WWW Traffic Modeling for HFC Networks

Dee Denteneer† & Verus Pronk†

Philips Research Laboratories
Prof. Holstlaan 4, 5656 AA Eindhoven, The Netherlands

Abstract

Traffic modeling for HFC networks differs from traffic modeling for queuing systems by a greater emphasis on individual connections and on short-range dependencies. These issues are illustrated by means of a model for WWW or web client traffic and a simulation of an HFC network that serves a number of web clients. These simulations offset the number of clients that an HFC network can serve against the quality of service they receive. In addition, these simulations compare the quality of service received by web clients to the quality of service received by comparable Poisson sources.

INTRODUCTION

Currently, teletraffic modeling is one of the hot topics of the telecommunications society and the body of work devoted to this topic is vast and rapidly expanding. It roots firmly in traffic measurements that establish that traffic on modern communication systems (e.g., an Ethernet or the Internet) differs in significant ways from assumptions about traffic that have been traditionally made in performance analysis.

With some simplification, one may say that the focus in this area is on models for *aggregate traffic* that exhibit *long-range dependence.* Interest in traffic due to a single source is secondary and motivated by a search for a physical explanation of this long-range dependence: often this explanation is formulated by means of heavy-tailed distributions (see [10], [14], [15], [18], or, more generally [1]). Here, long-range dependence means that time correlations between traffic loads exist for very long periods of time, i.e., they tail off hyperbolically in stead of exponentially as with the classical short-range dependent models. It is this characteristic of the traffic that is responsible for the excessive waiting times in queuing systems such as the Internet.

The relevance of these new models derives particularly from recent results in queuing analysis. In, e.g., [12], [5], and [17], it is shown that long-range dependence has an enormous impact on both waiting times and cell-loss probabilities in queuing systems. In the same vein, in [2], it is shown that waiting times in queues with heavy-tailed service times are considerably larger than waiting times in queues with light-tailed service times.

So traffic models for queuing systems rightly stress the long-range dependence of the traffic, possibly neglecting short-range dependencies.

Traffic modeling for HFC networks is different from traffic modeling for such queuing systems in two respects. First, in HFC networks, there is a clear relevance of short-range dependencies. Clearly, request mechanisms, such as multiple requests and piggybacking, make it plausible that packets generated 'close to each other' are relevant for throughput and delay. Bursts of traffic can effectively be dealt with, as it is not necessary to go through contention periods for each packet within a burst. Second, traffic modeling for HFC networks is more directly geared to the traffic generated by individual sources than to traffic form an aggregate of users, because it is the aggregation of single-user traffic itself that is the subject of investigation.

Analyzing the performance of HFC networks by means of simulation requires computer routines to generate traffic that mimics actual traffic in such systems. The sensitivity of HFC performance to traffic assumptions has, as yet, not been thoroughly investigated. However, studies that are available by now indicate that correct traffic models are of great concern and that both long-range dependencies and short-range dependencies are relevant.

As to the relevance of long-range dependence: in [8], the authors compare simulations with actually observed Ethernet traffic traces to simulations with artificial traffic, that was obtained by time-permuting these observed traces. Here, the traffic used in these two simulations is identical in one respect: the same data values are used in the simulations. However, the traffic streams differ in their time structure: the observed Ethernet traffic is long-range dependent, whereas the permuted traces are independent. Therefore, differences obtained in these simulations can be attributed to this time dependency. Their simulations show that the performance of an HFC network (measured in terms of average transmission delay) in case of the actual Ethernet traffic is much worse than the performance in case of the artificial traffic. They conclude that the correlation structure plays an important role, also for HFC networks.

In [6], the authors investigate the relevance of short-range dependencies: they compare the performance of HFC networks for traffic streams with exponential interarrival times to the performance for traffic streams with Pareto interarrival times. Again, the traffic assumption (exponential or Pareto) has an enormous impact on the outcomes of the simulations; the delays in simulations with Pareto interarrival times compare favorably with the delays in the simulations with exponential interarrival times.

Thus, there is sufficient evidence that traffic assumptions are relevant to the results (or may even determine these results) and that realistic models are needed. For this reason, we develop a model for web client traffic, which is then utilized, both to assess quality of service versus load and to further add to the current insights regarding sensitivity to traffic in HFC networks. It is the aim of this paper to further contribute to this sensitivity analysis by comparing the delay experienced by Poisson sources with the delay experienced by (very bursty) web client traffic.

