overview of the main theoretical problems in simulation is given in [62], while a review of simulation verification and validation is given in [63]. In [64], one of the best known practitioners of the whole science of simulation (Averill M. Law) gives a review of the key problems mentioned in the introduction to this paper: how to choose the warm-up period, the run lengths and the number of replications; [64] also reviews the question of how to avoid the major pitfalls in simulation output analysis.

Regarding packet network simulation, specifically many simulators have been developed by a research group, or occasionally an individual investigator, to address a specific target protocol or network type. Fewer tools address general, multi-protocol, network simulation and offer a comprehensive, advanced, simulation environment. In passing, however, we take this opportunity to mention an experimental testbed (EMULAB), and an Internet data collection activity ('Day in the life of the internet challenge'). EMULAB is a network testbed, hosted by the University of Utah, which provides interested researchers a flexible environment in which to '... develop debug, and evaluate their systems' [65]. Additionally, [66] reports 'A day in the life of the internet', which is a community-wide experiment in data measurement. In the first quarter of 2008, CAIDA (and the DNS Operations, Analysis, and Research Center, OARC) conducted their third 'Day in the life of the internet' data collection activity, and as previously targeted, a 48-h collection period. The 9th CAIDA-WIDE Workshop [67] was held to coordinate this event.

In our paper, we now review the focused theme of simulation for very large packet networks, including techniques for event reduction, federated simulation, issues associated with visualisation and variance reduction techniques, for example RESTART.

The challenge is to address a topic area that is complex, heterogeneous, distributed and continually changing, and to do this effectively, a general-purpose tool requires an active community supporting and further developing the tool. Common simulation platforms, along with benchmarking, facilitate the reproducibility of research results. This is a vital component of engineering endeavour, and enables efficient evaluation and comparison of for example alternative protocols or switch architectures. Two multi-protocol simulation tools, ns [15] and Opnet [68], stand out as examples of widely used platforms from open source and commercial stables, respectively. Both tools benefit from extensibility, and large user communities, with researchers contributing new models.

However, today's simulation tools are typically incapable of modelling large-scale networks. This lack of capability is not only a function of the model size (e.g. the number of hosts, routers, flows etc.) but also of the model state space. Researchers have argued that it is simply not feasible to build a network simulation of, for example, the entire Internet [69]. One conservative estimate in 2002 calculated that 100 s of Internet activity would require more than a year of CPU time, about 300 terabytes of main memory, and 1.4 petabytes of disk storage [70]. To add to the challenge, the Internet has been increasing in size faster than hardware platforms have been improving in performance. Even if it were possible to model, say, 100 s of Internet activity in a realistic timeframe, it is questionable whether it would be of any use unless huge numbers of such simulation experiments could be run to explore the underlying behavioural interactions.

There have been various parallelisation endeavours to address the simulation technology required to build large-scale network models that can process large volumes of events. For example, PDNS [71] is an extension to the publicly available tool ns, using a runtime infrastructure based on the High Level Architecture standard [72]. PDNS is a federated simulation approach with conservative synchronisation, interconnecting multiple copies of ns, in which each model different parts of the network. It has been tested on hundreds of processors with network models of a hundreds of thousands of nodes. However, as yet, there has not been corresponding work to address the efficient exploration of the huge state spaces involved. Effective large-scale simulation modelling needs to distribute the event processing over many co-operating processors, and also (1) to reduce the number of events by intelligent use of modelling techniques and (2) to identify which parts of the state space to focus in order to explore the key underlying behaviour interactions.

So, as well as computer power and simulation technology, abstraction is vital for speeding up simulations. In essence, this involves using a simpler model, with fewer events in the simulator, and identifying which input factors are the most important. The model either removes, or summarises,

doi: 10.1049/iet-com.2008.0111

details considered less important. For example, fluid simulations (also called rate-based simulations) neglect the small-scale fluctuations in packet queue behaviour and focus on source traffic rates, and burst scale queuing [73]. Early work on this, in the context of ATM technology, found that the speed-up attained depended primarily on the traffic characteristics, that is, the number of cells in a traffic burst [74]. In ns, the session level modelling option ignores queuing delays and replaces hop-by-hop packet forwarding behaviour with a precomputed (fixed) end-to-end delay [75]. Aggregation techniques also help to reduce the computational burden, either by reducing the number of events (e.g. flow aggregation in fluid simulation [73]), or by reducing the amount of state information that needs to be stored and acted upon. Typically, aggregate traffic models are used for cross-traffic: they can be combined with queue models to remove the need to handle background traffic events [76], or employed to remove the memory overhead of thousands of individual source models. In both cases, a significant challenge is to model accurately the response of the aggregate to individual packet drops.

As mentioned in Section 2, traffic elasticity causes many problems for performance evaluation. Simulation studies of elastic traffic are affected because the state-space is hugely increased, greatly increasing the time required to reach valid results from any simulation study. A significant challenge is to model accurately the response of an aggregate traffic stream to individual packet drops. Floyd and Paxson noted in 2001 that such reactive models were beyond the state of the art [69]; to the knowledge of the authors, this is still the case.

So, although the most obvious constraints on the feasibility of large-scale simulations are CPU time, computer memory and disk storage, there are also significant methodological challenges: reproducibility of experiments and validation of simulations in real world systems [77]. Even though a common simulator platform, with publicly available simulation scripts (recommended in [78]) are steps in the right direction, the task of reproducing and validating simulated data from large-scale scenarios is substantial. Partly, this is because of the properties of large-scale communications networks. As mentioned in Section 3, network topologies and traffic properties are difficult to characterise [79], and typically key aspects of the system are not directly observable, making them difficult to measure. But there can also be complex interactions between protocol layers and subtle couplings between different elements, problems that are also a considerable challenge in wireless Internet scenarios. Non-linear feedback mechanisms mean that slight changes in parameter values can cause sharp transitions in behaviour (phase effects) [80, 81]. Visualisation issues for large graphs (to represent connectivity and performance) are known to be NP-hard problems [82], and this clearly exposes another problem with large-scale network simulation experiments: how effectively can the state space be explored, particularly to investigate rare events, given the sheer scale of the problem space? The RESTART [83] technique increases the occurrence of rare events in small-scale (single queue) simulations by restarting the simulation in certain system states, but this approach could pose significant problems of coordination and state 'rollback' for distributed large-scale simulations. Ma and Schormans [84] considered aggregating simulation acceleration techniques (RESTART and Traffic Aggregation (TA), e.g. see [84]), and found that the overall (combined) speedup was better than a linear sum of the two considered separately.

