Application of the Stochastic Fluid Flow Model for Bottleneck Identification and Classification

Kurt Tutschku* and Markus Fiedler†

COST 279TD(02)046‡

Abstract

Network performance management is facing the challenge of provisioning advanced services with stringent delay and throughput requirements. For this reason, shortages of network capacity implying delay or loss, so-called bottlenecks, have to be identified and to be classified. The latter tasks imply the need for tractable analytical performance models. We identify the *stochastic fluid flow* model, which is based on bit rates and its statistics, as a possible candidate of being capable of describing qualitative behaviour of bottlenecks. In this work, we show how total and individual bit rate statistics at the output of a bottleneck are calculated via the *stochastic fluid flow* model. ¿From this, we deduce some general behaviours and classification criteria for bottlenecks.

Keywords: Bottleneck identification, bottleneck classification, bit rate statistics, stochastic fluid flow model, network performance management

1 Introduction

Advanced network services such as Voice-over-IP challenge IP network operation and IP traffic engineering considerably. They generate data streams which are increasingly sensitive to specific throughput and delay requirements. If the requirements are not met, the services may degrade substantially and might become of no use.

In packet-oriented networks, the traffic streams contend for finite resources provided by the communication network. They interfere with each other as well as with network entities. If the capacity limits of the network entities are too tight, the entities might change the characteristics of the data streams too rigorously. The requirements of the application are missed. A major task in network performance management is to eliminate this destructive behaviour. The first step in this procedure is to identify the limiting capacities, commonly denoted as *bottlenecks*. The bottlenecks have to be characterized in terms of location and bandwidth.

*Institute of Computer Science, University of Würzburg, Am Hubland, D-97074 Würzburg, Germany. Phone: +49 931 8886641, fax: +49 931 8886632, email: tutschku@informatik.uni-wuerzburg.de

†Blekinge Institute of Technology, Department of Telecommunications and Signal Processing, SE-371 79 Karlskrona, Sweden. Phone: +46 708 537339, fax: +46 455 385657, email: markus.fiedler@bth.se

‡This work has been sponsored by COST 279 through a Short Term Scientific Mission “On the importance of bandwidth statistics for network management” (Karlskrona, Sept 3–12, 2002).

Typical network performance management and capacity planning cycles [?] are based on observing the network state. The state is typically determined by flow, throughput or load measurements. The observed performance values are usually throughput time averages which are obtained on long time scales. For example the widely used performance monitoring tool MRTG considers observation intervals of five minutes as default [?]. On the other hand, many of today's network performance models for identifying bottlenecks are either not accurate enough or too complex for daily use. Most analytical bottleneck models focus on the analysis of packet delay and packet loss. The transfer of these results into throughput values, which can be directly observed, is in general difficult. As a consequence, many network administrators base their performance management actions on their experience rather than on performance models.

To facilitate model-based bottleneck identification, it is necessary to align measurement-based approaches with tractable analytical performance models. In particular, an analytical model should directly address throughput and load measurements. They should predict reliably the qualitative behaviour of the network.

Taking these requirements into account, the *stochastic fluid flow* model emerges as an interesting candidate for an analytical performance model for bottleneck identification. The fluid flow model considers time averages on small time scales for throughput. With this model, one can easily obtain analytical estimates on the impact of network entities on the throughput for data streams. The predicted throughput statistics appear to be easily measurable in practice as bit rate statistics for flows entering and leaving a bottleneck.

In this work we show how the *stochastic fluid flow* model can be used to predict the effects of bottlenecks on data streams. We describe how to calculate output rate densities for individual traffic flows and their superposition. From the numerical results, which are obtained from standard fluid flow analysis for Markov-Modulated Rate Processes (MMRP), we analyse how the bottleneck qualitatively changes the total and individual bit rate statistics of data streams.