The rest of the paper is organized as follows. In the next section, we present a model for web client traffic. The section thereafter describes the simulation configuration. We proceed by describing the results of simulating a typical HFC network that serves web clients, and end with some concluding remarks.

A MODEL FOR WEB CLIENT TRAFFIC

The aim of this section is to informally present a stochastic model that describes the traffic generated by a web client. Realization of this model resembles the actual traffic generated by web clients to such an extent that is not possible to distinguish artificial traces from actual traces.

Basically, the model has the following ingredients:

- The times at which packets are generated by the web client. This can equivalently be described by means of the interarrival times: the time between successive packets.
- The size of these packets.

The model then consists of the probability distribution functions that describe the variability of both interarrival times and packet sizes, and a description of the correlation between these. The familiar Poisson process is an example of a candidate model and substitutes an exponential distribution for the interarrival times, a constant distribution for the packet sizes, and assumes independence of all quantities.

However, the Poisson process falls short of our goals. Actual web client traffic is much more variable than a realization of a Poisson process and artificial traffic can easily be distinguished from actual traffic.

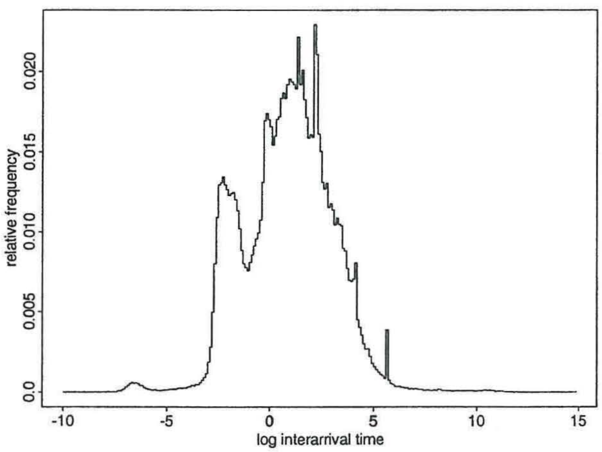

Figure 1: Histogram of the log interarrival times of successive page requests in the user process, based on a set of traces, selected from the UCB home IP usage study. The multi-modality of the histogram reflects the various time scales of the user process.

To understand the deficiencies of the Poisson process, observe that web client traffic is governed by two processes:

The user-process This is the process as perceived by users who are browsing the web. A web browser requests a succession of web-pages, e.g., by clicking with the mouse.

The TCP-process This is the actual packet exchange that goes on between the web client and the web server. After a user requests a page, one or several (as when a requested page contains several images) TCP connections are opened. The web client traffic in each connection consists of an *open connection*, an *information request*, a series of *acknowledgements*, and a *close connection*.

Each of these two processes has its own time scale, as users typically act much slower than computers. Hence, the existence of two time scales makes it untrue that just one, uni-modal probability distribution function will suffice to describe the time between successive events.

Now this argument can be extended. Again, the TCP-level process does not consist of a homogeneous generation of packets with identically distributed interarrival times. Rather, the traffic at this level consists of *flights* of packets. The time between packets in the same flight is largely determined by the speed with which a packet is handled by a computer; the time between successive flights is determined by the round-trip time of the network. Also, a user does not request pages at a constant rate. Rather, he will alternate between various states and these states can be characterized, e.g., as *actively navigating*, *thinking*, and *having a break*. The time between successive page requests will depend on the state the user is in. So, in order to describe the interarrival time at the user level, one probability distribution function is required for each such possible state.

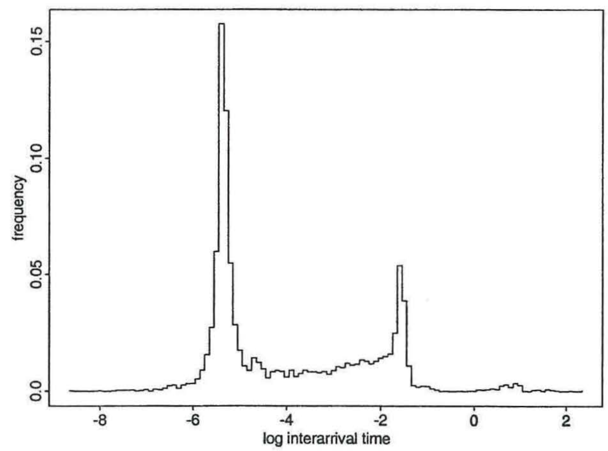

Figure 2: Histogram of the log interarrival times of successive packets in the TCP process, based on one trace from LBL-TCP-3. The multi-modality of the histogram reflects the various time scales of the TCP process.

