Fixed-Point Models for the End-to-End Performance Analysis of IP Networks

R. J. Gibbens†, S. K. Sargood, C. Van Eijl, F. P. Kelly†,
H. Azmoodeh, R. N. Macfadyen, N. W. Macfadyen

† Statistical Laboratory,
University of Cambridge, 16 Mill Lane, Cambridge, UK
Contact: R.J.Gibbens@statslab.cam.ac.uk
+44 (1223) 337945

British Telecommunications plc,
Adastral Park, Ipswich, UK
Contact: henny.azmoodeh@bt.com
+44 (1473) 647462

Abstract

This paper presents a new approach to modeling end-to-end performance for IP networks. Unlike earlier models, in which end stations generate traffic at a constant rate, the work discussed here takes the adaptive behaviour of TCP/IP into account. The approach is based on a fixed-point method which determines packet loss, link utilization and TCP throughput across the network. Results are presented for an IP backbone network, which highlight how this new model finds the natural operating point for TCP, which depends on route lengths (via round-trip times and number of resources), end-to-end packet loss and the number of user sessions.

1. Introduction

Internet Service Providers are now deploying large scale IP networks which offer a range of Quality of Service (QoS) metrics. These are either part of a Service Level Agreement, or can simply provide an indication of the network performance to customers [1,2]. Future developments of QoS will extend to multiple traffic classes, and there is therefore a growing need for end-to-end performance models to determine vital QoS metrics such as link utilisation, packet loss, one-way or round-trip delay, and throughput of TCP sessions.

The models presented in this work build on initial work in [3], where a fixed-point model was proposed for evaluating QoS on IP networks with Differentiated Services. Fixed-point models have been widely used in circuit switched networks in order to determine link utilisation and blocking probabilities for connections arriving at a resource [4]. As shown in [3], a similar approach can be applied in IP networks to simultaneously compute link utilisation and packet loss for all resources in the network.

Performance models for IP networks, including the one in [3], usually assume that end stations generate traffic at a constant rate. This open source model is representative of UDP traffic which generally does not back off in the event of packet loss (though applications are now being developed to be TCP-friendly and react favourably to congestion on the network [5]). However, most traffic on today's IP networks uses TCP/IP, which does react to packet loss by reducing its sending rate. This paper presents a fixed-point model for end stations that operate in TCP congestion avoidance mode and, hence, adapt their sending rate to the congestion in the network. The fixed-point is then a result of TCP sessions finding a natural level as the offered load to the network. This turns out to be dependent on several factors including the packet loss over a route, the route length (via the round trip time and the number of resources) and the mean number of TCP sessions supported.

Within the fixed-point method an appropriate single resource model should be used to determine link utilisation and packet loss per resource. The resource model used in this paper is based on [6], where resources are modelled as finite packet buffers or M/M/1/K queues, in order to evaluate the performance of IP Differentiated Services with two traffic classes. In [6] these are based on the Assured Forwarding Per-Hop-Behaviours [7], and are referred to as a Premium data class and a Best Effort data class, and are mapped to specified settings of the IP DSCodepoint [8]. In this work they will be referred to simply as high and low priority traffic.

An important assumption for the validity of a fixed-point approach is that traffic correlation across resources is negligible [9,10]. This independence assumption is usually satisfied when the network topology is well-connected, and traffic demands are aggregated such that the offered traffic to a resource can be modelled as a Markovian arrival process. Although it is now well established that traffic on both LAN and WAN IP networks can exhibit long-range dependent behaviour with both space and time scaling [11,12], suitable models to account for this are not readily available, or are too complex to be used in a performance model without resort to simulation. The critical importance of time-scale issues is discussed further in [13,14]. In [13] a case is considered where short-term correlation was found to have the dominant effect on performance. In [14] it is stressed that the relevant time-scale of interest will depend on the input traffic and, in addition, the properties of the network elements, such as finite buffer queues, and on the performance metric of interest. Therefore, the single resource model from [6], which assumes Poisson arrival statistics, is considered as a good starting point for the approach presented here.

This work attempts to meet the principal goals for "simulating" the Internet discussed in [15,16]: that of dealing with a large heterogeneous closed loop system by incorporating the adaptive behaviour of end-stations, and using appropriate resource models combined with the fixed-point approach to relate the packet and session levels respectively of the network. An interesting observation has been that, for the adaptive source case, the convergence of the fixed-point approach is highly sensitive to the buffer size used in the resource model, since this packet level feature has a direct impact on the dynamics of the closed loop system, and on traffic correlation over periods of the order of a round-trip time.

