Performance of TCP with Multiple Priority Classes

R. Vranken[1], R.D. van der Mei[2,3], R.E. Kooij[2], J.L. van den Berg[2,4]

[1]Technische Universiteit Eindhoven
Mathematics and Computer Science
Eindhoven, the Netherlands

[2]KPN Research
Expertise Group QoS Control
Leidschendam, the Netherlands
{R.D.vanderMei, R.E.Kooij, J.L.vandenBerg}@kpn.com

[3]Vrije Universiteit Amsterdam
Mathematics and Computer Science
Amsterdam, the Netherlands

[4]Universiteit Twente
Mathematical Sciences
Enschede, the Netherlands

Abstract

We consider the dimensioning problem for Internet access links carrying TCP traffic with two priority classes. To this end, we study the behaviour of TCP at the flow level described by a multiple-server Processor Sharing (PS) queueing model with two customer classes, where the customers represent flows generated by downloading Internet objects; the sojourn times represent the object transfer times. We present closed-form expressions for the mean sojourn times for high-priority customers and approximate expressions for the mean sojourn times of low-priority customers. The accuracy of the model is demonstrated by comparing results based on the PS model with "real" TCP simulation results obtained by the well-known Network Simulator. The experimental results demonstrate that the model-based results are highly accurate when the mean object size is at least 10 IP-packets, and the loss rate is negligible.

1. Introduction

Over the past decade the use of Internet services has experienced dramatic growth, and recently, Internet applications have evolved from standard document-retrieval functionality to advanced multimedia services. One of the main barriers to the continuing success of the Internet in the near future is Internet congestion, which in many cases leads to unacceptably long response times, particularly for real-time applications like the World Wide Web (WWW), PC banking and other on-line services. Consequently, the ability to somehow deliver QoS guarantees to the end users strongly enhances the competitive edge crucial for the business of Internet service providers (ISPs). Therefore, the Internet is currently migrating from best-effort type of service delivery technology towards a network able to support QoS-guarantees to the end users. The vast majority of the Internet traffic is transported over the Transfer Control Protocol (TCP) guaranteeing reliable transmission of data by means of a flow-control mechanism. To support QoS differentiation, the IETF has proposed the so-called `IP Diffserv` concept, defining priority classes at the IP layer.

In this paper we focus on the dimensioning of TCP/IP based networks supporting QoS differentiation. More specifically, we consider a single bottleneck link and address the question "How large should the bottleneck link capacity be to guarantee certain QoS objectives?" To this end, we describe the flow-level behaviour of TCP by a multiple-server Processor Sharing model with two customer classes, where the customers represent flows generated by downloading Internet objects and the sojourn times represent the object transfer times. The QoS metrics are the mean document transfer times for both high and low priority customer classes. Closed form expressions for these quantities have been obtained in our previous papers [1-2]. The validity of the use of our model is validated by comparing the model with TCP simulations. The results indicate that the model works well when the mean file size is at least the size of 10 IP packets, and the packet loss ratio is negligible. We also observe that the buffer size of only 50 packets, often used as a default setting for commercial Internet routers, leads to a significant degradation of Internet performance.

In the literature a significant number of papers focus on the performance of TCP/IP-based networks. Analysis of empirical traffic measurements has revealed that packet-level WAN traffic exhibits two intriguing properties: long-range dependence and self-similarity (cf. Park and Willinger [3] for an overview). At the flow level, Processor Sharing (PS) models provide a simple and effective way of modelling the elastic properties of traffic carried by transport protocols that adapt the transmission rate to the available capacity in the network. Nabe et al. [4] propose to use an M/G/1-processor sharing (PS) queuing model to discuss a design methodology of Internet access networks. Nunez et al. [5], Lindberger [6] and Bonald and Roberts [7] propose an M/G/C-PS queueing model to describe the flow-level behaviour for a link that can accommodate C times the peak rate of an individual source. Bonald et al. [8] show that the statistical bandwidth sharing performance under perfect fairness, such as in the case of PS-models, is insensitive to the flow-size distribution and the flow arrival process, given only that session arrivals are Poisson. Charzinski [9] defines the so-called fun factor as the expected value of the ratio between the data transfer time at full rate and the actually observed data transfer time. An attractive feature of PS-based models is that analytical results on their performance are available. We refer to an excellent paper by Cohen [10] for results on a very general class of PS-type of models. The validity of PS-based models is verified by Riedl et al. [11-12] by comparing the performance predictions based on the model by means of the Network Simulator, and by Beckers et al. [13] on the basis of experimental data in a real-life environment with customers with ADSL-based high-speed Internet access.

