Hybrid Simulation of a FIFO Queuing System with Trace-Driven Background Traffic

Ben Lauwens
Royal Military Academy
CISS/TELE
30, Renaissancelaan
Brussels, B-1000, BELGIUM
ben.lauwens@rma.ac.be

Jan Potemans
Katholieke Universiteit Leuven
ESAT/TELEMIC
10, Kasteelpark Arenberg
Heverlee, B-3001, BELGIUM
jan.potemans@esat.kuleuven.be

Bart Scheers
Royal Military Academy
CISS/TELE
30, Renaissancelaan
Brussels, B-1000, BELGIUM
bart.scheers@rma.ac.be

Antoine Van de Capelle
Katholieke Universiteit Leuven
ESAT/TELEMIC
10, Kasteelpark Arenberg
Heverlee, B-3001, BELGIUM
antoine.vandecapelle@esat.kuleuven.be

ABSTRACT

This paper introduces a novel hybrid packet-event / fluid-flow network simulation scheme. A packet-event simulation technique uses the arrivals and the departures of packets to model the queuing system. The applications that need fine-grained performance details, are simulated with an adapted event based approach. The impact of the background traffic on these foreground packets is simulated by virtual packets, which are created and put in the queue each time a foreground packet arrives. The calculation of the number and the size of these virtual packets is based on a fluid-flow approximation of the buffer occupation probability density function of the background stream. A *Many Sources Large Deviations* traffic descriptor is used to characterize the fluid-flow. A numerical evaluation of this hybrid simulation scheme is performed with a video streaming application as foreground traffic and measured network traces as background traffic.

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Probabilistic algorithms (including Monte Carlo); G.3 [**Probability and Statistics**]: Queuing Theory

General Terms

Simulation, Performance

Keywords

Hybrid Simulation, Large Deviations, Trace-Driven, FIFO Queuing System

1. INTRODUCTION

The discrete events for a packet-event network simulation are the arrivals of packets in a network node and the departure of packets on a carrier. A network can be modeled by queues having a limited buffer size. These queues are interconnected with links characterized by their transmission capacity. A packet-event network simulation gives a highly detailed view of what is happening in the network but it can become very slow if many events are generated. This is the case for high bandwidth links or for networks with a large number of devices. Parallel packet-event simulation can increase the scalability but this approach can not be used for high bandwidth links due to the sequential nature of a traffic flow on a link. The technique which uses events originating from the arrival or departure of packets, has to be abandoned if a higher simulation speed is required.

Fluid-flow based approaches for estimating the queuing behavior of network traffic can be used to accelerate the simulation. The level of detail is however less and precision is sacrificed for a faster run-time. Under fluid traffic, packets lose their identity; however, message identity is retained using the usual discrete-event methodology. Subtle protocol dynamics can not be studied using this technique. Network emulation systems which interact with applications running on a real network, are also not feasible.

The combination of packet-event simulations and fluid-flow approximations can be an answer to the challenge of getting packet level details in a reasonable amount of time. In the majority of scenarios, a simulation is carried out to get the performance indicators of a specific application when the data from that application is multiplexed with the background traffic on a network. A hybrid simulator separates the traffic into two classes. The packets of the foreground traffic, for which fine-grained performance details are needed, are simulated by an event-driven approach, while the background traffic, for which less detailed information is required, is approximated by a fluid-flow model.

These hybrid techniques in which packet-event simulation and fluid-flow approximations are combined, are a recent development. Some approaches use different simulators for the foreground traffic and the background stream [15], other separate the network into packet parts and fluid parts [5]. Two simulators [7, 12] integrate the Monte-Carlo simulation and the fluid model into the same simulator. Both use the behavior of the transport layer to model the fluid-flow. Traffic is simulated as an incompressible fluid, flowing among storage tanks (the buffers). For example open-loop (UDP) traffic is generated by an Exponential ON/OFF traffic generator and the closed-loop (TCP) source model is a simplified version of TCP Reno. A complex approach is needed to synchronize the foreground traffic and the background stream and to model the interaction of the the fluid-flow approximation with the packet-event simulation.

In this paper, an alternative hybrid simulation strategy is proposed. The foreground traffic is simulated in an adapted packet-event way. The impact of the background fluid-flow on the foreground traffic stream is modeled by putting virtual background packets in the packet-event queue, every time a foreground packet arrives. The number and size of the virtual background packets are estimated by sampling the buffer occupation probability density function of the background stream. The calculation of this background traffic descriptor is based on the Large Deviations theory. In articles [8, 9] the performance of different background calculation methods was evaluated: Heavy Traffic, Moderate Deviations and Large Deviations. These methods are all independent of the higher layer mechanics as TCP or UDP and can be computed using real traffic traces. The estimate based on the Large Deviations asymptotic has the most appealing characteristics for hybrid simulation and will be used in this paper as background traffic descriptor. In [13], the fluid-flow paradigm is extended with packet-event details. In our approach the event-driven simulation takes into account fluid-flow background streams.

In this paper all figures are based on the Bellcore LAN trace[1]. During this research many traffic traces have been investigated including traces of The National Laboratory for Applied Network Research[2]. No significant differences have been noted during analysis. The Bellcore LAN has been chosen as reference for long range dependency and self similar behavior, which is described in several papers [10, 14].

This paper is organized as follows. In section 2 the basic queuing model is detailed. The background traffic descriptor based on the Large Deviations theory is introduced in section 3 whereas section 4 deals with the hybrid simulation strategy. Packet-event simulations for a video streaming foreground application multiplexed with a real background trace, the Bellcore LAN trace, are compared to the hybrid simulation results in section 5.

2. QUEUING MODEL

2.1 Packet-Event Simulation

Packet-based data networks are easily modeled by queuing systems. Data is parceled up into packets and these are sent over wires. At network nodes where several wires[3] meet, incoming packets are inspected, queued up, and sent out over the appropriate wire. When the total number of traffic units, i.e. bytes or cells, reaches the buffer size, packets are discarded. The workload of the queue changes in a discontinuous way:

- if a new packet arrives in the network node and the number of traffic units in the queue plus the size of the new packet is less than the queue length, the load of the queue increases with the size of the packet, otherwise the packet is dropped;
- if a packet is processed, e.g. the last traffic unit is put on the wire, the following packet can leave the queue and the load of the queue decreases with the size of the packet which will be processed next.