These methodological challenges of reproducibility and validation have led researchers to consider scenario specification (e.g. topology and traffic models), measurement processing (e.g. snapshots of system state) and visualisation (e.g. of network dynamics). These developments need to be incorporated into a holistic methodological framework to support exploration of the large-scale state spaces to be found in today's network modelling challenges. This is a key aspect of what was reported in 2003 as the 'meta-grand challenge facing networking research' [77]: to develop new network theories, architectures and methodologies that will facilitate the development and deployment of the next generation of services and applications. So, the major challenges of packet network simulation modelling are: scenario scale and heterogeneity, measurement processing and visualisation, multi-level abstraction techniques, reproducibility and validation and partitioning of federated simulations. In spite of these issues being identified by the community some years ago, there has been relatively little co-ordinated progress. Most collaborative developments have focused around simulation technology, for example, the ns3 project that is incorporating a distributed simulation framework, and support for multi-core and 64 bit systems, into a successor to ns2 [85].

However, there have been some methodological developments. The research community has focused on deriving simulation scenario test suites for particular issues, for example, models for the evaluation of transport protocols being produced by the Transport Modeling Research Group in the Internet Research Task Force [86]. In measurement processing, the standardisation of Quality of Experience (QoE) measures such as the E-model R-factor [87] for voice service has paved the way for whole of network QoE visualisation. This has enabled network performance to be summarised in ways that aid fuller understanding of the underlying physical behaviour. For example, the interaction between routing and QoS mechanisms, and how together they cope with various forms of network degradation. This interaction has been compared in terms of the proportion of user connections meeting voice quality requirements as network load varies over a wide range. The shape and separation of QoE contours clearly indicate the effects of load balancing and QoS configuration [88].

There has also been progress in model abstraction for simulation acceleration. A time-stepped session-level technique, combined with both packet and fluid modelling abstractions, trades the number of events for simulation accuracy in scenarios of up to 1000 nodes [89]. An important contribution is the development of an analytical expression estimating simulation error (based on System Theory) that can be integrated into the supporting methodology [90]. Another aspect of model abstraction is the development of techniques to reduce the number of significant events to be modelled in packet queues. Much of the existing work has concentrated on FIFO queues, but recent developments have extended this to QoS-enabled configurations. Work is in progress on an analytical model that estimates the impact of the background flows on the foreground streams, taking into account GPS scheduler behaviour [91].

There is still plenty of scope for further research, particularly in the area of methodological support for large-scale packet network simulation. An obvious opportunity for multidisciplinary research is that of experimental design for packet network simulations. This is an important area that as yet has seen little coordinated research effort from the community. Tools such as Opnet are able to distribute a series of simulations to multiple machines to support parametric studies and replications. But, as far as we are aware, no current network simulation tools have built-in experimental design support to guide the process of exploring the model state space.

4 Measurement

Measurement is probably the least well explored of the three techniques. However, there is already a fine book on the subject [92].

4.1 Motivation – why measure?

Why measure the Internet? The motivation to measure the Internet, both its components and users, comes from a desire for further technical and social understanding. This is inevitable given that the Internet is now a routine and, indeed, integral part of many people's work and social lives.

The Internet may be measured to understand how it is used – a social motivation. Such a measurement may answer questions about the nation's youth and their use of social-networking sites, for example [93] (important if your new advertising target is that demographic), the latest network-based application [94, 95]. Also, perhaps, to provide an understanding of the popularity of sources of news and media; this could be critical in an analysis of editorial control in a media-ownership dispute. An interesting example of network measurement with social implications was when [96] was cited as legal evidence to illustrate the (lack of) impact of regulation upon the popularity of peer-2-peer networks.

Aside from socially motivated measurement, technical measurement is used to aid both engineering and commercial interests. The engineering application of measurement is an integral part of an optimisation process: both to provide a baseline of pre-optimised performance and to quantify any improvement. Network-provider examples exist both among network operators (http://www.nanog.org/subjects.html) demonstrating effective measurement-use and presentation, and among the researchers advancing specific topics to assist in network operation. For example, measurements provide accurate information for traffic engineering [97] and identification of anomalies [98]. A commercial incentive for measurement is also strong. For example, a significant number of company business-cases presume the Internet; we could not imagine a Google, an Amazon, or an eBay without the Internet.

The current state of the art in topology discovery still consists of either sending packets through the network and plotting their trajectory, for example, skitter (http://www.caida.org/tools/measurement/skitter) as used in [99], or reconstructing the router-perspective, for example, route-views (http://www.antc.uoregon.edu/route-views) or looking-glass (http://nitrous.digex.net). Despite rather primitive mechanisms, topology research, driven in part by its importance for a robust and reliable Internet, has both immediate value to network-operations and a significant research-momentum.

Despite the importance of measurement for auditing, security and understanding, this field remains a discipline in its infancy. Measurement is not a first-class objective for a router or switch; manufacturers have spent considerable energy to ensure that the packets move with the highest speed and accuracy; however, measurements are an add-on. The installation of ideal instrumentation on every link in a network is financially, legally and technically implausible, for a large network. So any measurement approach will rely on incomplete information. This means that measurement often requires invoking sampling theory, inference, and extrapolation, all with a fundamental understanding of what is being measured and to-what-end the measurements will be used; see Section 4.2.

Topology research has often been in concert with network-tomography, the correlation of network measurements to infer what we cannot measure directly. Network monitoring is expensive and impractical deployment for all links; however, understanding the use of a network and optimising the operation of a network requires accurate reproduction of traffic-matrixes [100, 101]. Simple link-utilisation is often not sufficient, it does not reveal the efficiency of the network, nor do such measurements tell the nature of the traffic. However, better measurements require more information, or improved methods to understand old information. In order that measurement infrastructure is not overwhelmed (and becomes a bottleneck), metrics are constructed from sampled data,

doi: 10.1049/iet-com.2008.0111

such as NetFlow (http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html) and equivalent compatible mechanisms. NetFlow summarises (sampled) packet data into records of individual flows; packet-counts, byte-counts, duration and so on. Unfortunately, sampling data affects the quality of information – for example, the precise recovery of the flow-duration distribution is affected: shorter flows may be under-sampled while larger flows may be over-sampled. The process of recovering flows-duration may require both a sophisticated model of the traffic-mix and the impact of sampling and while new approaches are routinely proposed, legacy systems means the fundamental problems will remain for some time [102]. To obtain more information from the network, we must measure the packet traffic.

The measurement of network packet traffic directly led to the breaking of a wide range of fundamental assumptions. Examples that informed directly by measurement include assumptions about topology (as we have seen in Section 2). Topologies were once presumed to have strong hierarchies following the original national telephone-companies. However, network topology has been shown to be both complex and poorly understood [103] (see also Section 2.3). Also, perhaps most famously was the early use of high-resolution measurement of Ethernet LAN traffic. [5] established that network packet traffic has self-similar characteristics. This resulted in the undermining of traffic modelled as a Poisson process, and had ramifications for everything from capacity planning and billing to any longer-term understanding of network use. While evaluation of the experience of real-customers is far more recent, it uses a range of tools from the analysis of passive network data [104, 105] through to end-user driven tools of the 'test my bandwidth' type (e.g. http://www.speedtest.net).