The remainder of this paper is organized as follows. In section 2, we discuss the problem is more detail, define the model and give exact and approximate expressions for the mean transfer times for both high and low priority customer classes. In section 3 the validity of both the model itself and the performance predictions based on the model are validated by means of the Network Simulator. Finally, section 4 contains several concluding remarks and addresses several challenging topics for further research.

2. Model description

2.1 Problem description

We consider a population of end users that can access information from application servers connected to the core network. The connectivity between the access network and the core network is provided via a bottleneck access link, see Figure 1. End users generate transaction requests to the application servers to download Internet objects (such as Web pages, emails or other documents) over TCP connections between the end users and the application servers. In this way, each transaction request generates one or more 'TCP flows' over the bottleneck access link. We distinguish two types of customers: high and low priority customers. In case of temporary overload of the access link, high priority customers have strict priority over low priority customers, and consequently, the low priority customers can only utilize the bandwidth 'left behind' by high priority customers.

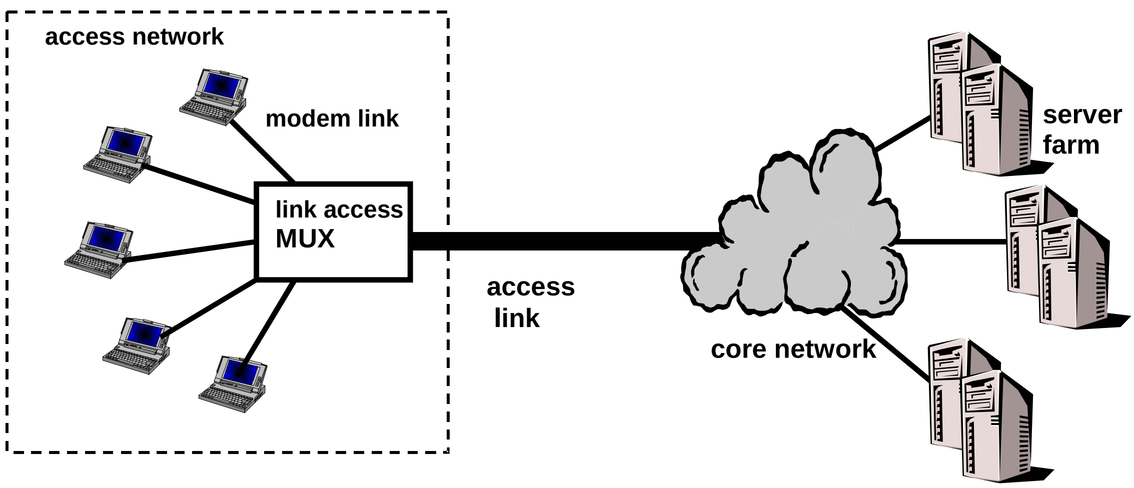

Figure 1. Access link connecting the access network to the core network.

For a cost-efficient design of their infrastructure, Internet Service Providers (ISPs) need to properly dimension the access links giving the subscribers access to the Internet: under-dimensioning generally leads to a less than acceptable QoS perceived by the subscribers, whereas over-dimensioning may be too expensive. Therefore, in this paper we focus on the dimensioning of the access link: "How much bandwidth is needed such that the object download times for both priority classes are acceptable?"

The packet-level dynamics are a complex interplay between many factors (such as the TCP window mechanism, the round-trip times, the buffer sizes, the maximum segment sizes), generally leading to highly complicated performance models. However, in the context of dimensioning access links, flow-level characterization of the TCP traffic, omitting many details of the packet-level TCP dynamics, is sufficient. Therefore, in this paper we consider a flow-level performance model for TCP with priority classes. The details of the model are discussed in the next subsection.