Let t_i^A be the arrival time of the ith packet of a traffic trace and t_i^L the time that the ith packet leaves the queue. The workload of a FIFO[4] queuing system with an infinite buffer changes

$$\text{if} \quad t = t_i^A \quad : \quad Q_{t_i} = Q_{t_{i-1}} + A_i \tag{1}$$

$$\text{if} \quad t = t_i^L \quad : \quad Q_{t_i} = Q_{t_{i-1}} - A_i \tag{2}$$

where Q_t is the workload of the queue at time t and A_i is respectively the size of the packet arriving at time t_i^A and the size of the packet leaving at time t_i^L. t_{i-1} is the previous changing time of the load of the queue. t_i^L can be calculated by evaluating the time at which the previous packet needs to be sent out

$$\text{if the line is idle} \quad : \quad t_i^L = t_i^A \tag{3}$$

$$\text{if the line is in use} \quad : \quad t_i^L = t_{i-1}^L + \frac{A_{i-1}}{C} \tag{4}$$

where C is the transmission capacity in traffic units of the network link. The equations (1, 2, 3 and 4) are the governing formulas for a packet-event simulation. A trace file containing the arrival times of the packets and the corresponding packet sizes of a measured traffic stream can be used as input for the event-driven simulation. A realistic network simulator will take into account the packet drops due to the limited buffer size.

The buffer occupation PDF[5], noted f_{PE}, is calculated for each buffer occupation level as the time that the buffer has this load level divided by the total simulation time T:

$$f_{PE}(Q = B) = \frac{1}{T} \sum_{t_i=0}^{T} (t_{i+1} - t_i)\, \mathbf{1}_B(Q_{t_i}) \tag{5}$$

where T is the total simulation time and $\mathbf{1}_B(Q_{t_i})$ the indicator function of the packet-event buffer occupation.

[1]The BC-pAug89.TL is a text file that contains a row with a time stamp and the IP packet size for each packet. The trace can be downloaded form The Internet Traffic Archive: http://ita.ee.lbl.gov/.

[2]The Network and Measurement team of the National Laboratory for Applied Network Research distributes several data sets: http://pma.nlanr.net/.

[3]A wireless network can be regarded as one generalized processor sharing queuing system for all the devices within transmission range. A paper is currently in preparation and will be submitted end 2007 about an extension of the theory to wireless base stations and ad-hoc networks.

[4]First In First Out

[5]Probability Distribution Function

$$\mathbf{1}_x(y) = \begin{cases} 1 & x = y \\ 0 & x \neq y \end{cases}$$

2.2 Lindley's Fluid-Flow Recursion

The basic FIFO queuing system can be quantified using the Lindley's recursion [3]:

$$Q_{t_n} = \left(Q_{t_{n-1}} + A_{t_n} - C_{t_n}\right)^+ \quad (6)$$

where x^+ denotes the positive part of x, i.e. $\max(x, 0)$. Q_{t_n} can be interpreted as the amount of work in the queue between time t_n and t_{n+1}, A_{t_n} as the arrival process, i.e. the number of traffic units from packets that arrive in the interval $]t_{n-1}, t_n]$ and $C_{t_n} = C \cdot (t_n - t_{n-1})$ the number of traffic units served between times t_{n-1} and t_n with fix transmission capacity C. If the arrival times t_n correspond to the arrival times t_i^A, the concept of a packet leaving the queue is abandoned and a continuous stream of traffic units is processed. If the time unit Δt, the difference between times t_{n-1} and t_n, is constant, the arrival process is also a continuous stream of traffic units and the idea of a packet entering the queue is also lost. A traffic flow is entering the queuing system and a traffic stream is leaving the system in a fluid way; a fluid-flow simulation is performed.

Using fix time units Δt, called fluid steps or bins, Lindley's recursion simplifies to

$$Q_n = (Q_{n-1} + A_n - C_{\Delta t})^+ \quad (7)$$

where Q_n can be interpreted as the load of the queue during fluid step $n \in \mathbb{Z}$, A_n the number of traffic units from packets that arrive in fluid step n and $C_{\Delta t} = C\Delta t$ as the constant number of traffic units served in one fluid step. Equation (7) is the most rudimentary formula describing the behavior of a fluid flow in a FIFO queuing system with an infinite buffer.

The fluid-flow buffer occupation PDF, f_{FF} can be calculated as the number of fluid steps the buffer occupation level has the corresponding level divided by the total number of fluid steps N:

$$f_{FF}(Q = B) = \frac{1}{N} \sum_{n=0}^{N} \mathbf{1}_B(Q_n) \quad (8)$$

where N is the total number of fluid steps and $\mathbf{1}_B(Q_n)$ the indicator function of the fluid occupation level.

The CCDF[6] of the buffer occupation can be calculated:

$$\Pr(Q > B) = 1 - \sum_{b=0}^{B} f(Q = b) \quad (9)$$

Figures 1 and 2 show the CCDF of the buffer occupation of a packet-event simulation (PE) and of several fluid-flow simulations (FF) with different fluid steps for the Bellcore LAN trace with a transmission capacity $C = 10$ Mbps. The fluid step, 1.214 ms, corresponds to the maximum packet size, 1518 byte, divided by the transmission capacity.

The CCDF of the buffer occupation of the fluid-flow simulation mimics the CCDF of the packet-event simulation if the fluid step is the maximum packet size divided by the transmission capacity or the fluid step is a multiple of this value. For smaller values the CCDF is higher for all occupation levels. The time to put a packet of the maximum size on the line is then longer than the bin duration. In the fluid flow model the packet is segmented in smaller ones. For values of the bin duration between the multiples, the CCDF is also between the CCDF of the multiples but the sharp edges are smoothed out. These edges correspond to a large amount of packets with a size equal to the bin occupation on that edge.

[6]Complementary Cumulative Distribution Function

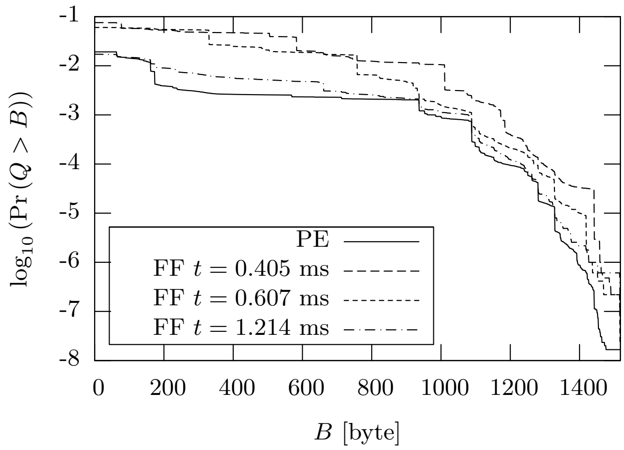

Figure 1: The CCDF of the buffer occupation Q of a packet-event simulation (PE) for the Bellcore LAN trace compared to the CCDF of the buffer occupation of fluid-flow simulations (FF) with different fluid steps for the same traffic trace. The transmission capacity equals 10 Mbps.