These states can be argued about theoretically, but they can also be observed in measurements: see Figures 1 and 2. They are both based on publicly available data, that can be found at the Internet traffic archives [7]. In Figure 1, the histogram of the interarrival times for requests of the user process are displayed as observed in the UCB home IP usage study: a collection of traffic traces that contain information on home IP usage by UC Berkeley students, faculty and, staff over a period of 18 days in Nov. 1996. Clearly, the histogram has several modes. Each mode reflects one of the time scales of the user process and each mode can be labeled with one of the user states. Figure 2 shows the histogram of the interarrival times of packets of the TCP process from the client to the server, as in LBL-TCP-3: a trace of two hours containing all wide-area TCP traffic between Lawrence Berkeley Laboratory and the rest of the world (see [15]). Again, the multimodality makes it clear that there are distinct time scales that play a role.

Our model can then be summarized as follows. The web client traffic can be characterized by means of a succession of states: either user states or connection states (within flight or between flights). The interarrival time of the packets depends on this state and can be estimated from measurement data by means of algorithms such as EM and Baum-Welch (e.g., see [11]). The duration of these states can also be estimated from these data, and in our model we have arrived at the following rules: user states are equally probable and form an independent succession. The total duration of the connection state and the duration of the between and within flight states can be inferred from the length of the file to be transferred, the round trip time of the network, and the computer speed.

Finally we note that the time between flights of packets corresponds roughly to the round trip time of the connection. This makes the models scalable: the round trip time can be shortened (artificially) so that faster networks or caching-techniques can be investigated.

SIMULATION CONFIGURATION

An HFC network with a single upstream and a single downstream channel is considered with a transmission capacity of approx. 3 and 30 Mbit/s, respectively. The downstream capacity is assumed sufficient not to form a bottleneck in the system. The round-trip delay is set to 2.6 ms., which includes transmission, propagation, interleaving, and processing delays. The transmission medium is assumed error-free.

Before their transmission upstream, application data is segmented into 64-byte *data cells* with a payload of 48 bytes, corresponding to that of an ATM cell. The overhead associated to this segmentation includes the various headers, but also physical-layer overhead. So, at the MAC layer, two immediately successive, 64-byte data cells can be considered as if they are transmitted back-to-back.

Access to the upstream channel is organized with a *request-grant* mechanism, whereby requests are sent in contention and resulting grants guarantee collision-free transmission. For contention resolution, a blocked, ternary tree algorithm is used ([3], [16]). Only after a tree has completed, a new tree is initiated.

The size of a request cell is one third the size of a data cell, so that the 'size' of each node in a tree corresponds to the size of one data cell.

The upstream transmission time is slotted and scheduled on a frame-by-frame basis. The length of a frame is 3 ms., corresponding to the transmission of at most 18 data cells in 18 slots. In each frame, at least a fixed number n of nodes of the tree *can* be scheduled. In case less than n nodes are actually scheduled, the remaining slots can be used for data cells. Conversely, if not all $18-n$ slots are used for data cells, more than n nodes in the tree are scheduled, if available.

Applications, that have requested and are still awaiting grants, are granted collision-free transmission of data cells in a round-robin, cell-by-cell fashion.

When an application runs out of pending grants, i.e., when a *request bound* is crossed, and additional cells have arrived for transmission since its last request, a new request is transmitted.

As a single web client only produces a moderate amount of upstream traffic, i.e., in the order of 1 or 2 ATM cells per second on average, a bulk Poisson source, generating single ATM cells, is used to consume the bulk of the upstream bandwidth. The remaining bandwidth is (partly) consumed by either a number of web clients or an equal number of *equivalent* Poisson clients. An equivalent Poisson client generates the same amount of ATM cells as a web client on average, but with exponentially distributed interarrival times. In this way, pairs of corresponding simulations were carried out.

For the very bursty web clients, however, cell rates observed in simulations differ substantially from the theoretical mean of 1.76/s: in simulations, which typically cover only a few minutes, sample mean rates for web clients were found that differed by more than a factor of 25. So, in order to obtain proper values for the mean rate of the equivalent Poisson clients, the total traffic produced by all web clients during a simulation was used to calculate the mean rate of each of the equivalent Poisson clients for the corresponding simulation.