2.2 Flow-level model for TCP with Priorities

In this section we model the flow-level dynamics of TCP supporting two QoS classes on a single bottleneck link of capacity C by an M/G/C Processor Sharing model with C identical servers and two priority classes. Customers represent flows of packets in the TCP network and the sojourn times represent the object transfer times. Moreover, it is assumed that the maximum transfer rate of each customer is 1. High priority customers have strict priority (pre-emptive resume) over low priority customers. High and low priority customers arrive according to independent Poisson processes with rates λ_H and λ_L, respectively. The service-time requirements of the high and low priority customers are generally distributed with finite first two moments β_H, $\beta_H^{(2)}$, β_L and $\beta_L^{(2)}$, respectively. The average load of the high and low priority classes is denoted by $\rho_H := \lambda_H\beta_H$ and $\rho_L := \lambda_L\beta_L$, and the total load of the system is denoted by $\rho_{H+L} := \rho_H + \rho_L$. The service process of high priority customers alternates between two modes: a *normal mode* and a *processor sharing mode*. The process is in normal mode when the number of high priority customers does not exceed C. In that case, each high priority customer occupies a single server and is served at unit rate. When the number of high-priority customers is larger than C, the system switches to a processor sharing mode. In that case, the total service capacity C is equally shared among the high priority customers: when there are $n_H \geq C$ high-priority customers in the system, each of these customers is served in a processor sharing fashion with rate C/n_H. Notice that customers are not buffered, and there is no customer blocking. The servers not used by the high priority customers are available for service of the low priority customers. The service process of low priority customers also switches between a normal mode and a processor sharing mode. The low priority service process is in normal mode if the total number of customers in the system does not exceed C; in that case, each customer is served by a single server at unit rate. The low priority service process switches to processor sharing mode when the total number of customers exceeds C: when there are C_L servers available for serving low priority customers and there are $n_L \geq C_L$ low priority customers in the system, then each of these customers is served at rate C_L/n_L. When all servers are occupied by the high priority customers (i.e., $n_H \geq C$), the service of the low priority customers is stopped; their service is continued as soon as the number of high priority customers becomes less than C. Recall that the priority rule is pre-emptive resume. The stability condition of the system is $\rho_{H+L} < C$. Throughout, it is assumed the system is stable and in steady state. Denote by S_H and S_L the steady state sojourn time of an arbitrary high priority and low priority customer respectively. In this paper, our main focus is on $E[S_H]$ and $E[S_L]$, the mean sojourn times of both high and low priority customers.

2.3 Analysis and Approximations

In previous papers [1-2] we have derived explicit results for the mean sojourn times of low and high priority customers. In particular, noting that the service of the high priority customers is not influenced by the presence of the low priority customers, the number of high priority customers in the system behaves as the number of customers in an M/G/C Processor Sharing model, which is a special case of the more general Processor Sharing type of models analyzed in Cohen [10]. From the results in that paper we obtain the following closed-form expression for $E[S_H]$: For $0 \leq \rho_H < C$,

$$E[S_H] = \frac{C^C}{\lambda_H C!\,A(C,\rho_H)}\left(\frac{C(\rho_H/C)^{C+1}}{1-\rho_H/C} + \frac{(\rho_H/C)^{C+1}}{(1-\rho_H/C)^2}\right) + \frac{1}{\lambda_H A(C,\rho_H)}\sum_{n=1}^{C}\frac{\rho_H^n}{(n-1)!},$$

with

$$A(C,\rho) := \frac{C^C}{C!}\frac{(\rho/C)^{C+1}}{1-\rho/C} + \sum_{n=0}^{C}\frac{\rho^n}{n!}.$$

Notice that the mean sojourn times for high priority traffic are *insensitive* with respect to the service-time distribution, in the sense that they depend only on the first moment β_H.

The analysis of sojourn times for low priority customers is more complicated. The complication is due to the fact that the service rate at which low priority customers are served fluctuates, depending on the variation in the number of high priority customers in the system. An exact mathematical analysis of the mean sojourn times appears to be possible only either if C=1 and service times for low-priority customers are exponential, or if C ≥1 and service times are exponentially distributed with means $\beta_H = \beta_L$. In the general case, however, an exact analysis does not seem to be possible. In [1-2] we have derived and validated the following approximation for the mean sojourn time for the low priority customers: For $0 \le \rho_{H+L} < C$,

$$E[S_L] \approx \frac{2\beta_L}{\lambda_L \beta_L^{(2)}}\left(\frac{\beta_{H+L}^{(2)}}{2\beta_{H+L}} f(\rho_{H+L}) - \frac{\beta_H^{(2)}}{2\beta_H} f(\rho_H)\right),$$