The paper is structured as follows. In Section 2 we define the meaning of the term *bottleneck* in this work. In Section 3 we describe a bottleneck analysis using the stochastic fluid flow model. Section 4 illustrates the findings of the previous section with a numerical example. Section 5 provide a brief overview on some related work and Section 6 provides a short summary and outlook to future work.

2 Bottleneck Characteristics

Packet-oriented networks rely on providing resources on-demand. Due to the the stochastic nature of the data streams, the finite resources of the network entities passed by these streams may be for some time or in general not sufficient to accommodate the requirements. This view provides directly a first definition of the term bottleneck:

Definition: A *bottleneck* within a communication facility is a temporary or permanent lack of capacity compared to the requirements of information streams to be handled. A bottleneck alters the characteristics of an information streams passing through it by introducing delay and/or loss.

The capacity of networks is in general determined by the the forwarding capacity of switches, the relaying capacity of routers, or the link capacity which is typically measured in *bits per second* (bps). In our investigation we are at first interested in bottlenecks with constant capacities and their qualitative impact on passing data streams. We distinguish between two types of bottlenecks: *a)* bottlenecks with small buffers; these kind of bottlenecks will introduce medium and limited delay but quite large packet loss in data streams, and *b)* bottlenecks with large buffers; these bottleneck are expected to cause large packet delays but are responsible for only small packet loss values.

In next section will outline how these types of bottlenecks can be modeled and identified with the *stochastic fluid flow* model.

3 Bottleneck Analysis Using the Stochastic Fluid Flow Model

Being one among other performance models, the stochastic fluid flow model has some appealing properties in the context of our study:

- The stochastic fluid flow model already works with rates, which primarily describe the intensity of particle stream, but may be measured in bps. Assuming a suitable discretization in time ΔT, fluid flow models describe the average flow of packets during ΔT pretty well.
- The stochastic fluid flow model focuses on the consequences of overload on queuing (macro behavior), but not on short-term queuing behavior, i.e. single packets having to wait for their turn to be transmitted (micro behavior).

In the sequel, we are looking at how bit rate distributions of data streams are changed when passing a bottleneck. We do this by investigating the flow of fluid particles through the buffer over time.

3.1 Basic Assumptions

Denote an individual traffic stream by index i. We assume σ states $S = 0, \ldots \sigma - 1$ during which the total input bit rate $r_s = \sum_i r_s^{(i)}$ and the bit rate contributions from individual traffic streams $r_s^{(i)}$ are constant. The mean state duration of the states are denoted by $\mathbf{E}[T_s]$ and the state probabilities by $\pi_s = \Pr\{S = s\}$, respectively.

The bottleneck parameters are given by the constant capacity C and the buffer size K. The buffer content is denoted by X with $0 \leq X < K$. In case of capacity shortage, both buffer and outlet are used by the participating streams in a fair-share manner, which is reflected by the factor $\gamma_s^{(i)} = \frac{r_s^{(i)}}{r_s}$. With respect to the buffer size, we focus on the extreme cases *infinite buffer* and *zero buffer*; the latter corresponds to a small buffer on packet level that allows for packet synchronization, but is not able to manage situations of temporary overload. However, finite buffer sizes are left for further study.

3.2 Infinite Buffer

Let the joint buffer content distribution be denoted by $\vec{F}(x)$, $F_s(x) = \Pr\{X \leq x \wedge S = s\}$ and its complement $\vec{G}(x)$, $G_s(x) = \Pr\{X > x \wedge S = s\}$. Depending on the source model, these values can be obtained by standard fluid flow analysis [**?**, **?**, **?**] et al. Of special interest in our context are

- the probability that the buffer is empty in state s, $F_s(0)$, which implies that the output bit rates equal to the input bit rates;
- the probability that the buffer is not empty in state s, $G_s(0)$, which implies that the output is saturated, i.e. the total output rate is equal the bottleneck capacity C, while the participating streams have to share this capacity amongst each other, receiving a share of $\gamma_s^{(i)} C$.