Commercial enterprises seek to engineer the most cost-effective networking solution for themselves and their customers. Ironically, commercial sensitivities preclude much direct knowledge. However, the need for performance analysis has led to the development of a range of software to evaluate specific-systems, from an early website tester [106], through to current commercial testers to evaluate video/audio media-stream servers (e.g. http://www.ixiacom.com and http://www.spirentcom.com). In this way, a content-provider may serve their customers better through sufficient servers. We have also seen a number of exciting new research offerings, platforms such as PlanetLab (http://www.planet-lab.org) to allow users to make their own measurement and evaluation among the infrastructure of PlanetLab and as offerings to the wider Internet, for example, Content Distribution (http://codeen.cs.princeton.edu). Other developments have included (research) monitoring infrastructures (http://lobster.ics.forth.gr/~appmon and http://www.internet2.edu/observatory) and even user-installed instrumentation (http://www.netdimes.org). From the early days for user-contributed measurements, it is clear that a range of approaches will be required to overcome the inherent difficulty in measurement.

In their seminal paper, Floyd and Paxson [69] provided several insights into why the Internet is difficult to measure (or simulate) and, thus, resistant to either modelling or predictive insight. Continual evolution means that the measurements of today's networks may not be valid for tomorrow. Global-scale networks affect traffic behaviour: measurements in one place do not represent results at another. Also, applications are rarely designed for measurement, and packet-level measurements often do not relate to application behaviour or performance. The reason for this lies with fundamental properties; however, that does not give us reason to surrender the need for information and thus the need for measurement. A National Academy of Sciences report on research horizons in networking identified (in 2001) the measurement of network infrastructure as a grand-challenge of critical importance for the computing community, in general, the network community, in particular [107] – evolution of speed, size, complexity and use suggests that this challenge remains far from being met.

4.2 Fundamental limits to measurement accuracy

A considerable amount of recent work has concentrated on developing probing techniques [108, 109]. There is also considerable ongoing work probing the core networks of some large networks [110, 111]. There are many variations of the basic probing scheme: single packet probes, packets pairs, triples etc., and variations regarding the patterns of probe delivery – Poisson patterns, deterministic patterns, etc. [112]. By using pairs (or longer string of packets), it has also been claimed that the available network capacity can be reliably estimated [113].

As discussed in Section 4.1, active monitoring by packet probing is essentially a performance sampling system [114]. Filsfil [115], at a Cisco hosted symposium on Measuring Internet Quality reports that magnitude of measurement error is currently very poorly understood. Roughly, the trade-offs are that accuracy increases with the number of samples taken (probes injected), but the overhead (bandwidth used) increases too. The most important limitations of active monitoring are:

- Inject too few probes and the sampling error will be too large – sampling error in probing

- Inject too many and you may adversely affect the performance you are trying to measure – probes interfering with measurements

Additionally, and this is less intuitive, there is an effect caused by the correlation between samples in measuring queuing type performance through queuing systems – error due to correlation between probed samples

Sampling error in probing: Achievable accuracy has been studied for both delay and loss. Obviously, the absolute error in the samples of probing delay through a queuing system increases with the variance of what's being measured – in this case the queuing delay. In [116–118], sampling error has been studied by assuming that the variance in the samples is well enough modelled by the variance of the static distribution of the queuing delays, which can be found using queuing theory. (There is good reason to believe that this may in practice turn out to be a rather conservative estimate for variance [119].)

Numerical results in these references show that measured mean delays can be prone to very considerable error. It has been discovered that even for static traffic in simple buffering scenarios, there are practical load limits beyond which measurement accuracy degrades very rapidly, owing to sampling error. These load limits may be well within the normal operating specification of packet networks, for example, 70% load on a VoIP access link. Precise numerical results are very situation-specific; a schematic representation (showing the fundamentals of the current understanding) is given in Fig. 6. Essentially, the relationship is that increasing load, increasing traffic burstiness and decreasing bandwidth all have the effect of making values returned by probed samples less accurate (the returned mean delays will have a higher absolute error).

Other recent research has addressed optimal packet probing patterns [120], discovering that traditional Poisson sampling is not always optimal. In [120], the authors then developed a more general class of distribution (Gamma renewal processes) that they show should minimise the mean-square error in the sampled data. Roughan [121] also gives a very thorough analysis and discussion of the possibilities available in varying the sampling pattern (rather than just the rate).

Probes interfering with measurements: Increasing the packet probing rate can be an unsatisfactory way to attempt to increase the accuracy of measured results, as the probe stream will interfere with the data traffic, skewing the results [122, 123]. In practice, this is most likely to be a problem only where probes are passing through a low bandwidth bottleneck. In [123], a simple example was analysed, which featured a link speed of 256 kb/s, and packet sizes of 40 and 1000 bytes, corresponding to typical voice and data packets (probes were of the same size as the traffic they were monitoring). Results do show that indeed it is possible to overload low rate links with probes, even with sensible probing rates, potentially damaging the measurement accuracy.

Error caused by correlation between the probed samples: The critical insight here is that packet delays are not actually independent. This can be seen from theory; also measurements of packet delays have shown correlations in practice [124]. So, when analysing the accuracy of mean packet delay estimates, correlations should be incorporated into the model.

This effect can also be visualised as an increase in the 'correlation scale' of the underlying queuing process that is directly related to the length of the busy period (of the packet buffer). The ends of busy periods form 'renewal points', so to measure the delay properties of any queue, it is a good idea to observe it over many such busy periods. However, the busy periods will increase in average duration with increasing load; so (with increasing load) it is sensible that measurements are spaced further apart to compensate. We have attempted to capture this idea in Fig. 7. These ideas are new (within the realm of packet level measurements); the key researcher is Roughan, see [119]

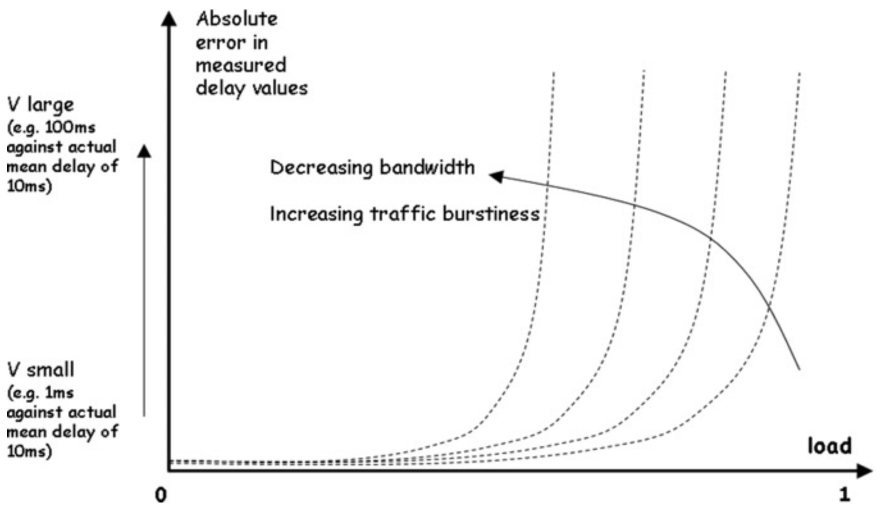

Figure 6 *Schematic representation of effect of decreasing bandwidth, increasing traffic burstiness and/or increasing load on measurement accuracy*