The step behavior of the simulated results is due to the distribution of the packet sizes in the Bellcore LAN trace. Figure 3 shows the complementary cumulative packet size distribution of the Bellcore LAN trace. A large step means a high probability that the corresponding number of traffic units are in the buffer. If a packet size is very probable and the number of packets in the buffer is small, the buffer occupation probability for the number of traffic units, that equals a highly probable packet size, will be very high. When multiple traffic traces are multiplexed, this step behavior will fade out.

2.3 Stationarity

Both PDFs, equation (5 and 8), are stable if the buffer occupation process Q_n is stationary. This is the case if the following assumptions are valid:

- $C_{\Delta t}$ is independent of n;
- the arrival fluid flow is stationary, i.e. $(A_{-n}, \ldots, A_0)$ has the same distribution as $(A_{-n-m}, \ldots, A_{-m})$ for every n and m;
- the mean traffic rate is lower than the capacity, i.e. $\mathbb{E}(A_n) \leq C$;
- the queue is empty at time $-\infty$.

The expression for Q_n can then be simplified [9] to:

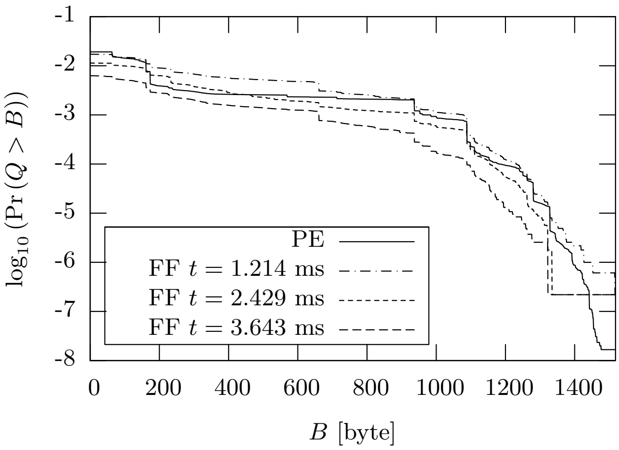

Figure 2: The CCDF of the buffer occupation Q of a packet-event simulation (PE) for the Bellcore LAN trace compared to the CCDF of the buffer occupation of fluid-flow simulations (FF) with different fluid steps for the same traffic trace. The transmission capacity equals 10 Mbps.

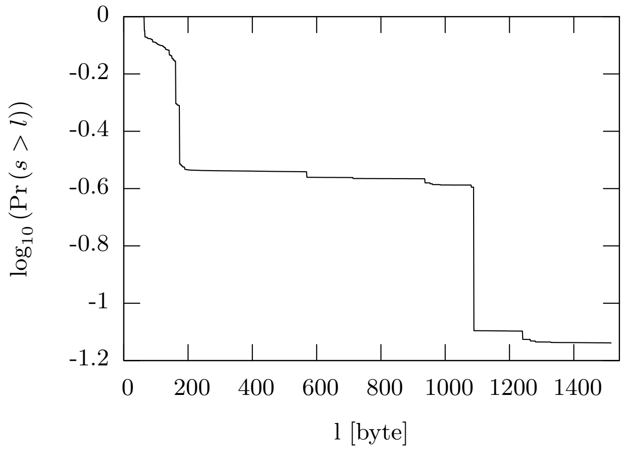

Figure 3: The complementary cumulative packet size distribution for the Bellcore LAN trace gives the probability that a packet size s is larger than a certain length l. The step behavior of the simulated graphs in figures 1 and 2 is due to the jumps in the distribution of the packet sizes in the traffic trace.

$$Q_n = \sup_{m \leq n} (S_{m,n} - m\Delta t C) \tag{10}$$

where C is the transmission capacity of the link, $\frac{t}{\Delta t} = m$ an integer number of bins, $S_{m,n} = \sum_{i=0}^{m-1} A_{n-i}$ the cumulative arrival process. The corresponding PDF, f_{SS}, is called the steady state distribution of the buffer occupation level.

$$f_{SS}(Q = B) = \frac{1}{N - m} \sum_{n=0}^{N-m+1} \mathbf{1}_B \left(\sup_{m \in \mathbb{N}_0} (S_{m,n} - m\Delta t C) \right) \tag{11}$$

Only in certain cases[7] this probability can be calculated analytically. For real traffic traces numerical approaches have to be used to find an estimate.

2.4 Rare-Event Approximation

In case of a high performance queuing network, it is very probable that the queue empties regularly. The probability that the queue load Q has a value B, is then small and the principle of the largest term can be used to move the supremum out of the probability. The PDF of the rare-event approximation, f_{RE}, becomes

$$f_{RE}(Q = B) = \sup_{m \in \mathbb{N}_0} \frac{1}{N - m} \sum_{n=0}^{N-m+1} \mathbf{1}_B (S_{m,n} - m\Delta t C) \tag{12}$$

The principle of the largest term stipulates that if a rare-event happens, it occurs in the most probable way. In a queuing system, this means that a specific time-scale, $t = m\Delta t$, dominates the build up of the queue load to a buffer occupation level B. The time parameter t can be interpreted as the most probable time starting from an empty queue that the buffer load reaches a number of traffic units B.

[7]For basic arrival distributions, i.e Poisson traffic, the calculation can be done but real traffic tends to behave differently [14].

Figure 4 shows the CCDF of the buffer occupation of fluid-flow simulations and of the rare event approximations for different transmission capacities C. The trace is the Bellcore LAN trace and the fluid steps corresponds to the maximum packet size, 1518 byte, divided by the corresponding transmission capacity. In case of a high transmission capacity the rare-event results are very similar to the fluid-flow simulations. The higher the transmission capacity, the lower the probability of a buffer overflow and the better the rare-event simulation. Moving the supremum out of the stationary PDF corresponds to the selection of the most probable path to overflow. For low transmission capacities the rare-event approximation considers only the path to overflow with the highest probability but the fluid-flow simulations counts all the different paths, which in this case have an impact on the buffer occupation. For higher transmission capacities, only one path to overflow has to be considered.

3. BACKGROUND TRAFFIC DESCRIPTOR

3.1 Large Deviations Approximation

The main problem with the trace-driven fluid-flow simulation is the large number of simulation runs needed to get a realistic traffic descriptor when several independent traffic streams are multiplexed. Each traffic trace can start in a random fluid step. The most natural approach is to generate the fluid arrivals for the different traffic streams in a cyclic way, e.g. fluid step A_N, the last of a traffic trace will be followed by fluid step A_0 of the same trace. The starting value of the fluid step for each traffic flow is randomized for each run. The final value is the mean of all the runs.

Large deviations theory gives a method to calculate exactly what happens when many independent traffic streams are multiplexed. The limiting regime of interest[8], the many flows regime, considers what happens when a queue is shared

[8]The large buffer asymptotic [3] can not be used due to the possible small values of B.