The rest of the paper is organised as follows. Section 2 discusses the fixed-point method both for the open and adaptive source case, and describes the encapsulation of the resource model of [6] in the fixed-point equations. A discussion of key issues and aspects of calibration in the modelling approach is also given. The working of the fixed-point model with adaptive sources is illustrated with a simple example. Section 3 discusses results for both approaches for a typical IP backbone, while Section 4 summarises and considers further work.

2. Fixed-Point Equations and Resource Models

2.1 Fixed-Point Equations with Open Source

Suppose that the offered arrival rates for high and low priority traffic streams on a directed route r of the network are α_{1r} and α_{2r} respectively. Then given loss probabilities L_{1j} and L_{2j} at each resource ($j = 1,2,...J$), and assuming independence, the reduced load at resource j is given by

$$\nu_{1j} = \sum_{r \ni j} \alpha_{1r} \prod_{j' \in r : p(j',r) < p(j,r)} (1 - L_{1j'}) \tag{1}$$

$$\nu_{2j} = \sum_{r \ni j} \alpha_{2r} \prod_{j' \in r : p(j',r) < p(j,r)} (1 - L_{2j'}) \tag{2}$$

where $p(j,r)$ gives the position of resource j on route r . The losses are generally dependent on ν_{1j} and ν_{2j} and can be expressed using

$$L_{1j}^{'} = L_{1j}\left(\nu_{1j}, \nu_{2j}\right) \tag{3}$$

$$L_{2j}^{'} = L_{2j}\left(\nu_{1j}, \nu_{2j}\right) \tag{4}$$

Repeated iteration (possibly damped) may be used to determine the fixed-point.

2.2 Fixed-Point Equations with Adaptive Source

Since the sources in this case generate not a given offered arrival rate, but one which depends on the end-to-end packet loss probability over a route, the fixed-point equations (1) and (2) must be solved with an adaptive load generated by each of the sources. This requires that the relationships be modelled in a different manner, with the mean number of TCP sessions being used as an input to determine the offered load, which is then solved self-consistently with the resource model for packet loss and utilisation over a given route. Before detailing the fixed-point approach used here, the key aspects of TCP protocol needed will be briefly discussed.

Several functional forms have been derived for the TCP source rate dependence on end-to-end packet loss probability, and two examples are

$$S\left(p_{ir}; T_{ir}\right) = \frac{1}{T_{ir}\sqrt{2 p_{ir} / 3}} \tag{5}$$

where T_{ir} is the round-trip time on route r , and which will be familiar in form to the classical TCP equation for throughput in congestion avoidance mode [17], and

$$S\left(p_{ir}; T_{ir}; T_0; W_{\max}; R_r\right) =$$
$$\min\left\{\frac{W_{\max}}{T_{ir}}, R_r, \frac{1}{T_{ir}\sqrt{2 p_{ir} / 3} + T_0 \min\left\{1, 3\sqrt{3 p_{ir} / 8}\right\} p_{ir} f\left(p_{ir}\right)}\right\} \tag{6}$$

with $f\left(p_{ir}\right) = 1 + p_{ir} + 2p_{ir}^{2} + 4p_{ir}^{3} + 8p_{ir}^{4} + 16p_{ir}^{5} + 32p_{ir}^{6}$

T_0 is the TCP timeout value and $W_{\max}$ is the maximum receive window size. R_r is a rate limit which may apply to a route r to limit the throughput of a TCP connection, e.g. this may be imposed by the bandwidth of the circuit connecting to the access router in the network. This equation also captures the important feature that above packet losses of ~1% TCP sessions in congestion avoidance become increasingly likely to suffer packet losses that result in timeouts rather than triple-duplicate ACKS [18].

For the fixed-point model the determination of the offered load to the network must be provided by considering the number of TCP sessions, which are then allowed to reach their natural operating state in congestion avoidance. Suppose that there are n_{ir} sessions present of type i on route r, and that the total offered load α_{ir} is related to p_{ir}, the end-to-end packet loss probability for class i traffic on route r, and to T_{ir} the round-trip time for a packet on the same route by

$$\alpha_{ir} = n_{ir} S(p_{ir}; T_{ir}) \tag{7}$$

From the independence assumption it follows that end-to-end packet loss probability p_{ir} is given by

$$1 - p_{ir} = \prod_{j \in r} (1 - L_{ij}) \tag{8}$$

where L_{ij} is the loss at resource j for traffic of type i.