IET Commun., 2009, Vol. 3, Iss. 6, pp. 887–905
doi: 10.1049/iet-com.2008.0111

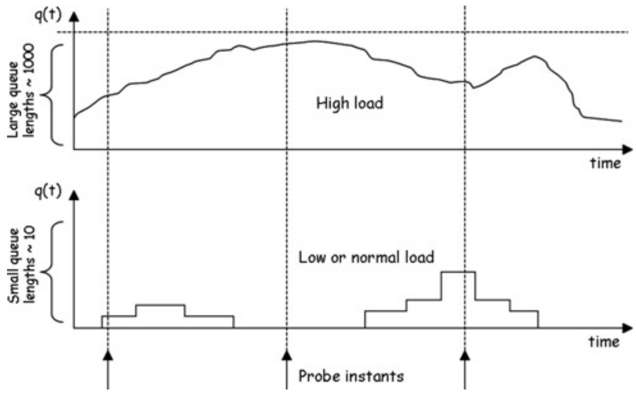

Figure 7 *Schematic representation of increasing probe-to-probe correlation as load increases*

for an explanation of the fundamental bounds that arise because of correlations.

Issues associated with loss probing specifically: In general, it will be harder to measure tail probabilities, for example, for packet loss probability or delay jitter, than it is to measure mean delays. Previous work has in part focused on the Internet [125, 126]. Other authors, [127], have used probed measurements to establish packet loss characteristics in wide area networks, in this case via three different aspects of loss rate based on active probing. In [122], the authors found that simple Poisson or deterministic probing is not well suited to measuring the duration or frequency of buffer overflow events over limited periods (in effect, busy hours); we return specifically to loss probing later. Queuing analysis has been used [128], to show that (for tail-drop queues) probes must be of similar length to the data packets, else the measured loss probabilities may easily be in error by many orders of magnitude.

One significant result to date [114] is that the burstiness of the loss process has a critical effect on measurement, resulting in the need for more samples than would otherwise be expected. In [128], an approximate analytical formula was developed, which relates the number of probes required (to achieve a certain level of absolute error) to aspects of the network and the applied traffic. From this formula, it can be seen that the number of probes required is:

- Inversely proportional to the (packet loss probability)^2
- Proportional to the square of the ratio (data rate of a single traffic source/output rate of the buffer)
- Proportional to $1/(1-\rho)^2$.

These approximate relationships provide an insight into the sensitivity of the measured loss probabilities to network

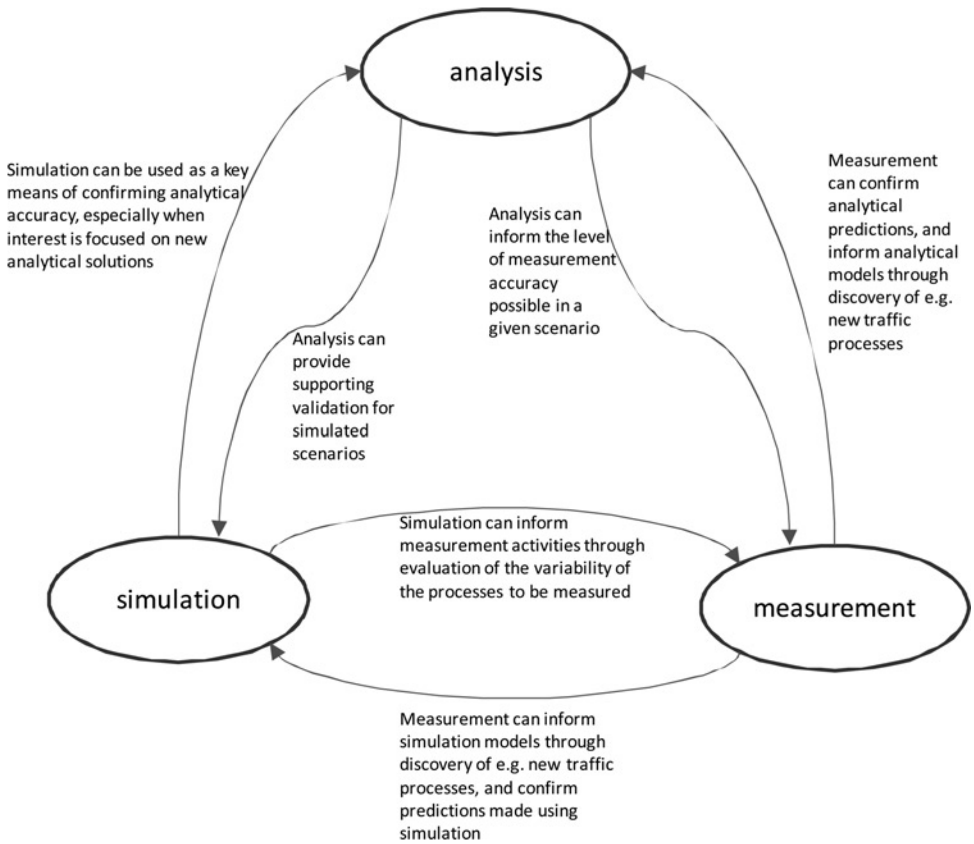

Figure 8 *A non-exhaustive representation of the relationships between analysis, simulation and measurement as used in network performance*

and traffic parameters. Low packet loss probabilities ($<10^{-3}$) may be un-measurable in practical terms if the traffic is bursty and the load is high. This confirms that hotspots should be avoided, as highly loaded nodes will both create loss and make it harder to measure. Low bandwidth links (or low bandwidth VPNs) should also be avoided where possible.

Problems associated with elastic traffic: We have already noted that traffic elasticity, essentially application response to the state of congestion in the rework, causes problems for performance evaluation, whether based on analysis or simulation. It is clear that measurements are also complicated by this. The fundamental bounds noted so far apply to traffic that is not reacting to downstream network conditions. When traffic reacts to the network state, it necessarily creates the circumstances in which the variance of the measured parameters is greatly increased. Indeed, it may then be best to argue that there is no 'static' variance at all.