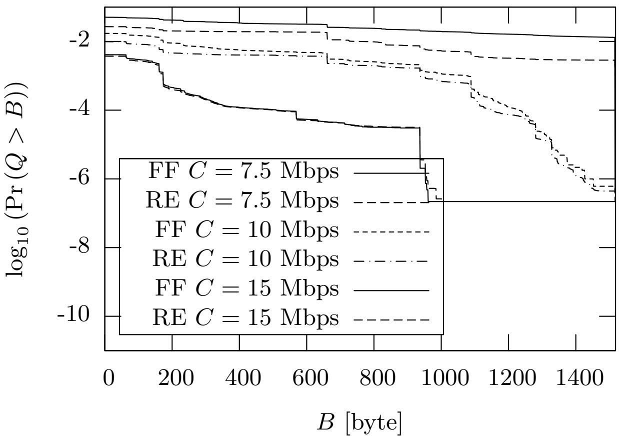

Figure 4: The CCDF of the buffer occupation Q of the fluid-flow simulations (FF) for the Bellcore LAN trace compared to the CCDF of the rare-event approximations (RE) for the same traffic trace. The transmission capacity varies between 7.5 Mbps (highest buffer occupation probability) and 15 Mbps (lowest buffer occupation probability).

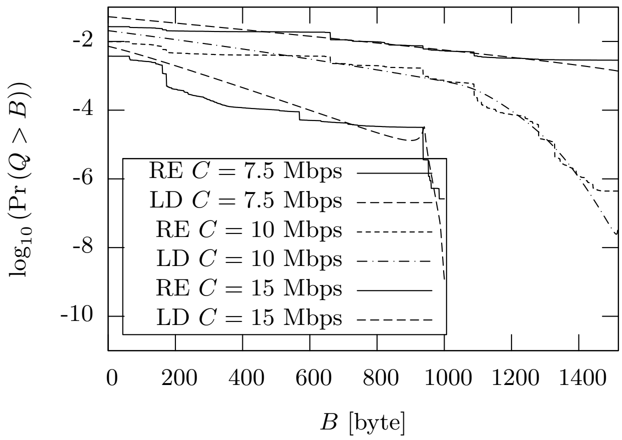

Figure 5: The CCDF of the buffer occupation Q of rare-event simulations (RE) for the Bellcore LAN trace compared to the CCDF of the large deviations calculation (LD) for the same traffic trace. The transmission capacity varies between 7.5 Mbps (highest buffer occupation probability) and 15 Mbps (lowest buffer occupation probability).

by a large number of independent traffic flows. This is certainly the case for a WAN or a large LAN background traffic stream.

Consider a single FIFO server queue, with N sources and constant service rate $C = cN$. Let $A_n^{(i)}$ be the number of traffic units arriving from source i during fluid step n. Assume that for each i, $\left(A_n^{(i)}, n \in \mathbb{Z}\right)$ is a stationary sequence of random variables, and that these sequences are independent of each other. The CCDF of the buffer occupation of the multiplexing of the different traffic streams in a FIFO queuing system for $N \to \infty$ can be estimated as [11]:

$$\Pr(Q \geq B) = \frac{1}{s_{\hat{t}}\sqrt{2\pi\sigma_{\hat{t}}^2 N}} \mathrm{e}^{-NI(b+ct)} \tag{13}$$

where $B = Nb$, $S_{m,n}^{(i)} = \sum_{j=0}^{m-1} A_{n-j}^{(i)}$, the moment generating function $\phi_{t=m\Delta t}^{(i)}(s) = \mathbb{E}\mathrm{e}^{sS_{m,n}^{(i)}}$ and the rate function

$$\begin{aligned} I(b+ct) &= \inf_{t:\frac{t}{\Delta t}\in\mathbb{N}_0} J_t(b+ct) \quad (14) \\ &= \inf_{t:\frac{t}{\Delta t}\in\mathbb{N}_0} \sup_{s\in\mathbb{R}^+} s(b+ct) - \sum_{i=1}^{n} n^i \log \phi_t^{(i)}(s) \end{aligned}$$

with N^i the number of flows from class i, $n^i = N^i/N$ and $\sigma_{\hat{t}}^2 = \frac{\mathrm{d}^2 J_t(b+ct)}{\mathrm{d}s^2} \left|(t = \hat{t}, s = s_{\hat{t}})\right.$. s_t can easily be found as

$$s_t = \arg_s \left(b + ct = \sum_{i=1}^{n} n^i \frac{\phi_t^{(i)'}(s)}{\phi_t^{(i)}(s)} \right) \tag{15}$$

and the expression for $\sigma_{\hat{t}}^2$ simplifies to

$$\sigma_{\hat{t}}^2 = \sum_{i=1}^{n} n^i \left(\frac{\phi_t^{(i)''}(s)}{\phi_t^{(i)}(s)} - \left(\frac{\phi_t^{(i)'}(s)}{\phi_t^{(i)}(s)} \right)^2 \right) \left|(t = \hat{t}, s = s_{\hat{t}})\right. \tag{16}$$

A linear search has to be performed to find the value of the optimum $\hat{t}$. This parameter, which is a multiple of the fluid step Δt, corresponds to the most probable buffer busy period before reaching an occupation level B starting from an empty buffer.

The estimate of equation (13) is exact for $N \to \infty$. The expressions can however be used when N is large but unknown, as for WAN traffic traces. References [1, 2] demonstrate that the equations (13, 14, 15 and 16) are valid for a traffic trace composed of a large number of independent and possible different traffic streams by putting $N = 1$ and performing the calculation on the complete trace.

Roughly speaking, this estimate gives an exponential approximation of the rare-event approach for the queuing behavior of multiplexed traffic streams where each traffic streams is composed of the packets from many independent sources.

Figure 5 shows the CCDF of the buffer occupation of the rare-event simulation and the large deviations calculation for the Bellcore LAN trace. The transmission capacity C is respectively 7.5 Mbps, 10 Mbps and 15 Mbps. The fluid step, 1.214 ms, corresponds to the maximum packet size, 1518 byte, divided by the transmission capacity. The large deviations approximation mimics very closely the rare-event results. For a small value of B the large deviations calculation overestimates the CCDF of the buffer occupation, for large values of B it underestimates the CCDF. In case of a high traffic load, e.g. a low transmission capacity, the logarithmic CCDF of the rare-event simulation is almost linear and the large deviations result, which is an exponential approximation, naturally fits the CCDF of the buffer occupation. In a low traffic load scenario the logarithmic CCDF is only piece-wise linear and the large deviations method has