In this paper, the parameter n_{ir} is generated as a vector of identically distributed independent random variables drawn from a Poisson distribution, with a typical sample size of ~1000 vectors. We note that the Poisson distribution for n_{ir} would arise from the M/G/∞ model for TCP sessions suggested by Paxson and Floyd [12]. For each sample vector the fixed-point technique relates the offered load at each resource to the packet loss probabilities for each class of traffic, and solves via a damped fixed-point calculation. The technique assumes that the packet loss along a route is cumulatively thinned as in (8).

Note that the M/G/∞ model for TCP sessions that was suggested in [12] is able to generate the long-range dependence in the number of sessions present observed by the authors over time scales of 0.1 seconds and longer. An M/G/1 processor sharing model has also been suggested for TCP sessions [19]: this model gives rise to a geometrically distributed number of TCP sessions and perhaps a much larger coefficient of variation for n_{ir}. Different models correspond to different assumptions about user behaviour. The M/G/1 model corresponds to TCP file transfers that slow down when congestion is encountered, but then take longer. The M/G/∞ model corresponds to users that reduce, rather than merely postpone, their aggregate load on the network at times of congestion.

2.3 Resource Models

Following an approach in [6], the resource model (at the packet level) is assumed to be that of an M/M/1/K queue with finite buffer size B, which rejects on arrival a low priority packet if there are T or more packets already in the buffer. Let i be the occupancy of the buffer, and suppose that arrival rates for high and low priority streams are ν_1 and ν_2 respectively. If the resource serves packets at the rate of C packets per second, then the state i can be modelled using a Markov chain with transition rates

$$\begin{aligned} q(i, i+1) &= (\nu_1 + \nu_2 : 0 \le i < T; \nu_1 : T \le i < B) \\ q(i, i-1) &= (C : 1 \le i \le B) \end{aligned} \tag{9}$$

The equilibrium distribution for the state i is given by

$$\pi_i = \pi_0 \prod_{k=1}^{i} \frac{q(k-1, k)}{q(k, k-1)} \tag{10}$$

where

$$\pi_0 = \left[1 + \sum_{i=1}^{B} \prod_{k=1}^{i} \frac{q(k-1, k)}{q(k, k-1)} \right]^{-1}. \tag{11}$$

The loss probabilities for high and low priority streams are then given by

$$L_1 = L_1(\nu_1, \nu_2) = \pi_B \tag{12}$$

$$L_2 = L_2(\nu_1, \nu_2) = \sum_{i=T}^{B} \pi_i \tag{13}$$

For a FIFO service discipline at the server, the mean delay of a packet accepted into the buffer when there are n packets already present is the sum of $(1+n)$ independent exponential random variables, each with mean duration C^{-1}. Thus the expected delays at a given resource for accepted high and low priority packets are given by (with a corrected normalisation factor to omit lost packets, amended from May [6])

$$E[D_1(\nu_1, \nu_2)] = \frac{\sum_{n=0}^{B-1} (1+n) \pi_n}{C \sum_{n=0}^{B-1} \pi_n} \tag{14}$$

$$E[D_2(\nu_1, \nu_2)] = \frac{\sum_{n=0}^{T-1} (1+n) \pi_n}{C \sum_{n=0}^{T-1} \pi_n} \tag{15}$$

Note that low priority packets are less likely to be accepted into the buffer. However, the ones that are accepted experience lower mean delays than accepted high priority packets, i.e.

$$L_2(\nu_1,\nu_2) \geq L_1(\nu_1,\nu_2) \tag{16}$$

but

$$E[D_1(\nu_1,\nu_2)] \geq E[D_2(\nu_1,\nu_2)] \tag{17}$$

In this work the resource model used has been the M/M/1/K queue (finite buffer) for both traffic classes, but the key assumption of independence allows flexibility here. Additionally, buffer management policies with a Weighted Random Early Detection (WRED) acceptance profile can also be used. This defines minimum and maximum thresholds between which the acceptance probability of packets arriving to the resource is a (linearly) increasing function of buffer position, and performance analysis of these is detailed in [6,20]. However, the simple drop threshold scheme adopted here provides adequate distinction between high and low priority packets, particularly under conditions of high offered load where this distinction is required. An alternative resource model which is of current interest uses Class Based Weighted Fair Queuing [21], where traffic in each class is fed into separate buffers. The service rate allocated to each class is weighted, and each buffer also has an acceptance profile.