5 Conclusions

At the end of this review of analysis, simulation and measurement in large-scale packet networks, it is important to see how the three methods may fit together to fully enable network performance evaluation. We try to show this circle of mutual support in Fig. 8 (although the examples given are not exhaustive).

Analysis can be easy to use, and can provide a very general understanding of performance in terms of the included parameters, and a good analytical solution can often be used to facilitate network performance optimisation. For example, in the simplest case, if (say) a buffer overflow probability is linked to traffic only through the load (as in some classical queue models), then the load is very easy to optimise in terms of that parameter (overflow probability). However, the overall network scenario of interest may have to be simplified to achieve analytical tractability, and this can be seen as a compromising rigour.

The next step might be to implement a more detailed simulation of the scenario, including many aspects that an analysis may have had to aggregate, or simply ignore. However, while simulation is routinely used to add detail to analytical studies, analysis is also frequently used to provide confirming validation for a new simulation tool or activity. Following simulation, an existing network (or experimental proto-network) may be examined for detailed measurements, to see how well these measurements support the earlier stages (using analysis and simulation).

However, although this way round seems the most obvious, it is not the only joint approach to performance evaluation. For example, measurement can be used to inform both analysis and simulation through the discovery of new processes (LRD being an obvious one). Also analysis (or simulation) can be used to bound the level of variability inherent in the network process that is to be measured, providing an estimate of for example, the number of samples required in the process of measurement. We have tried to represent a number of different possibilities for mutual interaction and support in Fig. 8.

6 Acknowledgments

The authors thank the anonymous reviewers whose comments have significantly contributed to improving this paper. John Schormans thanks the statisticians Prof. Steven Gilmour and Ben Parker.

7 References

[1] ROBERTS J.W.: 'Traffic theory and the internet', *IEEE Commun. Mag.*, 2001

[2] ROBERTS J.W. (ED.): 'Performance evaluation and design of multiservice networks', Final Report Action COST 224, Commission of the European Communities, Brussels, 1992)

[3] ROBERTS J.W., MOCCI V., VIRTAMO J. (EDS.): 'Broadband network teletraffic, final report of action COST 242', Commission of the European Communities (Springer, Berlin, 1996)

[4] PITTS J.M., SCHORMANS J.A.: 'Introduction to IP and ATM design and performance' (Wiley, 2000)

[5] LELAND W., TAQQU M., WILLINGER W., WILSON D.: 'On the self-similar nature of ethernet traffic (extended version)', *IEEE/ACM Trans. Networking*, 1994, **2**, (1), pp. 1–15

[6] SCHULZRINNE H.: 'The IETF internet telephony architecture and protocols', *IEEE Network*, May/June, 1999

[7] POLYZIOS C.A., PURDY K.H., YANG P.-F., SHRADER D., SINNREICH H., MENARD F., SCHULZRINNE H.: 'From POTS to PANS: a commentary on the evolution to internet telephony', IEEE Internet Comput., May/June 1999

[8] FU X., SCHULZRINNE H., BADRER A., *ET AL.*: 'NSIS: a new extensible IP signalling protocol suite', *IEEE Commun. Mag.*, 2005

[9] SHIN S., SCHULZRINNE H.: 'Experimental measurement of the capacity for VoIP traffic in IEEE 802.11 WLANs'. IEEE INFOCOM 2007

[10] KHO W., BASET S.A., SCHULZRINNE H.: 'Skype relay calls: measurements and experiments'. IEEE INFOCOM 2008

[11] GARRETT M.W., WILLINGER W.: 'Analysis, modeling and generation of self-similar VBR video traffic'. Proc SIGCOMM 1994, pp. 269–280

doi: 10.1049/iet-com.2008.0111

[12] HARPER P.J.: 'IP traffic aggregation: headroom needed for QOS and application to TDM transport'. Paul Julian Harper Turin Networks, 1415 N McDowell Ave, Petaluma, CA 9495, USA. Email: pharper@turinnetworks.com.paper available from http://www.choicecomp.com/TrafficAggregation/IpAggregation.pdf

[13] CAO J., RAMANAN K.: 'A Poisson limit for buffer overflow probabilities'. Proc. INFOCOMM, 2002

[14] WILLINGER W., TAQQU M.S., SHERMAN R., WILSON D.V.: 'Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level', *IEEE/ACM Trans. Networking*, 1997, **5**, (1), pp. 71–86

[15] SHELUHIN O.I., SMOLSKIY, SERGEY, OSIN, REW: 'Self-similar processes in telecommunications' (John Wiley and Sons Ltd, 2007) [NS2-07] The Network Simulator – ns2 (ns-2.32 released 3 Sep 2007) [Online], Available http://www.isi.edu/nsnam/ns

[16] FLOYD S., JACOBSON V.: 'Random early detection gateways for congestion avoidance', *IEEE/ACM Trans. Networking*, 1993, **1**, (4), pp. 397–413

[17] GEVROS P., CROWCROFT J., KIRSTEIN P., BHATTI S.: 'Congestion control mechanisms and the best effort service model', *IEEE Network*, 2001, **15**, (3), pp. 16–26

[18] BONALD T., MAY M., BOLOT J.-C.: 'Analytic evaluation of RED performance', *IEEE INFOCOM*, 2000, **3**, pp. 1415–1424

[19] PITTS J.M., WANG X., YANG Q., SCHORMANS J.A.: 'Excess-rate queueing theory for M/M/1/RED with application to VoIP QoS', *Electron. Lett.*, 2006, **42**, (20), pp. 1188–1189

[20] VERMA D.: 'Supporting service level agreements on IP networks' (MacMillan Technical Publishing, 1999), ISBN 1-57870-146-5

[21] REEVE, BILTON, HOLMES, BROSS: '21CN networks and systems for BT in the 21st century', *IEE Commun. Engineer*, Oct/Nov 2005

[22] FIEMS D., WALRAEVENS J., BRUNEEL H.: 'Performance of a partially shared priority buffer with correlated arrivals'. Book Series Lecture Notes in Computer Science (Springer, Berlin/Heidelberg, 2007), ISSN 0302-9743 (Print) 1611-3349 (Online), vol. 4516/2007, ISBN 978-3-540-72989-1, pp. 582–593

[23] KLEINROCK L.: 'Queueing systems, volume 2' (Wiley-Interscience, 1976, 1st edn.)