where

$$\beta_{H+L} := \frac{\lambda_H}{\lambda_H+\lambda_L}\beta_H + \frac{\lambda_L}{\lambda_H+\lambda_L}\beta_L, \beta_{H+L}^{(2)} := \frac{\lambda_H}{\lambda_H+\lambda_L}\beta_H^{(2)} + \frac{\lambda_L}{\lambda_H+\lambda_L}\beta_L^{(2)},$$

and, for $0 \le \rho < C$,

$$f(\rho) := \frac{C^C}{C!\,A(C,\rho)}\left(\frac{C(\rho/C)^{C+1}}{1-\rho/C} + \frac{(\rho/C)^{C+1}}{(1-\rho/C)^2}\right) + \frac{1}{A(C,\rho)}\sum_{n=1}^{C}\frac{\rho^n}{(n-1)!},$$

with $A(C, \rho)$ as defined above.

This approximation has been extensively tested by simulation and has been shown to be very accurate for a broad range of the parameter values, see [1-2]. Note, that if high and low priority customers have the same service requirement distribution, the approximation shows insensitivity with respect to this distribution in the sense that it only depends on the mean service requirement.

3. Validation

In the previous sections we have presented expressions for mean sojourn times for both high and low priority customers in a M/G/C-PS system. Because we want to apply the M/G/C-PS model for the prediction of TCP file download times, we will compare the model with TCP simulation results.

3.1 Experimental set-up

We apply our proposed model to the network scenario depicted in Figure 1, where the access link has a capacity of 10 Mbps, while each user has a modem rate of 1 Mbps.

In our simulations we have used version 2 of the Network Simulator (NS, cf. [14]), which is designed to simulate file transfers in a TCP/IP network environment. We have assumed a total number of 1000 customers, which are represented by 1000 TCP sink nodes. The same number of TCP sources are representing servers from where files are being downloaded by the customers. For the validation of our model we only consider downstream traffic. The traffic sources used TCP Tahoe, with maximum window size 20 packets and packet size 1500 bytes. The above scenario suffices to test the model in case there is only one priority class. If there are two priority classes then two types of traffic sources are defined and each type of traffic source has its own TCP sink (destination). We have implemented this multi priority scenario by using the so-called class based queueing command predefined in NS.

3.2 Single priority case

Although we are interested in an Internet based download session with QoS differentiation we will study the single priority case first. Because the M/G/C-PS model does not take all the dynamics of TCP into account (e.g. slow-start and the effect of packet loss) we study for which parameter settings the M/G/C-PS model is an accurate approximation in the single priority class case.

3.2.1 The impact of buffer size

Figure 2 shows the average file download time as a function of the load, both for our model and for simulation results with a buffer size of 50 and 5000 packets. A buffer size of 50 packets is often used as a default setting for commercial Internet routers. With a buffer size of 5000 packets we model a very large buffer. In Figure 2 the file sizes are assumed to be exponentially distributed with a mean of 15kB. Figure 3 shows similar results for the case that the file size follows a Pareto distribution with the same mean and a shape parameter of 2.2, leading to a squared coefficient of variation of 2.27. This choice of the shape parameter is based upon real data analysis, see Nabe et al. [4].

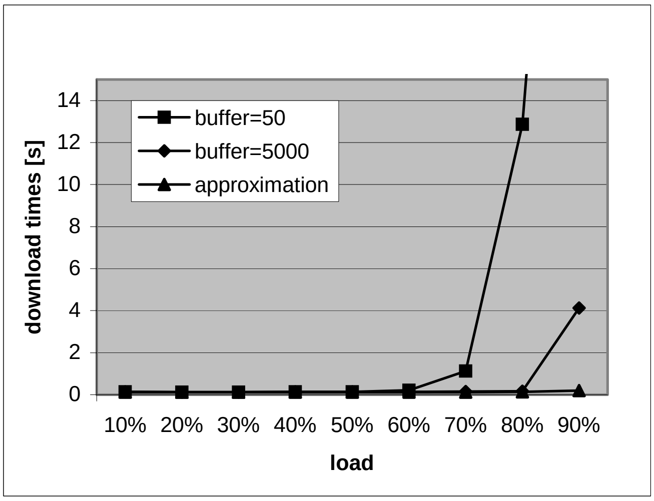

Figure 2. Average download times: exponential file size, mean = 15kB. Access link rate C = 10 Mbps.