It is well understood that it is difficult to cope with long-range dependence by packet level buffering. On the other hand, buffering is extremely helpful in coping with short time scale packet level fluctuations caused by the chance superposition of packets from different streams. Further, the amount of buffering provided directly influences the form of the packet-level traffic offered to the network, via the interaction of packet drops with TCP's congestion avoidance algorithm. Note also that as the threshold for low priority packets increases, so does the average delay for high priority packets. In this study we have used relatively small thresholds: this sacrifices some resource utilisation, but reduces both delays and the possibilities for oscillatory behaviour.

2.4 Calibration

The fixed-point technique coherently integrates three simpler models that represent respectively: random fluctuations in n_{ir}, the number of sessions competing for use of the network; the packet level stochastic behaviour of queues in the resource model stated initially in equations (3) and (4) and developed therein; and the adaptive behaviour of TCP captured in (5) and (6).

The three models integrated by the fixed-point technique however are each capable of calibration from measurement. The key parameter of n_{ir} is its coefficient of variation, which may depend upon various aspects of user behaviour or the service discipline. The resource model of May [6] may be parameterised by the size of bursts of packets [20]. Finally the load on a route may come from open sources, such as short transfers completed by TCP sessions in slow start, as well as from those in congestion avoidance: the proportion of load from each is the key measurement needed to calibrate the model.

Currently sessions which last only a few round-trip-times are known to be an important (and currently dominant) class of traffic on the Internet today [22,23]; a direct result of users' browsers downloading Web pages using the HTTP/1.0 protocol. This application level protocol launches ~4-6 TCP sessions concurrently, with a single TCP session for each object of the Web page. These sessions largely remain in slow-start and terminate before they are able to reach congestion avoidance phase. However future versions of HTTP/1.1 incorporate amendments to use pipelining and persistent sessions [24], which will provide better throughput due to a single persistent TCP session to download all objects on a Web page. The performance of both long [18] and short file transfers [25] has been addressed for TCP sessions; the fixed-point models reported in this work can therefore provide an *upper-estimate* of TCP throughput which may be achieved.

2.5 Example Network with Adaptive Sources

The fixed-point model is applied to a simple network example to illustrate the method in a manner which has a straightforward solution. The network consists of two resources (shown as circles) and three routes, with a rate limit R_1 applied to sessions on route 1.

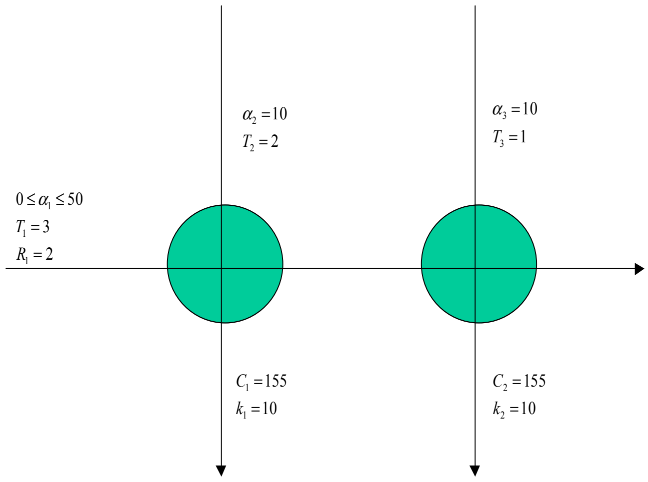

Figure 1. Network with two resources and three routes, with rate limit for sessions on route 1.