[24] MAERTENS T., WALRAEVENS J., BRUNEEL H.: 'A modified HOL priority scheduling discipline: performance analysis', *Eur. J. Oper. Res.*, 2007, **180**, (3), pp. 1168–1185

[25] KUMAR A., MANJUNATH D., KURI J.: 'Communication networking – an analytical approach' (Elsevier, 2004)

[26] PADHYE J., *ET AL.*: 'Modelling TCP throughput: a simple model and its empirical validation'. Proc. SIGCOMM'98

[27] SASTRY N., SOLLINS K., CROWCROFT J.: 'Architecting citywide ubiquitous WiFi access'. The Sixth Workshop on Hot Topics in Networks (HotNETsVI), 2007

[28] BIANCHI G.: 'Performance analysis of the IEEE 802.11 distributed coordination function', *IEEE J. Selected Areas Commun.*, 2000, **18**, (3), pp. 535–547

[29] LEUNG K.-C., LI V.O.K.: 'Transmission control protocol (TCP) in wireless networks: issues, approaches, and challenges', *IEEE Commun. Surveys*, 4th quarter, 2006, **8**, (4), pp. 64–79

[30] CHAKRAVORTY R., KATTIN S., CROWCROFT J., PRATT I.: 'Flow aggregation for enhanced TCP over wide-area wireless', *IEEE INFOCOM*, 2003

[31] ZHU H., LI M., CHLAMTAC I., PRABHAKARAN B.: 'A survey of quality of service in IEEE 802.11 networks', *IEEE Wireless Commun.*, August, 2004, pp. 6–14

[32] MONDRAGON R.J.: 'Topological modelling of large networks', *Philos. Trans. A, Roy. Soc., Lond.*, doi: 10/1098/rsta.2008.0008

[33] CHA M., CHOUDHURY G., YATES J., SHAIKH A., MOON S.: '2006 case study: resilient backbone network design for IPTV services'. Workshop on IPTV Services over WWW, 23 May 2006 (Edinburgh, Scotland, UK, 2006)

[34] FLOYD S., KOHLER E.: 'Internet research needs better models'. Proc. HorNets-I, October 2002

[35] Cooperative Association for Internet Data Analysis (CAIDA): http://www.caida.org/

[36] FALOUTSOS M., FALOUTSOS P., FALOUTSOS C.: 'On power–law relationships of the internet topology', *Comput. Commun. Rev.*, **29**, pp. 251–262 (DOI 10.1145/316188.316229)

[37] FREEMAN M.L.: 'A set of measures of centrality based on betweenness', *Sociometry*, 1977, **40**, (35), pp. 35–41

[38] FUKS H., LAWNICZAK A.T.: 'Performance of data networks with random links', *Math. Comput. Simul.*, 1999, **51**, pp. 103–119

[39] ZHAO L., LAI Y.-C., PARK K., YE N.: 'Onset of traffic congestion in complex networks', *Phys. Rev. E*, 2005, **171**

[40] FUKS H., LAWNICZAK A.T., VOLKOV S.: 'Packet delay in data network models', *ACM Trans. Modelling Comput. Simul.*, 2001, **11**, (3), pp. 233–250

[41] GUIMERA R., DIAZ-GUILERA A., VEGA-REDONDO F., CABRALES A., ARENAS A.: 'Optimal network topologies for local search with congestion', *Phys. Rev. Lett.*, 2002, **89**, (24), pp. 248–701

[42] HOLME P.: 'Edge overload breakdown in evolving networks', *Phys. Rev. E*, 2002, **66**, 036119, 2992

[43] WOOLF M., ARROWSMITH D., MONDRAGON R., PITTS J.M.: 'Optimization and phase transition in a chaotic model of data traffic', *Phys. Rev. E*, 2002, **66**, 056106

[44] PASTOR-SATORRAS R., VAZQUEZ A., VESPIGNANI A.: 'Dynamical and correlation properties of the Internet', *Phys. Rev. Lett.*, 2001, **87**, 258701 (DOI 10.1103/PhysRevLett.87.258701)

[45] PASTOR-SATORRAS R., VESPIGNANI A.: 'Evolution and structure of the internet – a statistical physics approach' (Cambridge University Press), 2004

[46] NEWMAN M.E.J.: 'Assortative mixing in networks', *Phys. Rev. Lett.*, 2002, **89**, 208701 (DOI 10.1103/PhysRevLett.89.208701)

[47] COSTA L., DA F., RODRIGUES F.A., TRAVIESO G., VILLAS BOAS P.R.: 'Characterization of complex networks: a survey of measurements', *Adv. Phys.*, 2007, **56**, pp. 167–242 (DOI 10.1080/00018730601170527)

[48] TANGMUNARUNKIT H., GOVINDAN R., JAMIN S., SHENKER S., WILLINGER W.: 'Network topology generators: degree-based vs. structural'. Proc. 2002 Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications (ACM Press), pp. 147–159 (DOI 10.1145/633025.633040)

[49] ALVAREZ-HAMELIN J.-I., DALL'ASTA L., BARRAT, A., VESPIGNANI A.: 'k-core decomposition: a tool for the visualization of large scale networks', arXiv.org:cs/0504107. 2005

[50] RODRIGUES F.A., VILLAS BOAS P.R., TRAVIESO G., COSTA L., DA F.: 'Seeking the best Internet model', arXiv:0706.3225v1

[51] BARABASI A.-L: 'The architecture of complexity – from network structure to human biology', *IEEE Control Syst. Mag.*, August 2007

[52] SIGANOS G., FALOUTSOS M., FALOUTSOS P., FALOUTSOS C.: 'Power laws and the AS-level Internet topology', *IEEE/ACM Trans.*, Networking, 2003

[53] CHANG H., JAMIN S., WILLINGER W.: 'What causal forces shape internet connectivity at the AS-level?', Technical Report CSE-475-03, EECS Dept., Univ. of Michigan, 2003

[54] BU T., TOWSLEY D.: 'On distinguishing between internet power law topology generators'. Proc. INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE, 2002), vol. 2, pp. 638–647 (DOI 10.1109/INFCOM.2002.1019309)

[55] DOROGOVTSEV S.N., MENDES J.F.: 'Evolution of networks: from biological nets to the internet and www' (Oxford University Press, Oxford, 2003)

[56] KRIOUKOV D., CHUNG F., CLAFFY K.C., FOMENKOV M., VESPIGNANI A., WILLINGER W.: 'The workshop on Internet topology (WIT) report'. ACM SIGCOMM'2007 Computer Communication Review (CCR) 37, 69–73. (http://www.caida.org/publications/papers/2006/wit/wit.pdf)

[57] ZEGURA E., CALVERT K., BHATTACHARJEE S.: 'How to model an Internet network'. Proc. IEEE INFOCOM The Conference on Computer Communications (IEEE, 1996), vol. 2, pp. 594–602

[58] DOAR M.: 'A better model for generating test networks'. In 'Communications: the key to global prosperity'. Proc. Global Telecommunications Conference, GLOBECOM '96, pp. 86–93 (DOI 10.1109/GLOCOM.1996.586131)

[59] WINICK J., JAMIN S.: 'Inet-2.0 internet topology generator', Tech. Report UM-CSE-TR-456-02, Univ. Michigan, Ann Harbour

[60] ZHOU S., MONDRAGON R.J.: 'Accurately modelling the Internet topology', *Phys. Rev. E*, 2004, **70**, 066108 (DOI 10.1103/PhysRevE.70.066108)

[61] WILLINGER W., GOVINDAN R., JAMIN S., PAXSON V., SHENKER S.: 'Scaling phenomena in the Internet: critically examining criticality', *Proc. Natl. Acad. Sci. USA*, 2002, **99**, pp. 2573–2580 (DOI 10.1073/pnas.012583099)

[62] GOLDSMAN D.: 'Introduction to simulation'. Proc. 2007 Winter Simulation Conf.