Table 1 lists the simulation settings. The codes WC and EP stand for web client and equivalent Poisson, respectively. The codes *low* and *high* stand for low and high aggregate load, respectively. Simulations WC and EP *low* give a joint data load of approx. 69%, of which 2.7% is jointly generated by the web clients and equivalent Poisson clients, respectively. In simulations WC and EP *high*, these figures are 84.8% and 11.5%. For the *low* simulations, at least one node of the tree can be (and is) scheduled, corresponding to a load of 5.6%. For the *high* simulation, at least two nodes can be scheduled, corresponding to a load of 11.1%.

simulation	*nodes/ frame*	*sources*		*cell rate*	*request bound*
WC *low*	≥ 1	1	bulk Poisson	4000/s	150
		100	web clients	1.64/s	15
EP *low*	≥ 1	1	bulk Poisson	4000/s	150
		100	eq. Poisson	1.64/s	15
WC *high*	≥ 2	1	bulk Poisson	4400/s	150
		400	web clients	1.72/s	15
EP *high*	≥ 2	1	bulk Poisson	4400/s	150
		400	eq. Poisson	1.72/s	15

Table 1: Simulation settings. Note that the web-client cell rates are time averages observed in simulations and these deviate from the long-term average of 1.76/s.

RESULTS

In this section, we present some results obtained in a series of simulations. Aim of these simulations is twofold.

- To compare the QoS that web clients receive with the QoS received by equivalent Poisson clients.
- To quantitatively offset the number of web clients that can be served by an HFC network against the QoS that they will receive.

Figure 3 illustrates time series of individual cell transmission delays (CTDs) for the bulk Poisson traffic, web client traffic and equivalent Poisson traffic relating to the *high* simulations.

The figure shows a number of notable differences. First, the CTD of the aggregate web-client traffic is significantly lower on average. Second, the bulk Poisson traffic CTD in WC *high* is lower and less variable than in EP *high*.

However, what strikes the most is that the differences are caused by a change in only a relatively small portion of the total traffic.

The large differences in mean CTD are to be contributed to the influence of the more bursty behavior of the web clients on primarily the contention resolution process. This bursty behaviour generally causes successive web-client cells to be generated by a relatively small number of web clients as compared to the uncorrelated generation of successive cells by the equivalent Poisson clients. As a result, fewer web clients with larger requests will contend in a single tree, causing less delay in getting the requests to the scheduler.

For a more detailed analysis, consider Figures 4 and 5, which illustrate cumulative distribution functions (CDFs) of the various CTD series. Figure 4 shows this for the bulk Poisson traffic for all simulations. Figure 5 only considers the WC simulations and shows the CDFs of both web-client and bulk Poisson traffic.

First, Figure 4 shows the unsurprising fact that cell transmission delay increases with increasing load: the CDFs of the *low* simulations lie to the left of the CDFs of the *high* simulations. Second, it can