The fixed-point equations for this network are given by

$$p_1 = \Pr\left(n_1 \min\left\{ R_1, \frac{1}{T_1\sqrt{2(p_1+p_2)/3}} \right\} + \frac{n_2}{T_2\sqrt{2p_1/3}}; C_1; k_1 \right) \quad (18)$$

$$p_2 = \Pr\left(n_1 \min\left\{ R_1, \frac{1}{T_1\sqrt{2(p_1+p_2)/3}} \right\} + \frac{n_3}{T_3\sqrt{2p_2/3}}; C_2; k_2 \right) \quad (19)$$

For this example the packet loss probabilities are summed over routes (for simplicity). They are determined at each of the resources together with the 5-percentiles of throughputs on each of the routes, using a Monte Carlo technique. The mean number of sessions α_1 on route 1 is varied to determine each point of the graph, using 1000 "fixed-point computations" each of the network. In Figure 2 the packet loss probability is seen to increase with mean number of sessions for both resources, resource 2 experiencing a higher packet loss than resource 1. In Figure 3 the rate limit for connection throughput on route 1 applies when the mean number of sessions is sufficiently low, but as the mean number of sessions increases, the throughputs for all routes decrease. Route 3 achieves the highest throughput due to the lower value of round trip time T_3.

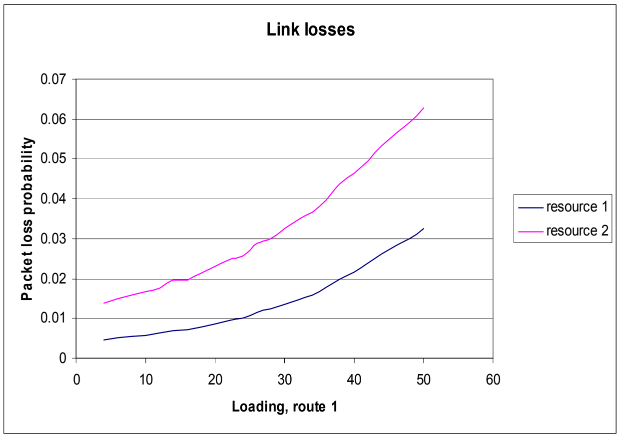

Figure 2 Packet Losses for resources

3. Modelling IP Backbone Network Topologies

In this section the fixed-point models are applied to the performance analysis of network topologies representative of IP backbones, and this is achieved by incorporating the techniques into design tools for these topologies which also determine the routes. The open source model is considered first and results are used as a reference for the network performance (link utilisation and packet loss for each of the traffic classes) when there is a "bulk" traffic demand. The TCP adaptive source model then considers network performance (with TCP throughput as the additional performance measure) with and without rate limits. Rate limits are introduced by considering dial-up users with 56K modems, and broadband users with 1Mbps modems, though any access rate limit can be used. The main objective is to determine the performance just as the "congested" links are beginning to show sufficiently high packet losses that capacity upgrades would be required. Initial results for this network topology are given which provide an indication of the general flexibility of this approach, while further work to examine robustness and sensitivity is needed, including validation with network measurements.

In these IP networks, links which experience transient traffic bursts and cause link utilisation to exceed unity would cause packet loss, and because of the adaptive nature of TCP protocol in congestion, this would result in TCP sessions backing-off (with Reno remaining in congestion avoidance if a triple duplicate ACK is received by the source, or backing off to slow-start in the case of a timeout), thereby bringing the utilisation back down to a value below unity. While the open-source model self-consistently captures the relation between link utilisation and packet loss, the adaptive source model is necessary to properly account for TCP behaviour, and the actual network performance turns out to be sensitive to many parameters including the mean number of sessions, the round-trip time, time-out value, number of resources along a route, the end-to-end packet loss probability and also the buffer size.