[63] SARGENT R.: 'Verification and validation of simulation models'. Proc. 2007 Winter Simulation Conf.

[64] LAW A.M.: 'Statistical analysis of simulation output data: the practical state of the art'. PROC. 2007 Winter Simulation Conference

[65] http://www.emulab.net/

[66] http://www.caida.org/projects/ditl/

[67] http://www.caida.org/workshops/wide/0801/

[68] Opnet Technologies, Inc. (2008), Opnet Modeler, [Online], Details available: http://www.opnet.com/

[69] FLOYD S., PAXSON V.: 'Difficulties in simulating the Internet', *IEEE/ACM Trans. Networking*, 2001, **9**, (4), pp. 392–403

[70] RILEY G., AMMAR M.: 'Simulating large networks: how big is big enough?'. Proc. First Int. Conf. on Grand Challenges for Modelling and Simulation, January 2002, San Antonio

[71] RILEY G.F.: 'Large-scale network simulation: how big? how fast?'. 11th IEEE/ACM Int. Symp. on Modelling, Analysis and Simulation of Computer Telecommunications Systems, 2003, MASCOTS 2003, 12–15 October 2003, pp. 116–123

[72] KUHL F., DAHMANN J., WEATHERLY R.: 'Creating computer simulation systems: an introduction to the high level architecture' (Prentice Hall, 1999), ISBN 0130225118

[73] LIU B., FIGUEIREDO D.R., GUO Y., KUROSE J., TOWSLEY D.: 'A study of networks simulation efficiency: fluid simulation vs. packet-level simulation', INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. *Proc. IEEE*, 2001, **3**, pp. 1244–1253

[74] PITTS J.M.: 'Cell-rate modelling for accelerated simulation of ATM at the burst level', *IEE Proc. Commun.*, 1995, **142**, (6), pp. 379–385

[75] BRESLAU L., ESTRIN D., FALL K., *ET AL.*: 'Advances in network simulation', *Computer*, 2000, **33**, (5), pp. 59–67

[76] MA A., SCHORMANS J., CUTHBERT L.: 'Aggregation technique for networks with power law traffic and application to accelerated simulation', *IEE Proc. Commun.*, 2003, **150**, (3), pp. 177–183

[77] LIEBEHERR J., SHROFF N. (EDS.): 'NSF workshop on fundamental research in networking', [Online], 2003, Available: http://www.cs.virginia.edu/~jorg/workshop1/

[78] HEIDEMANN J., MILLS K., KUMAR S.: 'Expanding confidence in network simulation', *IEEE Network Mag.*, 2001, **15**, (5), pp. 58–63

[79] ZHOU S., MONDRAGON R.J.: 'Towards modelling the internet topology – the interactive growth model'. 18th International Teletraffic Congress, Berlin, September 2003

[80] FUKUDA K., TAKAYASU H., TAKAYASU M.: 'Origin of critical behaviour in Ethernet traffic', *Phys. A*, 2000, **287**, pp. 289–301

[81] MATHIS M., SEMKE J., MAHDAVI J.: 'The macroscopic behaviour of the TCP congestion avoidance algorithm', *ACM SIGCOMM Comput. Commun. Rev.*, 1997, **27**, (3)

[82] HERMAN I., MELANCON G., MARSHALL M.S.: 'Graph visualization and navigation in information visualization: a survey', *IEEE Trans. Visualization and Computer Graphics*, 2000, **6**, (1), pp. 24–43

[83] GARVELS M.J.J., KROESE D.P.: 'A comparison of RESTART implementations'. Simulation Conf. Proc., Winter, vol. 1, 13–16 December 1998, pp. 601–608

[84] MA A., SCHORMANS J.A.: 'Fast, stable simulation of power-law packet traffic using concatenated acceleration techniques', *IEE Proc. Commun.*, 2005, **152**, (4), pp. 420–426

[85] The ns-3 project, [Online], Details available http://www.nsnam.org/

[86] ANDREW L., MARCONDES C., FLOYD S., *ET AL.*: 'Towards a common TCP evaluation suite'. Proceedings of International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet), March 2008

[87] The e-model, a computational model for use in transmission planning. ITU-T Recommendation G.107. May 2000

[88] YANG Q., PITTS J.M.: 'Guaranteeing enterprise VoIP QoS with novel approach to DiffServ AF configuration'. IEEE Int. Communication Conf. (ICC'07), Glasgow, UK, 24–28 June 2007

[89] LU S., SCHORMANS J.A.: 'SSIM: a multi-resolution fluid traffic simulator for MANETs'. 3rd Int. Conf. on Mobile Computing and Ubiquitous Networking (ICMU2006), October 2006

[90] LU S., SCHORMANS J.A.: 'A time-stepped approach for accelerated simulation of mobile ad-hoc networks', *IET Commun.*, 2008, **2**, (5), pp. 609–620

[91] AMARADASA V., SCHORMANS J.A.: 'Accelerated simulation of a GPS scheduler using traffic aggregation'. 4th Int. Working Conf., Performance Modelling and Evaluation of Heterogeneous Networks (HETNETs), 2006, Ilkley, Yorkshire, UK

[92] CROVELLA M., KRISHNAMURTHY B.: 'Internet measurement: infrastructure, traffic and applications' (Wiley, 2006)

[93] MISLOVE A., MARCON M., GUMMADI K.P., DRUSCHEL P., BHATTACHARJEE B.: 'Measurement and analysis of online social networks'. Proc. 7th ACM SIGCOMM Conf. on Internet Measurement, San Diego, California, USA, 24–26 October 2007

[94] CHA M., KWAK H., RODRIGUEZ P., AHN Y., MOON S.: 'I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system'. Proc. 7th ACM

SIGCOMM Conf. on Internet Measurement, San Diego, California, USA, 24–26 October 2007

[95] SVOBODA P., KARNER W., RUPP M.: 'Traffic analysis and modeling for world of warcraft'. IEEE Int. Conf. on Commun., 24–28 June 2007

[96] KARAGIANNIS T., BROIDO A., BROWNLEE N., CLAFFY K.C., FALOUTSOS M.: 'Is P2P dying or just hiding?'. IEEE Globecom, Dallas, TX, Nov 2004