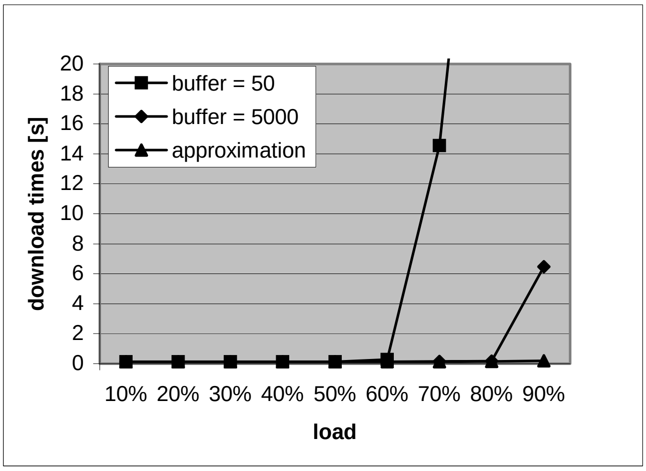

Figure 3. Average download times: Pareto file size, mean = 15kB. Access link rate C = 10 Mbps.

The results presented in Figures 2 and 3 demonstrate that the buffer size plays an important role for the download times. In fact, a buffer size of only 50 packets, often used as a default setting for commercial Internet routers, leads to a significant degradation of Internet performance. Notably, in order to allow fast downloads under heavy load, a large buffer size is needed. Under that condition the M/G/C-PS model forms an accurate model for predicting TCP based download times. This can be explained as follows. The M/G/C-PS model does not take the effect of packet loss on TCP into account. For a small buffer (50 packets) packet loss typically starts to occur at loads of 60%.

3.2.2 The impact of mean file size

Figure 4 shows the relative error of the approximations based on the M/G/C-PS model as a function of the load. The file sizes are Pareto-distributed with shape parameter 2.2, while the mean file size takes the values 15 kB, 150 kB and 1.5 MB. The buffer size is 5000 packets.

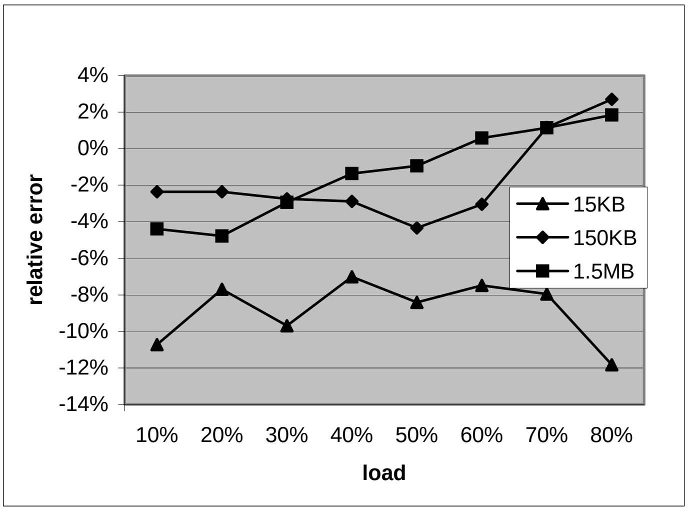

Figure 4. Relative error of M/G/C-PS model with respect to NS simulations. Pareto file sizes. Access link rate C = 10 Mbps.

Based upon Figure 4, and similar simulations, we conclude that the M/G/C-PS model leads to a very accurate estimate of average download times, if the load does not exceed 80%, the average file size is larger than 15 kB and the buffer is sufficiently large. The condition on the average file size can be formulated more generically: the average file size is at least 10 packets.

For an average file size of 15kB the M/G/C-PS model underestimates the download time by about 10%. This underestimation is due to neglecting the TCP slow-start phase. The results also show that the model becomes better with increasing file size because the influence of the slow-start algorithm diminishes with increasing file size.