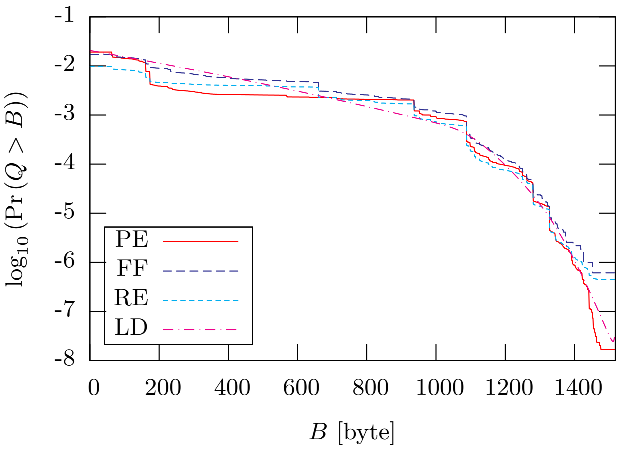

Figure 6: The CCDF of the buffer occupation Q for a packet-event simulation (PE), a fluid-flow simulation (FF), a rare-event approximation (RE) and a large deviations calculation (LD) for the Bellcore LAN trace with a transmission capacity of 10 Mbps and a fluid step of 1.214 ms.

more difficulties tracking the CCDF of the rare-event simulation. This can produce non physical results as shown for $C = 15$ Mbps and $B \approx 900$ byte.

The inaccuracy of the Large Deviations method has three main causes:

- the use of a fluid step smooths out the sharp edges of the CCDF of the packet-event simulations, which are due to the statistics of the packet size distribution of the trace; a wrong choice of bin can ruin the fluid-flow simulation;
- the rare-event approximation only considers the most probable path to overflow; the best results are obtained for a small traffic load, e.g. a high transmission capacity compared to the mean traffic rate[9];
- the large deviations calculation gives an exponential approximation to the rare event results which can be considered as piece-wise linear; the higher the load the better the approximation.

The first and last point can easily be remedied by considering the multiplexing characteristics of the large deviations approach.

Figure 6 shows the CCDF of the buffer occupation for a packet-event simulation, a fluid-flow simulation, a rare-event approximation and a large deviations calculation for the Bellcore LAN trace with a transmission capacity of 10 Mbps and a fluid step of 1.214 ms, corresponding to the maximum packet size, 1518 byte, divided by the transmission capacity.

3.2 Multiplexing traffic streams

It might be expected that when bursty and smooth traffic are multiplexed on a link that the bursty traffic will be smoothed out and the smooth traffic stream will become more bursty. Indeed, this happens in a router with a small number of inputs. But in the *Many Flows Scaling* regime[10] this is not the case. In other words, the individual traffic flows do not depend on the traffic mix at the router as long as the queue empties regularly with high probability. This is known as *decoupling* [16] and allows for the introduction of a intuitive descriptor of the stochastic properties of an individual traffic stream, called the *Effective Bandwidth* ($\alpha_t(s)$) of the fluid-flow [6].

$$\alpha_t(s) = \frac{1}{st} \log \phi_t(s) \tag{17}$$

In the large deviations limit, it makes sense to talk about the effective bandwidth of a single flow through a network as long as in each network device the service rate is higher than the mean arrival rate. In that case, the effective bandwidth of the departure flow at the last device will exactly be the same as the effective bandwidth of the arrival flow at the first device. The effective bandwidth is additive for independent sources and can also be understood in terms of admission regions. Suppose there are mN flows with effective bandwidth $\alpha_t(s)$ and nN flows with effective bandwidth $\beta_t(s)$. For what values of m and n does the system meet the quality of service constraint γ? In reference [1] the quality constraint is defined as

$$\Pr\left(Q^N \geq Nx\right) < e^{-\gamma N} \tag{18}$$

and the admissible region is

$$\bigcap_{t>0} \left\{ m, n : \exists s > 0 : m\alpha_t(s) + n\beta_t(s) < c + \frac{b}{t} - \frac{\gamma}{st} \right\} \tag{19}$$

This equation shows that the effective bandwidth can be used to characterize the trade-off between flows of different types.

The space-scale s is a parameter for the degree of multiplexing [4]. The more s approaches 0, the more the multiplexing will be efficient and the effective parameter tends to the mean rate of the traffic stream at the appropriate time-scale. If s becomes large, the flows will not multiplex very well and the effective bandwidth will be close to the peak rate of the traffic stream at the corresponding time-scale. For a fixed t, the effective bandwidth is strictly convex going from the mean rate to the peak rate of the traffic flow at time-scale t. Figure 7 gives a surface plot of the effective bandwidth of the Bellcore LAN trace. Ripples in the t-direction correspond to periodic components in the traffic stream.

Figure 8 shows the CCDF of the buffer occupation for a packet-event simulation of a randomized multiplexing of 25 times the Bellcore LAN traces and the CCDF of the large deviations approximation for the same traffic. The transmission capacity is 100 Mbps and the basic fluid step is 0.1214 ms, corresponding to the maximum packet size, 1518 byte, divided by the transmission capacity. The plotted CCDF of the buffer occupation for the packet-event simulation is the

[9]A mean traffic load of 40% saturates a typical WAN easily. This paper considers only non saturated networks. A Heavy Traffic approach can be used for saturated networks as explained in [8, 9].

[10]Experimental setups have shown that already for a rather small number of flows $N = 3$ for two traffic classes, the departure flows are mostly decoupled [16].

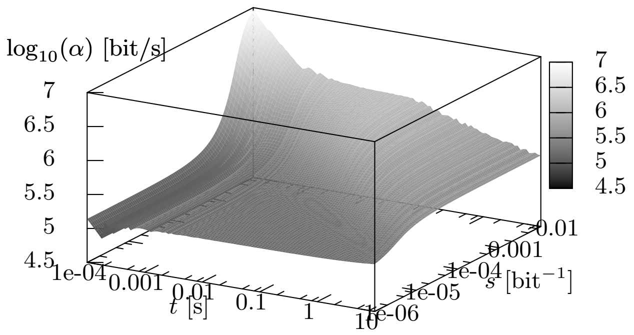

Figure 7: The effective bandwidth of the Bellcore LAN trace as a function of s the space or multiplexing parameter and t the time scale.

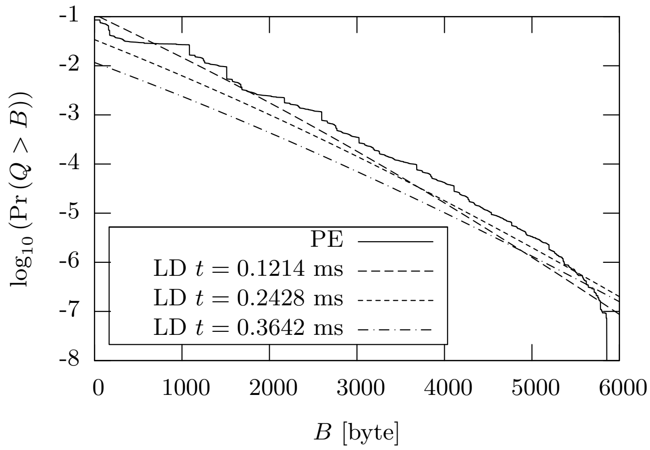

Figure 8: The mean CCDF of the buffer occupation Q for a packet-event simulation (PE) of a random multiplexed stream of 25 Bellcore LAN trace and the CCDF of the large deviations (LD) approximation for the same traffic mix with a transmission capacity of 100 Mbps and different fluid steps.

mean value of 25 independent simulation runs with randomized start indices for the 25 traffic streams. The calculation is done for the complete duration of the trace and indices corresponding to a time larger than the trace length are obtained modulo the number of packets in the trace. So the traffic is considered to be cyclic with a period equal to the trace length.