[97] ZHANG Y., ROUGHAN M., LUND C., DONOHO D.L.: 'Estimating point-to-point and point-to-multipoint traffic matrices: an information-theoretic approach', *IEEE/ACM Trans. Networking*, 2005, **13**, (5), pp. 947–960

[98] LAKHINA A., CROVELLA M., DIOT C.: 'Mining anomalies using traffic feature distributions'. ACM SIGCOMM 2005, Philadelphia, September 2005

[99] HUFFAKER B., PLUMMER D., MOORE D., CLAFFY K.C.: 'Topology discovery by active probing'. Proc. IEEE Computer Society Symp. on Applications and the Internet (Saint) Workshops, Washington, DC, 2002

[100] VARDI Y.: 'Network tomography: estimating source-destination traffic intensities from link data', *J. Am. Stat. Assoc.*, 1996, **91**, (433), pp. 365–377

[101] DUFFIELD N.G.: 'Network tomography of binary network performance characteristics', *IEEE Trans. Inform. Theory*, 2006, **52**, pp. 5373–5388

[102] ESTAN G., KEYS K., MOORE D., VARGHESE G.: 'Building a better NetFlow'. Proc. SIGCOMM '04 Porkland, Oregon

[103] OLIVEIRA R., PEI D., WILLINGER W., ZHANG B., ZHANG L.: 'In search of the elusive ground truth: the internet's AS-level connectivity structure'. In Proc. ACM SIGMETRICS'08, May 2008

[104] ANEROUSIS N., CACERES R., DUFFIELD N., *ET AL*.: 'Using the AT&T Labs PacketScope for internet measurements, design, performance analysis'. AT&T Services and Infrastructure Performance Symp., November 1997

[105] MOORE A., HALL J., KREIBICH C., HARRIS E., PRATT I.: 'Architecture of a network monitor'. Passive & Active Measurement Workshop 2003 (PAM2003), 2003

[106] BARFORD P., CROVELLA M.: 'Generating representative web workloads for network and server performance evaluation'. Proc. ACM SIGMETRICS '98, 5 May 1997

[107] PATTERSON D.A., CLARK D.D., KARLIN A., *ET AL*.: 'Computer Science and Telecommunications Board Committee on Research Horizons in Networks – looking over the fence at networks: a neighbor's view of network research' (National Academy of Sciences, Washington, DC, 2001)

[108] PASZTOR A., VEITCH D.: 'A precision infrastructure for active probing'. Proc. PAM2001 Workshop on Passive & Active Measurements, Amsterdam, April 2001

[109] SHANLUNOV S., TEITELBAUM B., KARP A., BOOTE J.W., ZEKAUSKAS M.: 'A one-way delay measurement protocol (OWAMP)', Internet Draft, May 2003

[110] FRALEIGH C., MOON S., LYLES B., *ET AL*.: 'Packet-level traffic measurements from the sprint IP backbone', *IEEE Network*, November/December 2003, **17**, (6), pp. 6–16

[111] JOHN W., TAFVELIN S: 'Analysis of internet backbone traffic and header anomalies observed'. ACM Internet Measurement Conf., 2007

[112] PÁSZTOR A., DARRYL VEITCH D.: 'Active probing using packet quartets'. Proc. 2nd ACM SIGCOMM Workshop on Internet measurement, Marseille, France, 2002, pp. 293–305

[113] DOVROLIS C., RAMANATHAN P., MOORE D.: 'Packet-dispersion techniques and a capacity-estimation methodology', *IEEE/ACM Trans. Networking*, 2004, **12**, (6), pp. 963–977

[114] BARFORD P., SOMMERS J: 'Comparing probe-based and router-based packet-loss measurement', *IEEE Internet Comput.*, 2004, **8**, (5), pp. 50–56

[115] FILSFILS C., WARD, D.: 'Measuring internet quality', see http://www.cisco.com/web/about/ac50/ac207/proceedings/MIQ.ppt#486,4, Application View (seen 05-02-2007)

[116] SCHORMANS J.A., TIMOTIJEVIC T.: 'Evaluating the accuracy of active measurement of delay and loss in packet networks'. 6th IFIP/IEEE Int. Conf. on Management of Multimedia Networks and Services, September 2003, Queens University, Belfast

[117] TIMOTIJEVIC T., SCHORMANS J.A.: 'Bandwidth overhead of probe technology guaranteeing QoS in packet networks', *IEE Electronics Lett.*, 2003, **39**, (10), pp. 816–818

[118] TIMOTIJEVIC T., LEUNG C.M., SCHORMANS J.A.: 'Accuracy of measurement techniques supporting QoS in packet based Intranet and Extranet VPNs', *IEE Proc. Communications Special Edition on VPNs and Broadband Technology*, 2004, **151**, (1), pp. 89–94

[119] ROUGHAN M.: 'Fundamental bounds on the accuracy of network performance measurements'. ACM SIGMETRICS'05, 6–10 June 2005, Banff, Canada

[120] BACCELLI F., MACHIRAJA S., VEITCH D., BOLOT J.: 'On optimal probing for delay and loss measurement'. Proc. IMC'07, San Diego, California, October 2007

doi: 10.1049/iet-com.2008.0111

[121] ROUGHAN M.: 'A comparison of Poisson and uniform sampling for active measurements', *IEEE Selected Areas Commun.*, 2006, **24**, (12), pp. 2299–2312

[122] SOMMERS J., BARFORD P., DUFFIELD N., RON A.: 'Improving accuracy in end to end packet loss measurements'. SIGCOMM'05, 21–26 August, Philadelphia, USA

[123] SCHORMANS J.A., PITTS J.M.: 'So you think you can measure IP QoS?'. The 2nd IEE Telecommunications QoS Conference 2004, Savoy Place, 2–3 March 2004

[124] ANDREN J., HILDING M., VEITCH D.: 'Understanding end-to-end internet traffic dynamics'. Proc. Globecom'98, Sydney, Australia, 1998, pp. 1118–1122

[125] BOLOT J.: 'End-to-end packet delay and loss behavior in the Internet'. Proc. ACM SIGCOMM '93, San Francisco, September 1993

[126] PAXSON V.: 'End-to-end internet packet dynamics'. Proc. ACM SIGCOMM '97, Cannes, France, September 1997

[127] ZHANG Y., DUFFIELD N., PAXSON V., SHENKAR S.: 'On the constancy of Internet path properties'. Proc. ACM SIGCOMM Internet Measurement Workshop '01, San Francisco, November 2001

[128] HASIB M., SCHORMANS J.A., TIMOTIJEVIC T.: 'Accuracy of packet loss monitoring over networked CPE', *IET Commun.*, 2007, **1**, (3), pp. 507–513

doi: 10.1049/iet-com.2008.0111