Please, observe that $F_s(0) + G_s(0) = \pi_s$. Depending on the drift $d_s = r_s - C$, we may identify the following types and sets of states:

- overload state $s \in \mathcal{S}^o$: the drift is positive, i.e. $d_s > 0$, the buffer fills up and is never empty. In this case, $F_s(0) = 0$, $G_s(0) = \pi_s$ is valid.
- Underload state $s \in \mathcal{S}^u$: the drift is negative, i.e. $d_s < 0$, the buffer content sinks until the buffer is empty. In this case, usually, $F_s(0) \neq 0$ and $G_s(0) \neq 0$ is valid.
- Indifferent state $s \in \mathcal{S}^i$: the drift is zero, i.e. $d_s = 0$, which means that the buffer content is frozen. In this case, usually, $F_s(0) \neq 0$ and $G_s(0) \neq 0$ is true.

The total set of states is denoted as $\mathcal{S} = \mathcal{S}^o \cup \mathcal{S}^i \cup \mathcal{S}^u$.

Let state s begin at time t_s and at buffer level x_s with the first particle arriving to the buffer. This particle leaves the buffer at time t_s^{o}, i.e. experiences a delay τ_s^{o}. At time $t_{s'} = t_s + T_s$, state s ends at a buffer level $x_{s'}$. The last particle from state s leaves the buffer at $t_{s'}^{\mathrm{o}}$, i.e. with a delay of $\tau_{s'}^{\mathrm{o}}$. Fluid particles from state s that have been entering the buffer during T_s need $\tau_s = t_{s'}^{\mathrm{o}} - t_s^{\mathrm{o}}$ to escape from the buffer. The latter time interval may consist of two sub-intervals:

- The time interval τ_s^ν denotes the time during which fluid originating from state s flows through the outlet while the buffer is *non-empty*.
- The time interval τ_s^ε denotes the time during which fluid originating from state s flows through the outlet while the buffer is *empty*.

We define the following probabilities:

- ν_s denotes the joint probability that buffer is non-empty and the fluid leaving the buffer originates from state s; the link is saturated.
- ε_s denotes the joint probability that buffer is empty and and the fluid leaving the buffer originates from state s; the link is not saturated.

3.2.1 Overload State

Figure 1 shows how the buffer content changes over time during a typical overload state. The buffer content rises ($x_{s'} > x_s$) because of $d_s > 0$. Particles originating from state s pass the outlet during the time $\tau_s = \tau_s^\nu > T_s$. From Figure 1, we see that $\tau_s^\nu = t_{s'} - t_{s'}^\star$, which leads to

$$\tau_s^\nu = \frac{r_s}{C} T_s \,, \tag{1}$$

which is valid independently of x_s. During this time, the buffer puts out fluid at maximal bottleneck speed C, while traffic stream i gets a share of $\gamma_s^{(i)} C$.

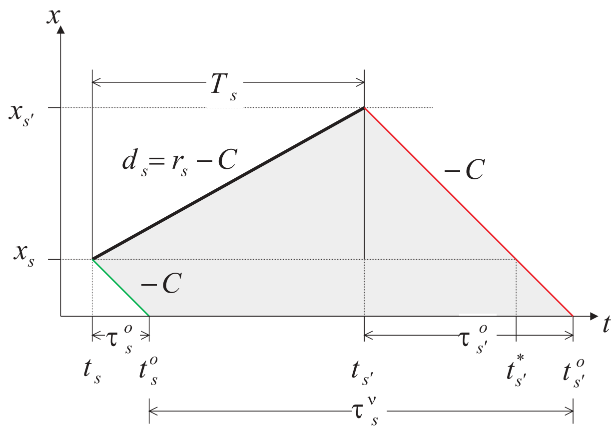

Figure 1: Buffer behaviour in overload state.