For a value of B around 5800 byte the packet-event simulation has a steep drop in the CCDF. This is caused by the minimum time between packets, the maximum packet length and the measurement of the traffic trace on a real network with a finite queue size. There is no significant difference between the 25 packet-event simulation runs for values of B less than 5000 byte. For larger values of B, the irregular drop in the CCDF for the different simulations makes the statistics unreliable.

The infimum calculation of the large deviations approximation prefers the plots corresponding to a fluid step of 0.1214 ms for values of B smaller than 3725 byte and to a fluid step of 0.2428 ms for values of B larger than 3725 byte.

4. HYBRID SIMULATION

4.1 Simulation Strategy

In a hybrid simulation scheme, the foreground traffic is handled in an adapted packet-event way. The workload of the queue still changes discontinuously on each event, but a number of virtual packets are fed into the queue before the foreground packet is added. The virtual packets are introduced to take the background traffic into account without doing a packet event simulation of both foreground and background traffic streams. The amount of network traffic due to the foreground stream has to be small compared to the background traffic.

The following strategy is used:

- if a new packet arrives in the network node and the number of traffic units already in the queue and the number of traffic units due to the background stream plus the size of the new packet are less than the queue length, the packet is queued after the virtual background packets, otherwise the packet is dropped;
- if a packet is processed, e.g. the last traffic unit is put on the wire, the following packet, real or virtual, can leave the queue and the load of the queue decreases with the size of the last packet.

Let t_i^A be the arrival time of the ith packet of the foreground traffic trace and t_i^L the time that the ith packet, real or virtual, leaves the queue. The workload of a queuing system with a finite buffer size changes depending on the arrival times, the leaving times, the buffer size B and $\hat{Q}_{t_i} = Q_{t_{i-1}} + A_i + V_i$

$$\text{if} \quad t = t_i^A : \quad Q_{t_i} = \hat{Q}_{t_i} \quad \text{if } \hat{Q}_{t_i} \leq B \tag{20}$$

$$\text{if} \quad t = t_i^A : \quad Q_{t_i} = \min\left(Q_{t_{i-1}} + V_i, B\right) \quad \text{if } \hat{Q}_{t_i} > B \tag{21}$$

$$\text{if} \quad t = t_i^L : \quad Q_{t_i} = Q_{t_{i-1}} - P_i \tag{22}$$

where Q_t is the workload of the queue, A_i the size of the arriving packet on time t_i^A and P_i the size of the leaving packet, real or virtual, on time t_i^L. t_{i-1} is the previous changing times of the load of the queue. V_i is the number of traffic units due to the virtual packets. Every virtual packet is scheduled before the arriving packet. t_i^L can be calculated considering the time needed to put a packet on the line

$$\text{if the line is idle} \quad : \quad t_i^L = t_i^A \tag{23}$$

$$\text{if the line is in use} \quad : \quad t_i^L = t_{i-1}^L + \frac{P_{i-1}}{C} \tag{24}$$

where C is the transmission capacity in traffic units of the network link.

The number of virtual traffic units V_i is calculated by sampling the buffer occupation PDF of the background stream to get the sampled value $\widetilde{V}_i$:

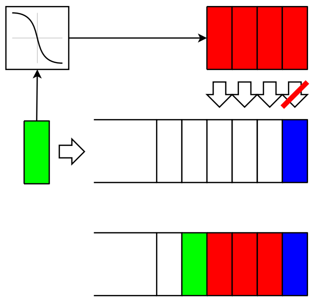

Figure 9: Drawing of the hybrid simulation strategy. A foreground packet (green) enters a queuing system. A number of virtual packets (red) is created depending on the sampled value of the PDF of the buffer occupation of the background traffic stream, the PDF of the background packet sizes and the number of traffic units already in the system (blue).

$$\begin{aligned} \text{if} \quad \widetilde{V}_i > Q_{t-1} \quad &: \quad V_i = \widetilde{V}_i - Q_{t-1} && (25) \\ \text{if} \quad \widetilde{V}_i \leq Q_{t-1} \quad &: \quad V_i = 0 && (26) \end{aligned}$$

If the number of virtual traffic units V_i is less than the maximum packet size of the background stream, only one virtual packet with size V_i is created. It this is not the case, the packet sizes are calculate by sampling the PDF of the packet sizes of the background traffic. The last virtual packet has a size equal to the remaining traffic units once the number of traffic units left is less than the maximum packet size. Figure 3 shows a plot of the PDF of the packet sizes of the Bellcore LAN trace.

Figure 9 shows a drawing of the hybrid buffer strategy. A foreground packet enters the queuing system. In this example 3 virtual background packets are created corresponding to the sampled value of the buffer occupation PDF of the background traffic minus the number of traffic units from the packet that was already in the queue. The size of the virtual packets is obtained by sampling the packet sizes PDF of the background traffic stream. The virtual packets are put into the queue before the foreground packet. The foreground packet sees a buffer occupation of the sum of the traffic units of the 3 virtual packets and the existing packet in the system.

4.2 Background PDF

The large deviations approximation of equation (13) gives an estimate for the CCDF of the buffer occupation due to the background stream. The PDF of the buffer occupation can be obtained by subtraction of the CCDF for a queue load $B-1$ and B.

$$f_{LD}(Q = B) = \Pr(Q > B) - \Pr(Q > B + 1) \qquad (27)$$

The CCDF of the buffer occupation in the large deviations approximation, equation (13), can directly be obtained from the background trace, e.g. a list of pairs (t_i, A_i) consisting of an arrival time, t_i, and a packet size, A_i. The moment generating function $\phi_t(s) = \mathbb{E}e^{sS_t(\tau)}$ is needed for the calculation of the CCDF.