3.2.3 The insensitivity of the average download time with respect to the file size distribution

The M/G/C-PS model predicts that the average download time is insensitive with respect to the file size distribution, see Section 2. Table 1 shows the simulation results for deterministic, exponential and Pareto file size distributions. The mean file size is 150 kB and for the Pareto case the shape parameter is again set to 2.2.

load	deterministic	exponential	Pareto
0.1	1.20	1.20	1.23
0.2	1.20	1.21	1.23
0.3	1.20	1.21	1.23
0.4	1.20	1.21	1.25
0.5	1.21	1.23	1.26
0.6	1.22	1.26	1.27
0.7	1.26	1.33	1.27
0.8	1.34	1.49	1.41
0.9	1.70	6.77	4.76

Table 1. Testing the insensitivity of the single priority class with respect to the service time distribution.

We conclude from Table 1 that the download times are indeed fairly insensitive to the file-size distribution when the load is no more than 80%.

3.2.4 Conclusion for single priority case

The aim of this section was to validate the use of the M/G/C-PS model in the setting of TCP based Internet downloads, in case there is only a single priority class of customers. We have found that the M/G/C-PS model forms an accurate model for predicting TCP based download times when the following conditions are met:

- All users have the same access rate
- No packet loss occurs: this can be realised by choosing an appropriately large buffer
- The mean file size is at least 10 IP packets
- The load on the link does not exceed 80%

We also found that when the above conditions are met, the insensitivity property for download times with respect to file size distribution holds. Notice that for the single priority case, similar results on the validation of using M/G/C-PS to model TCP downloads by using the Network Simulator have also been reported by Riedl et al. [11-12].

3.3 Multiple priority case

In the previous sections we have validated the results of our model for the mean sojourn time for the single priority case. In the multiple priority case the download times of the high priority customers are not influenced by the presence of low priority customers. Hence the validation of the accuracy of our model for the high priority customers follows from the previous sections. The aim of this subsection is to validate our model for the low priority customers.

3.3.1 Simulation versus model

In this section two numerical experiments are presented. Again the core link has a capacity of 10 Mbps and the access link capacity is 1 Mbps. In this section we set the buffer size equal to 5000 packets. The simulations are performed at different loads on the access link which are obtained by varying the arrival rates.

The first experiment assumes that the low priority files are exponentially distributed. The high priority files follow a Pareto distribution with shape parameter 2.2. We assume that the high and low priority files have equal mean arrival rate and mean file size. Therefore both priority classes equally contribute to the load. The mean file sizes considered are 15 kB and 150 kB, respectively. We determine the mean download times for the low priority sessions as a function of the load per server for the two mean file sizes, both through the Network Simulator and through our model, see Figure 5.

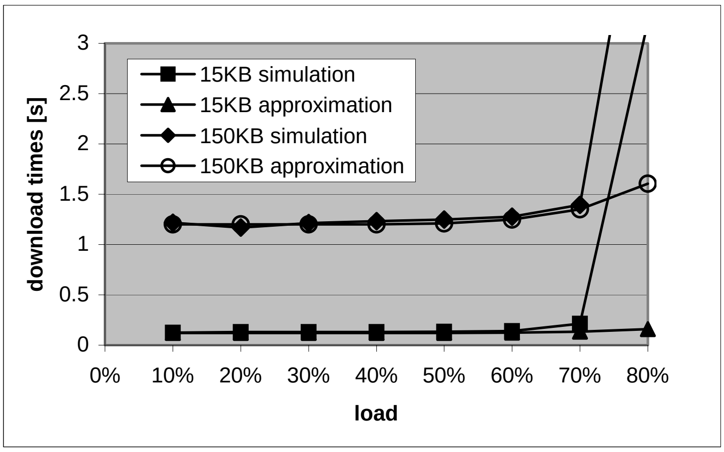

Figure 5. Mean download times for low priority customers: exponential low priority file size. Access link rate C = 10 Mbps.

Next we assume that the high and low priority files both follow a Pareto distribution with shape parameter 2.2. Again high and low priority files have equal mean arrival rate and mean file size, thus equally contributing to the load. The mean file sizes considered are 15 kB and 150 kB. The comparison between the outcome of the Network Simulator and our model is depicted in Figure 6.