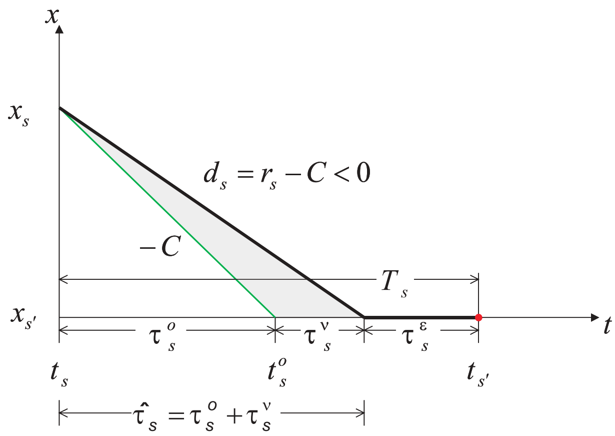

Figure 2: Buffer behaviour in underload state.

We now proceed by taking average over time, which leads to

$$\mathbf{E}[\tau_s^\nu] = \frac{r_s}{C}\,\mathbf{E}[T_s] \ . \tag{2}$$

As state s appears with probability π_s, the probability that the link is saturated from fluid from state s is given by

$$\nu_s = \frac{r_s}{C}\,\pi_s \quad \forall s \in \mathcal{S}^o \,, \tag{3}$$

while due to $F_s(0) = 0$, there is no chance that the input flow gets to the output right away:

$$\varepsilon_s = 0 \quad \forall s \in \mathcal{S}^o \,. \tag{4}$$

3.2.2 Underload State

Figure 2 illustrates how the buffer content changes over time during a typical underload state. We obtain:

$$\tau_s^\nu = \frac{r_s}{C}\,\tau_{s'} \ ; \tag{5}$$

$$\tau_s^\varepsilon = T_s - \tau_{s'} \ . \tag{6}$$

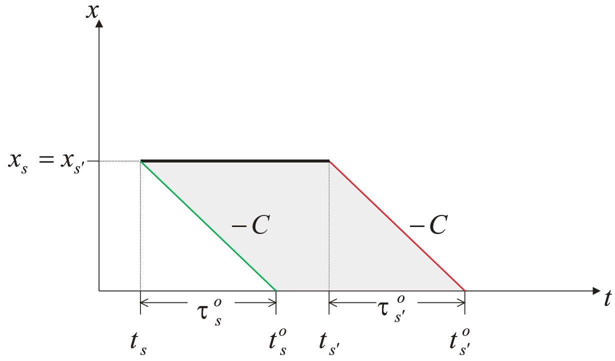

Figure 3: Buffer behaviour in indifferent state.

Taking expectations leads to:

$$\mathbf{E}[\tau_s^\nu] = \frac{r_s}{C}\left(\mathbf{E}[\tau_s^o] + \mathbf{E}[\tau_s^\nu]\right); \tag{7}$$

$$\mathbf{E}[\tau_s^\varepsilon] = \mathbf{E}[T_s] - \mathbf{E}[\tau_s^o] - \mathbf{E}[\tau_s^\nu]\,. \tag{8}$$

As probabilities, we obtain the contribution to the probability that the link is saturated as

$$\nu_s = \frac{r_s}{C}\,G_s(0) \quad \forall s \in \mathcal{S}^u \tag{9}$$

and the contribution to the probability that the input rate is visible at the output as

$$\varepsilon_s = F_s(0) \quad \forall s \in \mathcal{S}^u\,. \tag{10}$$

3.2.3 Indifferent State

In a state s with an input rate matching the capacity, the buffer content remains unchanged, see Figure 3. We see right away that $\tau_{s'}^o - \tau_s^o = t_{s'} - t_s = T_s$. Taking expectations and going over to probabilities, we have to distinguish two cases in which the buffer content remains at its level:

1. $x_s > 0$: The buffer remains non-empty:

$$\nu_s = G_s(0) \quad \forall s \in \mathcal{S}^i\,. \tag{11}$$

2. $x_s = 0$: The buffer remains empty ($\tau_{s'}^o = t_{s'}, \tau_s^o = t_s$):

$$\varepsilon_s = F_s(0) \quad \forall s \in \mathcal{S}^i\,. \tag{12}$$

3.3 Zero Buffer

Things get much simpler in the bufferless case, as there is no chance of transferring any fluid into the domain of an adjacent state due to $x_s = x_{s'} = 0$. We reuse the some of the notions introduced before in the sense that we merely relate them to whether the link is saturated or not.