$$\phi_t(s) = \frac{1}{T-t}\int_{\tau=0}^{T-t} e^{sS_t(\tau)} d\tau$$

For a fixed time step t, the cumulative arrival process $S_t(\tau)$ and also $\phi_t(s)$ are piece-wise constant functions which jump at the time epochs that a packet enters or leaves the recording interval t. Let t_i^A be the entering time of the ith packet of the traffic trace in the fluid step t and t_i^L the leaving time of the same packet. $S_t(\tau)$ changes

$$\begin{aligned} \text{if} \quad \tau_j = t_i^A \quad &: \quad S_t(\tau_j) = S_t(\tau_{j-1}) + A_i && (28) \\ \text{if} \quad \tau_j = t_i^L \quad &: \quad S_t(\tau_j) = S_t(\tau_{j-1}) - A_i && (29) \end{aligned}$$

where A_i is the size of the arriving packet on time t_i^A and τ_{j-1} is the previous jump time. The list of pairs $(\tau_j, S_t(\tau_j))$ can be transformed in a convenient new list of pairs (v_t, θ_{t,v_t}) for the calculation of $\phi_t(s)$ where

$$\theta_{t,v_t} = \sum_{\tau_j=0}^{\tau_{\max}} \tau_j \mathbf{1}_{v_t}\left(S_t(\tau_j)\right) \qquad (30)$$

and $\tau_{\max} = \max(\tau_j)$. Only the total time θ_{t,v_t} that $S_t(\tau_j) = v_t$ is important for the calculation of $\phi_t(s)$. The moment generating functions and derivatives in the variable s reduces to

$$\begin{aligned} \phi_t(s) &= \frac{1}{T-t}\sum_{v_t=0}^{v_{\max,t}} \theta_{t,v_t} e^{sv_t} && (31) \\ \phi_t^{'}(s) &= \frac{1}{T-t}\sum_{v_t=0}^{v_{\max,t}} \theta_{t,v_t} v_t e^{sv_t} && (32) \\ \phi_t^{''}(s) &= \frac{1}{T-t}\sum_{v_t=0}^{v_{\max,t}} \theta_{t,v_t} v_t^2 e^{sv_t} && (33) \end{aligned}$$

where $v_{\max,t} = \max(S_t(\tau_j))$. This simplifies the calculation of s_t, in equation (15), and of σ_t^2, in equation (16). It also reduces the memory consumption by not storing the list of pairs $(\tau_j, S_t(\tau_j))$.

5. NUMERICAL RESULTS

To validate this novel hybrid simulation concept, a test scenario is evaluated in which 25 Bellcore LAN traces, the background traffic, are multiplexed over a 100 Mbps and are fed into a queuing system together with a variable number n of video streams. These video streams, the foreground traffic, are H.264 coded and are combined with a corresponding voice stream that uses a Siren14 codec. The Bellcore LAN trace consists of one million packets with a mean rate of 1.105 Mbps captured in 3.143 ks. One stream of video and voice has a mean rate of 298 kbps and has a duration of 329 s. During this time an interview was filmed in a busy street. The video is streamed and captured on a unsaturated LAN.

The packet-event simulation has a simulation time, e.g. the simulated time, that equals the capture time of the Bellcore trace. During the simulation the video stream is continuously played and all the multiplexing is done in a cyclic way with a random start index of the traces. During one simulation 50 million Bellcore LAN events and $746410 \cdot n$

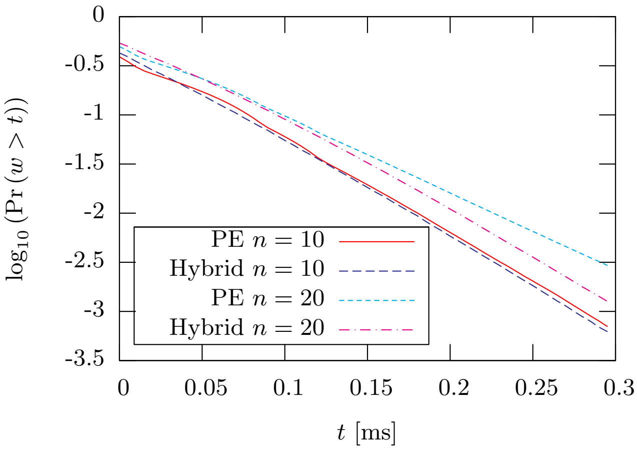

Figure 10: The CCDF of the waiting time w of a packet-event simulation and a hybrid simulation for different values of the number of foreground video streams n. The background traffic is composed of 25 randomly multiplexed Bellcore LAN traces. The transmission capacity equals 100 Mbps.

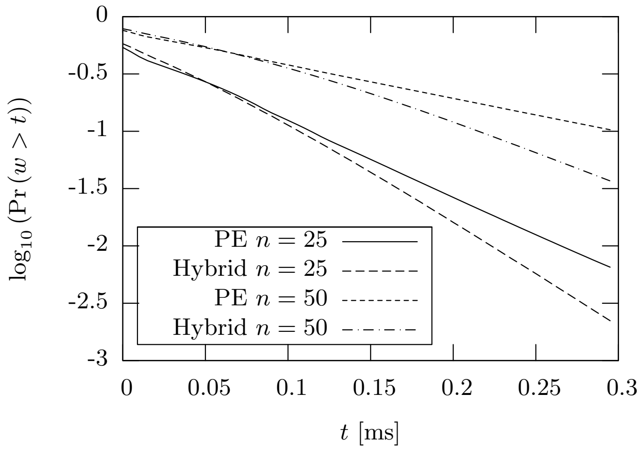

Figure 11: The CCDF of the waiting time w of a packet-event simulation and a hybrid simulation for different values of the number of foreground video streams n. The background is composed of 25 randomly multiplexed Bellcore LAN traces. The transmission capacity equals 100 Mbps.

video stream events are handled. The plotted results, figures 10 and 11, are the mean values of 25 independent runs with randomized start indices.

The hybrid simulation has the same simulation time to ensure that both simulation have the same accuracy which equals $1/(746410 \cdot n)$. Only ten runs had to be carried out before a 95% confidence interval could be obtained.

The waiting time w of a video stream packet in the queuing system is chosen as performance indicator. The waiting time corresponds to the time a packet is queued up in the buffer before being transmitted. Figures 10 and 11 show the CCDF of the waiting times for different values of n[11]. For small values of the mean waiting time the hybrid approach overestimates the CCDF with a small amount for all values of n. This is directly related to the slightly higher value of the CCDF of the buffer occupation for the large deviations calculation compared to the packet-event simulation in figure 8. For larger values of the waiting time the hybrid method underestimates the CCDF of the waiting time. If n equals 20 or larger, there is a significant difference between the hybrid simulation and the complete packet-event simulation for high values of the waiting time. The buffer occupation which is seen by an incoming foreground packet in the adapted packet-event simulation, is only composed of background packets. For a high value of n the buffer occupation distribution is no longer only dependent on the background stream but also the foreground packets have to be included.