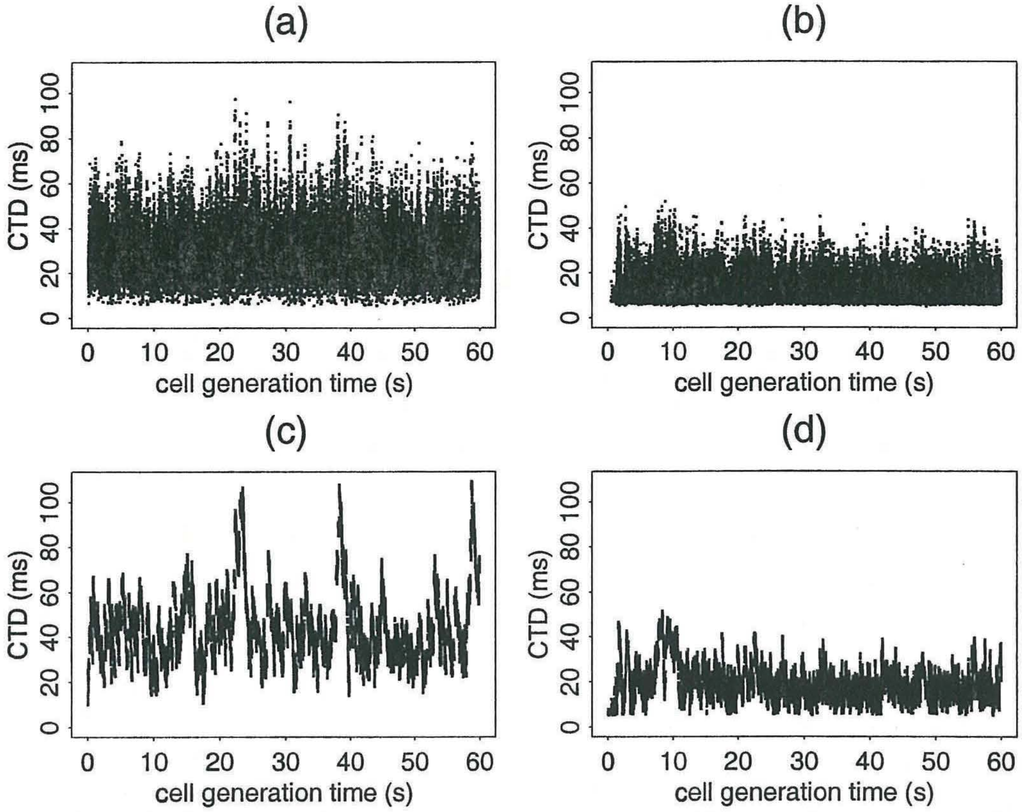

Figure 3: Time series of cell transmission delays (CTDs) for the *high* simulations: (a) equivalent Poisson client traffic, (b) web-client traffic, (c) bulk Poisson traffic that accompanies equivalent Poisson client traffic, and (d) bulk traffic that accompanies web-client traffic.

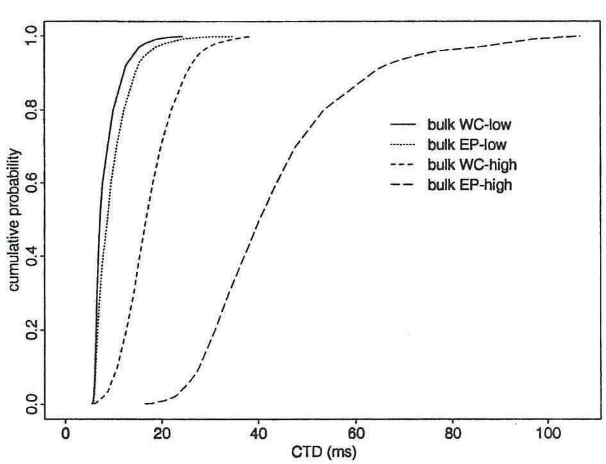

Figure 4: Cumulative distribution functions of the cell transmission delay (CTD) experienced by the bulk Poisson traffic in all simulations.

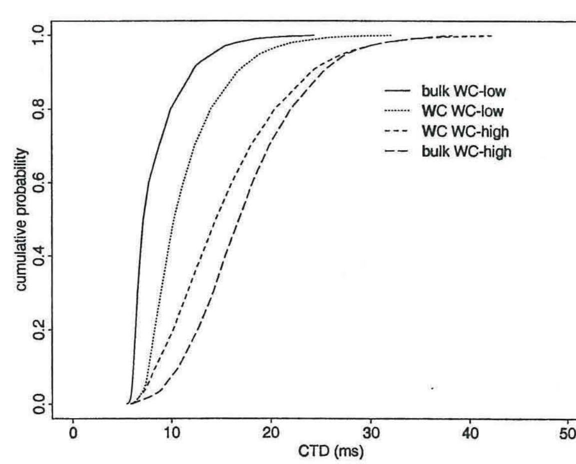

Figure 5: Cumulative distribution functions of the cell transmission delay (CTD) experienced by both the bulk Poisson and web-client traffic in the WC simulations.

be observed that the CDFs of the WC simulations lie to the left of those of the EP simulations. This indicates that transmission delay in case of web-client traffic is lower than in the corresponding case of equivalent Poisson traffic. Equivalent Poisson clients generate requests of size 1 in general, whereas web clients tend to generate larger requests, making the contention process more efficient.

Figure 5 compares the delay experienced by web clients and the bulk Poisson traffic. It shows that bulk Poisson traffic experiences less delay in the *low* simulation, but more delay in the *high* simulation than the accompanying web-client traffic. The difference in the *low* simulation is caused by the contention process: web clients will generally have to contend with the bulk Poisson source during the contention process, whereas the latter will often be the only contender. The difference in the *high* simulation is caused by the build-up of the queue for the bulk Poisson source. Similar behaviour was observed in the EP simulations.