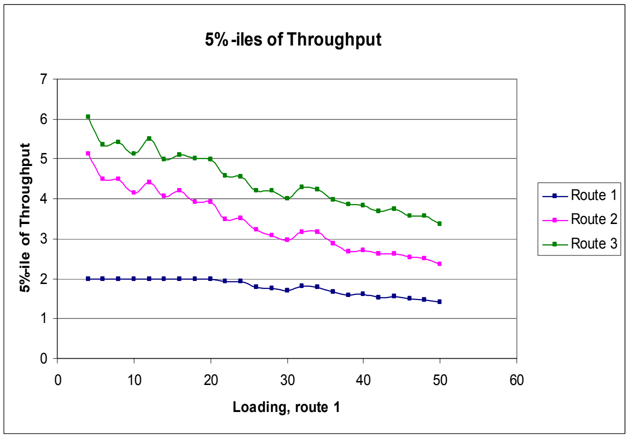

Figure 3. Throughputs (5%iles) for the Routes

The example network topology is given in Figure 4. It consists of 18 core routers, 9 access routers and 86 uni-directional links (including 100Mbps LAN access-core router interconnects, and 'long-distance' core router - core router WAN links which are STM-1s). The topology reflects a sparse-mesh which has Points of Presence in both the USA and Europe, with core router back-up. The routing between access routers for traffic demands is computed off-line. There are 41 routes (between access routers) with a hypothetical and non-zero offered traffic load varying from 7.5 to 60 Mbps, and the split between high and low priority traffic is 30:70. The traffic demands on the network are used differently in the open and adaptive source cases, in the latter there is mean connection bandwidth which is used to determine the number of TCP sessions, and is based here on a mean throughput of 300kbps for high priority traffic and 100kbps for low priority traffic. From these indicative throughputs and the hypothetical traffic demand, the mean number of sessions is determined, from which 1000 samples are taken based on a Poisson distribution, as detailed in Section 2. Round trip delays for a route are taken to be the total propagation delays, and are therefore the same for both traffic classes, since interface bandwidths are sufficiently high that propagation delays dominate router queuing delays. For links of 2Mbps and below, queuing delays will become the more important component of the overall round-trip delay (RTD) and can be easily accomodated into the fixed -point equations. Round-trip delays for routes vary from 6ms to 130ms. Time-out values are taken to be ~8*RTD as an estimate over measurement data in [18], but can be independently set. Buffers are assumed to be 32 packets, with a mean packet size of 500Bytes, and the low priority traffic discard threshold is taken to be 16.

3.1 Open Source Model

The results for the low priority traffic with the open source model show that overall for this traffic demand, link utilisation and packet loss are relatively low on most links. A small number (only 3) "congested" links contribute packet loss >1% which, through the thinning effect along the routes, causes route losses of up to ~30% for low priority traffic between Washington and European nodes. Packet losses otherwise are very low for other low priority traffic and all high priority traffic over the routes.

3.2 Adaptive Source Model

In this model the network topology used is the same as for the open-source model, and the hypothetical traffic demands are indicative of users accessing a server located near to the principal PoPs. This demand is re-cast into the mean number of user-sessions given a defined mean session rate for high and low priority traffic, from which the TCP throughput of a long-lived session which has reached congestion avoidance can be determined. This TCP throughput may itself be access limited and the model in Equation (6) allows access limits to be introduced as illustrated in the Example network. Three cases are compared: the open-source model, TCP without rate limits (i.e. set to infinity), and TCP with a rate limits of 1Mbps.

Figure 4 shows the links on the network which have the highest utilisation for the adaptive source case, where packet loss for low priority traffic exceeded 1% and for these links utilisation exceeded 80%. The congested directions of the links are indicated with arrows. Losses for high priority traffic in almost all cases were negligible due to the buffer discard scheme used. The following were observed:

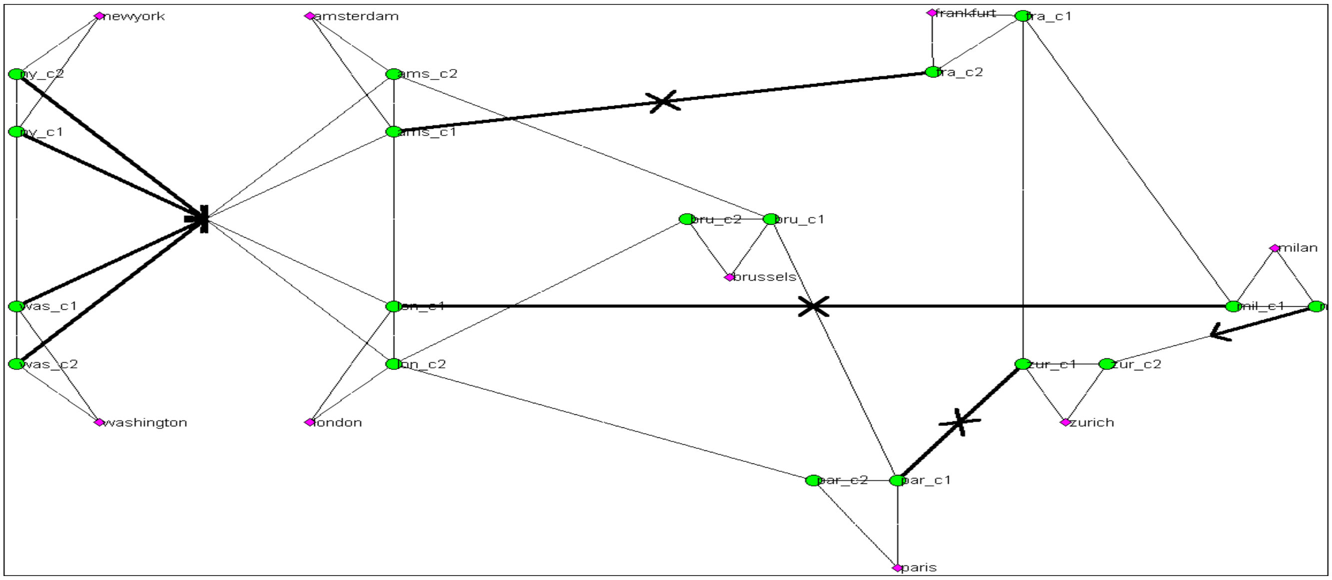

Figure 4. Network topology for Fixed-Point Models

- Links had a higher utilisation for the case of TCP without rate limits than with a 1Mbps rate limit, and for this buffer depth there were more links with a high utilisation than in the open-source case. This of course is largely a result of the open-source offered loads assumed.

- The link from mil_c2 to zur_c2 had the second highest loss of high priority packets at 0.5%. This link is used by three routes: New York to Zurich, London to Zurich and Milan to Zurich. Of these, Milan to Zurich route has the highest end-to-end throughput of high priority traffic in the entire network. That one of the most congested links should form a part of a route with such a high throughput is explained by the effect of the round-trip times, which for the routes from New York and London are 114ms and 34ms respectively, whereas from Milan it is only 6ms seconds. The very short round-trip time explains the high mean throughput per connection, of ~25 and 7 times that for New York and London respectively. This type of interaction between TCP dynamics and the network topology would be hard to discover without a coherent fixed-point model of the network taking into account the adaptive nature of the load.
- The Milan to Zurich route serves to illustrate an interesting phenomenon when rate limits are imposed at the sources. Imposing a rate limit of 1 Mbps on all sources causes the mean throughput of high priority traffic to fall by a factor of 4, but the mean throughput of low priority traffic to rise by a factor of ~10. This is an extreme example, but the effect is apparent on other routes: if the high priority traffic is limited not by congestion in the network but by a rate limit imposed at the source then there will be more network resource, and hence higher throughputs, available to the lower priority traffic.
- Lowering the 1Mbps access rate limit to 56kbps (corresponding to today's dial-up modems) and keeping the mean number of sessions the same along a route, there was very low packet loss and link utilisation throughout the network. The access rate limits here prevented the backbone network links reaching congestion levels by limiting the offered traffic.

4. Summary

Fixed-point models have been developed for IP networks which allow the adaptive nature of TCP sources to be accounted for. Based on an input matrix of a mean number of sessions per route and per traffic class, the parameters of packet loss, link utilisation and mean TCP throughputs have been determined. The key assumption in the model has been that of independence of packet inter-arrivals at a resource, and the fixed-point solution is determined as a natural network operating point for a resource based on an M/M/1/K (i.e. finite buffer) queue. Generalisation to other resource models has been discussed. This approach accounts for the adaptive nature of the end-stations, and the model has been applied to topologies typical of IP backbones today. Interesting features of the effect of topology on performance show up when packet losses along a route are computed using the method of thinning, and show that TCP throughput over many routes can be determined by the level of utilisation of the "weakest" link in the chain. Validation of this approach based on network measurements[e.g. as in 18,25] is now essential to ensure that design and performance guidelines can be provided to network planners, however the overall approach captures the essential feature of adaptive sources operating with TCP protocol in congestion avoidance.

In summary the aim of this work has been to capture in a coherent model, the various individual sub-processes, which together result in the end-to-end performance of large-scale (many nodes and routes) networks with adaptive traffic sources. This is realised through the development of a coherent modelling framework in terms of existing, though primitive and simplistic, modelling proposals for the sub-processes (M/M/1/K packet level resource models, M/G/∞ connection level models) and use of the independent resource assumption.

Improved sub-process models, while desirable for accuracy, would introduce overwhelming complexity at this stage. The degree to which they are necessary is problematic since comparison with simulation is frustrated by the large-scale nature of the network problems that need to be addressed. The examples illustrate the potential power of such a simplified end-to-end network model, and accordingly help motivate large-scale simulation studies and deeper studies of the sub-process modelling assumptions.