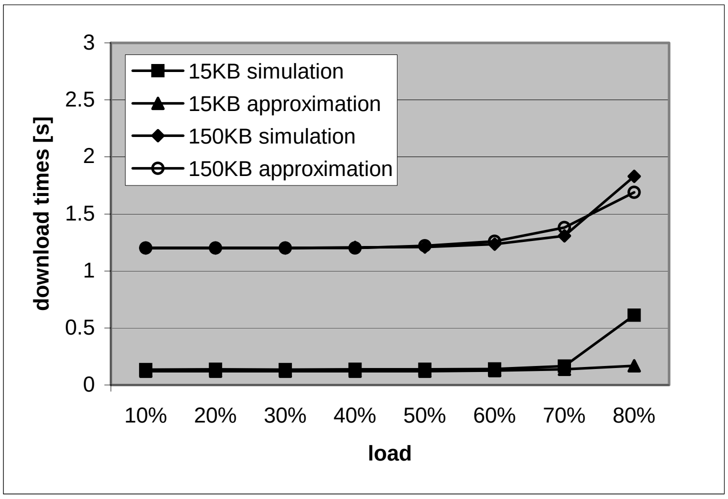

Figure 6. Mean download times for low priority customers: Pareto low priority file size. Access link rate C = 10 Mbps.

Based upon Figures 5 and 6, and similar simulations, we conclude that our model is able to predict average download times for low priority customers very well, as long as the load does not exceed 70%. Of course, it is assumed that the buffer is sufficiently large, such that no packet loss occurs.

3.3.2 The case of identical file size distributions for high and low priority customers

It is quite realistic to assume that the file size distributions for high and low priority customers are identical, as they all represent Internet downloads. It follows from the expressions presented in section 2.3 that according to our model the mean sojourn time is insensitive with respect to the file size distribution if the squared coefficient of variation of high and low priority files are equal. In this section we will conduct an experiment to validate this.

To this end, we assume that high and low priority files follow the same distribution, with mean file size 150kB. Again high and low priority files equally contribute to the load. In the experiment the file size distribution is deterministic, exponential or Pareto with shape parameter 2.2. This implies a squared coefficient of 0, 1 or 2.27 respectively. The results of the simulations are visualized in Figure 7.

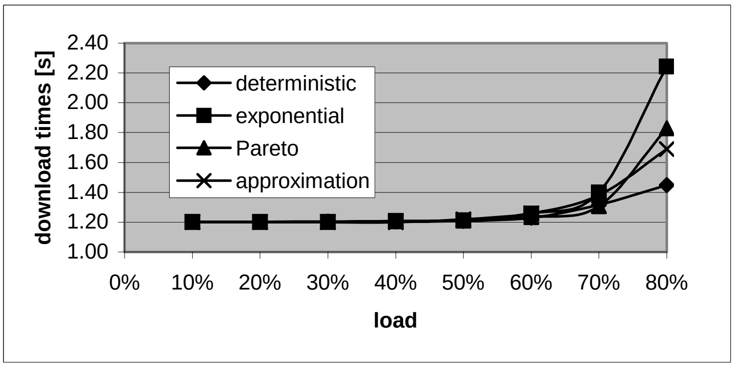

Figure7. Average download time: identical distributions for high and low priority customers. Access link rate C = 10 Mbps.

Based upon Figure 7, and similar simulations, we conclude that if the high and low priority files are identically distributed, then the average download times for low priority customers is not very sensitive with respect to the file size distribution.

3.3.3 Conclusion for multiple priority case

The aim of this section was to validate the use of our model in the setting of TCP based Internet downloads, in case of the occurrence of both high and low priority file downloads. We have found that our model forms an accurate model for predicting TCP based download times for the low priority customers when the following conditions are met:

- All users have the same access rate
- No packet loss occurs: this can be realised by choosing an appropriately large buffer
- The mean file size, both for high and low priority customers, is at least 10 IP packets
- The load on the link does not exceed 70%.

We also found that if, in addition, the file sizes of high and low priority are identically distributed then the mean down load time for low priority files is not very sensitive with respect to the file size distribution.

Note that because we are able to predict average download times with our model, it allows us to dimension the capacity of the access link.

4. Summary and topics for further research

In this paper we have proposed to use the M/G/C-PS model to describe the flow-level dynamics of TCP- based Internet object downloads in the occurrence of two priority classes. Validation experiments with the Network Simulator have demonstrated that the performance predictions based on the model are highly accurate when all users have the same access rate, the mean object size is at least the size of 10 IP packets, the link load does not exceeds 70-80% and the packet loss rate is negligible.