To this comparison of the delays experienced by the various sources, we should add, however, that they also depend on the scheduler in use. In the current simulations, we have used a simple round-robin scheduler. Using fair schedulers, such as weighted fair queuing (see [4] or [13]) or weighted round-robin (see [9]), may significantly alter the results.

CONCLUSIONS

In this paper, we have illustrated the importance of accurate traffic modeling for HFC networks. Using Poisson processes only to describe traffic does not give a clear picture of HFC network performance. There is a clear need for application-specific traffic models for an accurate prediction of QoS versus load in service scenario studies. Most notable in this context is the need to also consider short-range dependencies in traffic, as well as single sources, as they significantly influence the contention-resolution process in HFC networks.

References

[1] Adler, R., R. Feldman, and M.S. Taqqu [1998], *A practical guide to heavy tails. Statistical techniques for analyzing heavy-tailed distributions*, Birkhäuser, Basel.

[2] Boxma, O.J., and J.W. Cohen [1998], The M/G/1 queue with heavy-tailed service time distribution, *IEEE Journal on Selected Areas in Communication*, Vol. 16, No. 5, pp. 749-763.

[3] Capetanakis, J.I. [1979], Tree algorithms for packet broadcast channels, *IEEE Transactions on Information Theory*, Vol. 25, no. 5, pp. 505-515.

[4] Demers, A., S. Keshav, and S. Shenker [1990], Analysis and simulation of a fair queuing algorithm, *Internetworking: Research and Experience*, vol. 1, no. 1, pp. 3-26.

[5] Erramilli, A., O. Narayan, and W. Willinger [1996], Experimental queuing analysis with long-range dependent packet traffic, *IEEE/ACM Transactions on Networking*, Vol. 4, No. 2, pp. 209-223.

[6] Harpantidou, Z., and M. Paterakis [1998], Random multiple access of broadcast channels with Pareto distributed packet interarrival times, *IEEE Personal Communications*, April, pp. 48-55.

[7] `http://www.acm.orig/sigcom/ITA/`

[8] Ivanovich, M.V., and M. Zukerman [1998], Evaluation of priority and scheduling schemes for an IEEE 802.14 MAC protocol loaded by real traffic, *Proc. Infocom '98/97.*

[9] Katevenis, M., S. Sidiropoulos, and C. Courcoubetis [1991], Weighted round-robin cell multiplexing in a general-purpose ATM switch chip, *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 1265-1279.

[10] Leland, W., M. Taqqu, W. Willinger, and D. Wilson [1994], On the self-similar nature of Ethernet traffic (extended version), *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1, pp. 1-15.

[11] McLachlan, G., and T. Krishnan [1997], *The EM Algorithm and Extensions*, Wiley.

[12] Norros, I. [1995], On the use of fractional Brownian motion in the theory of connectionless networks, *IEEE Journal on Selected Areas in Communications*, Vol. 13, No. 6, pp. 953-962.

[13] Parekh, A.K., and R.G. Gallager [1993], A generalized processor sharing approach to flow control in integrated services networks: the single-node case, *IEEE/ACM Transactions on Networking*, vol. 1, no. 3, pp. 344-357.

[14] Paxson, V. [1994], Empirically derived analytical models of wide-area TCP connections, *IEEE/ACM Transactions on Networking*, Vol. 2, No. 2, pp. 316-336.

[15] Paxson, V., and S. Floyd [1995], Wide-area traffic: the failure of Poisson modeling, *IEEE/ACM Transactions on Networking*, Vol. 3, No. 2, pp. 226-244.

[16] Tsybakov, B.S., and V.A. Mikhailov [1978], Free synchronous packet access in a broadcast channel with feedback, *Probl. Peredachi Inf.*, vol. 14, no. 4, pp. 32-59 [*Probl. Inform. Trans.*, vol. 14, pp. 259-280, 1978].

[17] Tsybakov, B., and N.D. Georganas [1997], On self-similar traffic in ATM queues: definitions, overflow probability bound, and cell delay distribution, *IEEE/ACM Transactions on Networking*, Vol. 5, No. 3, pp. 397-409.

[18] Willinger, W., M. Taqqu, R. Sherman, and D. Wilson [1997], Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level (extended version), *IEEE/ACM Transactions on Networking*, vol. 5, No. 1, pp. 71-86.

†The authors can be contacted at
{dentenee,pronk}@natlab.research.philips.com