Table 1 shows the proportion of the mean rate of the foreground video stream to the total traffic rate as a function of n, the number of foreground video streams with $\rho = \frac{\mu_{\text{video}}}{\mu_{\text{Bellcore}}+\mu_{\text{video}}}$ the traffic proportion of the foreground

n	ρ	R	D
5	0.051	0.0695	0.0054
10	0.097	0.1298	0.0065
15	0.139	0.1829	0.0132
20	0.177	0.2299	0.0270
25	0.212	0.2718	0.0451

Table 1: Comparison of the traffic proportion of the foreground video stream to the total traffic rate ρ, the number of generated events R, and the mean quadratic difference over 60 samples of the logarithmic CCDFs, D between the hybrid simulation and the corresponding complete packet-event simulation as a function of the number of foreground video streams.

video stream to the total traffic rate, $R = \frac{\#\text{H}}{\#\text{PE}}$ the proportion of generated events and $D = \overline{(CCDF_H - \text{CCDF}_{\text{PE}})^2}$ the mean quadratic difference over 60 samples of the logarithmic CCDFs between the hybrid simulation (H) and the corresponding complete packet-event simulation (PE). The background traffic is composed of 25 randomly multiplexed Bellcore LAN traces and the transmission capacity equals 100 Mbps. For a traffic proportion of 10%, only 13% of the events are needed to get a mean quadratic difference of the logarithmic values of the CCDF of less than 1%. The run time of the large deviations calculation is moderate[12] compared to the run time of one packet-event simulation. If the number of runs is counted for the packet-event simulations to converge in a 95% confidence interval, the unoptimized hybrid method in the test scenario realizes a 20 fold speed up compared to the complete packet-event simulation.

6. CONCLUSION

[11]Values of n larger than 25 clearly violate the hypothesis that the mean rate of the foreground traffic has to be small compared to the mean rate of the background traffic. Simulations for higher values have been carried out to obtain the limits of the working domain of the hybrid method.

[12]30% of one run of the packet-event simulation in the test scenario

In this paper, a novel hybrid simulation technique is detailed and evaluated. The results of a large deviations fluid flow analysis of the background traffic stream is integrated in the packet-event simulation of the foreground traffic. The large deviations calculation gives the PDF of the buffer occupation for a x-times multiplexed background traffic. This calculation depends on the transmission capacity. During the packet-event simulation, the obtained PDF can be sampled to add virtual packets with a combined size equal to the sampled value, each time a foreground packet arrives in the queuing system. This give the opportunity to easily integrate this novel technique in an existing simulator as ns-2 or J-Sim without the overhead and the complexity to synchronize between the fluid-flow simulator and the packet-event simulator. To obtain accurate results, the mean rate of the foreground stream has to be small compared to the mean rate of the background traffic. Good results are obtained when the mean rate of foreground traffic is less than 10% of the mean rate of the background traffic stream.

Both captured traces or synthetic traffic can be used as input for the large deviations calculation and the packet-event simulation. The only requirement for the background streams is that they are composed of many independent traffic streams so that the *Many Source Large Deviations* approximation is valid. In the examples the Bellcore LAN trace, which contains long-range dependent and self-similar traffic, is used with satisfactory results.

In a rudimentary queuing system, the number of discrete events is reduced and the total run time of the hybrid simulation is only 5 % of the run time of a packet-event simulation. Recent traffic traces contain billions of packets and similar reductions can mean a great speed up of the simulations without sacrificing the detailed packet-event results for the foreground applications.

In a future paper, the results of the application of this hybrid method to a large number of recent traffic traces as background streams and to several foreground applications will be presented. Not only the mean waiting time of a foreground packet in the queue will be considered but also other performance indicators, e.g. the jitter of the packets of the foreground traffic and the drop rate.

Topics for further research also include:

- integrating the traffic of the foreground application into the large deviations calculation to eliminate the constraint that the mean rate of the foreground traffic has to be small compared to the mean rate of the background stream;
- simulating a more complicated network topology, e.g. tandem queues;
- integrating the hybrid method in an existing packet-event simulator;
- extending the FIFO queue to priority queuing and generalized processor sharing; this allows to consider the impact of quality of service constraints of the background stream;
- adapting the model to include wireless base-stations and ad-hoc networks based on a generalized processor sharing approach where each device is considered to have one queue and the medium is shared between all the devices.

7. REFERENCES

[1] C. Courcoubetis, V. Siris, and G. Stamoulis. Application of the many sources asymptotic and effective bandwidths to traffic engineering. *Telecommunication Systems*, 12:167–191, 1999.

[2] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a switch handling many traffic sources. *Journal of Applied Probability*, 33:886–903, 1996.

[3] A. Ganesh, N. O'Conell, and D. Wischik. *Big Queues*. Springer, 2004.

[4] R. Gibbens. Traffic characterization and effective bandwidths for broadband network traces. In *Stochastic Networks: Theory and Applications*, volume 4 of *Royal Statistical Society Lecture Notes Series*, pages 169–180. Oxford University Press, 1996.

[5] Y. Gu, Y. Liu, and D. Towsley. On integrating fluid models with packet simulation. In *IEEE INFOCOM'04*, volume 4, pages 2856–2866, 2004.

[6] F. Kelly. Notes on effective bandwidths. In *Stochastic Networks: Theory and Applications*, volume 4 of *Royal Statistical Society Lecture Notes Series*, pages 141–168. Oxford University Press, 1996.

[7] C. Kiddle, R. Simmonds, C. Williamson, and B. Unger. Hybrid packet/fluid flow network simulation. In *PADS'03*, pages 143–152, 2003.

[8] B. Lauwens, B. Scheers, and A. Van de Capelle. Characterization of background traffic in hybrid network simulation. In *NATO IST-062*, 2006.

[9] B. Lauwens, B. Scheers, and A. Van de Capelle. Buffer occupation probability of trace-driven background stream in hybrid simulation. In *IEEE ICC 2007*, June 2007.

[10] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Trans. Netw.*, 2:1–15, 1994.

[11] N. Likhanov and R. Mazumdar. Cell loss asymptotics in buffers fed with large number of independent stationary sources. In *IEEE INFOCOM'98*, volume 1, pages 339–346, 1998.

[12] B. Melamed, S. Pan, and Y. Wardi. Hybrid discrete-continuous fluid-flow simulation. In *SPIE*, volume 4526, pages 263–270, 2001.

[13] B. Melamed, S. Pan, and Y. Wardi. HNS: A Streamlined Hybrid Network Simulator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 14(3):1–27, 2004.

[14] V. Paxson and S. Floyd. Wide-area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, pages 226–244, 1995.

[15] G. F. Riley, T. M. Jaafar, and R. M. Fujimoto. Integrated fluid and packet network simulations. In *MASCOTS'02*, pages 511–518, 2002.

[16] D. Wischik. The output of a switch or effective bandwidths for networks. *Queueing Syst.*, 32:383–396, 1999.