References

1. UUNet Service Level Agreements at http://www.uu.net/customer/sla/
2. AT&T Network Performance at http://ipnetwork.bgtmo.ip.att.net/
3. R. J. Gibbens, S. K. Sargood, F. P. Kelly, H. Azmoodeh, R. N. Macfadyen, N. W. Macfadyen, "An Approach to Service Level Agreements for IP Networks with Differentiated Services, " Royal Society Discussion Meeting on Network Modelling in the 21st Century, December 1999 at http://www.statslab.cam.ac.uk/~richard/research/topics/royalsoc1999/
4. K. W. Ross, " Multi-Service Loss Models for Broadband Telecommunication Networks," Springer-Verlag, 1995.
5. TCP friendly WWW site at http://www.psc.edu/networking/tcp_friendly.htm l
6. M. May, J. C. Bolot, A. Jean-Marie, C. Diot, "Simple Performance Models of Differentiated Services Schemes for the Internet," Proceedings of INFOCOM'99, New York, March 1999.

7. J. Heinamen, F. Baker, W. Weiss, J. Wrocklawski, "Assured Forwarding PHB," RFC 2597, June 1999
8. K. Nichols, F. Baker, S. Blake, D. Black, " Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474, Dec. 1998
9. D. Bertsekas and R. Gallagher, "Data Networks, " Ch. 5, Prentice-Hall, 1992
10. L. Kleinrock, "Communication Nets," Dover, 1964
11. W. Leland, M. Taqqu, W. Willinger, D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," IEEE/ACM Transactions on Networking, Vol. 2 No. 1, pp. 1-15, Feb. 1994.
12. V. Paxson & S. Floyd, "Wide Area Traffic: The Failure of Poisson Modelling," IEEE/ACM Transactions on Networking, Vol. 3 No. 3, pp. 226-244, June 1995.
13. B. K. Ryu and A. Elwalid, "The Importance of Long-range Dependence of VBR Video Traffic in ATM Traffic Engineering: Myths and realities," ACM Computer Communication Review, vol. 26, pp. 3-14, Oct. 1996.
14. M. Grossglauser and J. Bolot, "On the Relevance of Long Range Dependence in Network Traffic," IEEE/ACM Transactions on Networking, 1998.
15. V. Paxson and S. Floyd, "Why We Don't Know How to Simulate the Internet," Proceedings of the 1997 Winter Simulation Conference, December 1997.
16. W. Willinger and V. Paxson," Where Mathematics Meets the Internet," Notices of the American Mathematical Society, Vol.45, No.8, pp. 961-970, August 1998
17. M. Mathis, J. Semke, J. Mahdavi, T. Ott, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", Computer Communication Review, Vol. 27, No.3, July 1997.
18. Padhye, J., Firoiu, V., Towsley, D., and Kurose, J., Modeling TCP Throughput: A Simple Model and its Empirical Validation, UMASS CMPSCI Tech. Report TR98-008, Feb. 1998.
19. D. Heyman, T.V. Lakshman, A.L. Neidhart, " A New Method of Analysing Feedback-Based Protocols with Applications to Engineering Web Traffic over the Internet," AT&T Technical Report, 1999
20. M. May, T. Bonald, J. C. Bolot, "Analytical Evaluation of RED Performance," to be presented at Infocom'2000
21. Cisco's version of Class Based Weighted Fair Queueing can be found from http://www.cisco.com
22. K. Claffy, G. Miller, K. Thompson, "The Nature of the Beast: Recent Traffic Measurements from an Internet Backbone," INET '98 Conference, April 1998
23. K. Thompson, G. J. Miller, R. Wilder, "Wide Area Internet Traffic Patterns and Characteristics, " IEEE Network, Vol. 11, No. 6, Nov./Dec. 1997
24. H. Frystyk Nielsen, J. Gettys, A. Baird-Smith, E. Prud'hommeaux, H. Lie, C. Lilley, " Network Performance Effects of HTTP/1.1, CSS1, and PNG," Proc. ACM Sigcomm., 1997
25. N. Cardwell, S. Savage, T. Anderson, "Modelling TCP Latency," to be presented at IEEE Infocom 2000, March 2000.