Finally we address a number of topics for further research.
First, the use of the PS-type models relies on the assumption that the customer population is homogeneous in the sense in case of overload that the bandwidth is equally shared among all customers. In practice, however, customer populations are heterogeneous, caused by differences in parameters such as the modem speeds, the settings of the maximum TCP window sizes, the round-trip times for TCP acknowledgments and the TCP flavor. It is an interesting topic for further research to refine the model to take into account the heterogeneity of the customer population and to develop and validate approximations for the Internet object transfer times for the refined model.
Second, in the present paper QoS differentiation was implemented by considering two priority classes, where high-priority customers have strict priority over low-priority customers. This priority scheme may lead to ‘starvation’ of low-priority customers, which may be undesirable in some applications. In fact, in many situations it seems to be more appropriate to implement QoS differentiation by means of Weighted Fair Queueing (WFQ), i.e., by putting weights on the relative amount of bandwidth that both priority classes receive. Investigation of the performance implications of such a WFQ-based scheduling scheme is a challenging topic for further research.
Finally, the performance predictions based on the model tend to be rather inaccurate when the objects have the size of only a few packets. We suspect that in those situations the transfer times are primarily determined by contention at the link buffers, rather than the fair sharing of bandwidth assumed in the model. Refinement of the model to incorporate the impact of link buffers is an interesting topic for further research.

5. References

1. J.L. van den Berg, R.D. van der Mei, B.M.M. Gijsen, M.J. Pikaart and R. Vranken (2001), Processing Times for Transaction Servers with Quality of Service Differentiation, Proc. Of Measuring, Modelling and Evaluation of Computer and Communications Systems, 12-14 september 2001, Aachen, 241-252.
2. R.D. van der Mei, J.L. van den Berg, R. Vranken and B.M.M. Gijsen (2001), Sojourn times in processor sharing systems with priorities, submitted for publication.
3. K. Park and W. Willinger eds. (2000). *Self-Similar Network Traffic and Performance Evaluation* (Wiley, New York).
4. M. Nabe, M. Murata and H. Miyahara (1998). Analysis and modeling of World Wide Web traffic for capacity dimensioning of Internet access lines. *Performance Evaluation* 34, 249-271.

5. R. Nunez, H. van den Berg and M. Mandjes (1999), Performance Evaluation of Strategies for Integration of Elastic and Stream Traffic, Proceedings ITC 16, Edinburgh.
6. K. Lindberger (1999). Balancing Quality of Service, pricing and utilization in multiservice networks with stream and elastic traffic. In: *Proceedings ITC16* (Edinburgh, UK), 1127-1136.
7. T. Bonald and J. Roberts (2000). Performance of bandwidth sharing mechanisms for service differentiation in the Internet. In: *Proceedings 13th ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management* (Monterrey, September), 22-1 - 22-10.
8. T. Bonald, A. Proutiere, G. Regnie and J.W. Roberts (2001). Insensitivity results in statistical bandwidth sharing. In: *Teletraffic Engineering in the Internet Era* (J. Moriera de Souza, N.L.S. Fonseca and E.A. de Souza e Silva eds.), 125-136.
9. J. Charzinski (2001). Measured HTTP performance and fun factors. In: *Teletraffic Engineering in the Internet Era* (J. Moriera de Souza, N.L.S. Fonseca and E.A. de Souza e Silva eds.), 1063-1074.
10. J.W. Cohen (1979). The multitype phase service network with generalized processor sharing. *Acta Informatica* 12, 245-284.
11. A. Riedl, M. Perske, T. Bauschert and A. Probst (2000). Dimensioning of IP access networks with elastic traffic, Networks 2000 (Toronto, September).
12. A. Riedl, T. Bauschert, M. Perske and A. Probst (2000). Investigation of the M/G/R processor sharing model for Dimensioning of IP access networks with eleastic traffic. Proceedings First Polish-German Teletraffic Symposium PGTS 2000 (Dresden, September), 97-106.
13. J.V.L. Beckers, I. Hendrawan, R.E. Kooij and R.D. van der Mei (2001). Generalized processor sharing models for Internet access lines. Proceedings of IFIP Conference on Performance Modeling and Evaluation of ATM & IP networks (Budapast, juni 2001), 101-112.
14. http://www.isi.edu/nsnam/ns/