3.3.1 Overload State

The capacity of the bottleneck is shared by all participating streams:

$$\nu_s = \pi_s \quad \forall s \in \mathcal{S}^o \tag{13}$$

$$\varepsilon_s = 0 \quad \forall s \in \mathcal{S}^o \tag{14}$$

3.3.2 Underload and Indifferent State

There is enough capacity for all streams so that the input rates appear unchanged at the output:

$$\nu_s = 0 \quad \forall s \in \mathcal{S}^u \cup \mathcal{S}^i \tag{15}$$

$$\varepsilon_s = \pi_s \quad \forall s \in \mathcal{S}^u \cup \mathcal{S}^i \tag{16}$$

3.4 Bit Rate Statistics

Denote R^{o} as the total output bit rate and $R^{(i),\mathrm{o}}$ as the contribution of stream i.
Theorem 1: The joint probabilities ν_s and ε_s of all states sum up to one.

$$\sum_{s \in \mathcal{S}} \nu_s + \varepsilon_s = 1 \tag{17}$$

The proof is trivial for the zero buffer case: $\sum_{s \in \mathcal{S}} \pi_s = 1$. Unfortunately, this is not necessarily the case for the infinite buffer, as $F_s(0)$ and $G_s(0)$ are seldomly available in closed form.

Total Output Bit Rate

Finally, the bit rate statistics for the total output bit rate is composed as follows:

$$\Pr\{R^{\mathrm{o}} = r\} = \sum_{s \in \mathcal{S}^u \cup \mathcal{S}^i : r = r_s} \varepsilon_s + \sum_{s \in \mathcal{S} : r = C} \nu_s \tag{18}$$

Theorem 2.1: As there is no loss in the infinite buffer case, the expectations of the total bit rate is the same at input and output:

$$\mathbf{E}[R^{\mathrm{o}}] = \mathbf{E}[R] \tag{19}$$

Individual Output Bit Rate

$$\Pr\{R^{\mathrm{o}} = r\} = \sum_{s \in \mathcal{S}^u \cup \mathcal{S}^i : r = r_s^{(i)}} \varepsilon_s + \sum_{s \in \S : r = \gamma_s^{(i)} C} \nu_s \tag{20}$$

Theorem 2.2: As there is no loss in the infinite buffer case, the expectations of the individual bit rates of stream i are the same on input and output:

$$\mathbf{E}\left[R^{(i),\mathrm{o}}\right] = \mathbf{E}\left[R^{(i)}\right] \quad \forall i \tag{21}$$

4 Numerical Example

4.1 Description

Up to now, we did not assume more than a generic *stochastic fluid flow* model from which we can derive the probabilities $F_s(0)$ and $G_s(0)$ for each state s, respectively. To be able to get some indications on the qualitative impacts of bottlenecks on traffic streams and their bit rate statistics through a numerical example, we assume that the aggregate

r	$\Pr\{R^o = r\}$				
	$C = 11$	$C = 9$	$C = 7$	$C = 5$	$C = 3$
1	5.63×10^{-2}	5.63×10^{-2}	5.63×10^{-2}	5.63×10^{-2}	5.63×10^{-2}
2	1.88×10^{-1}	1.88×10^{-1}	1.88×10^{-1}	1.88×10^{-1}	1.88×10^{-1}
3	2.82×10^{-1}	2.82×10^{-1}	2.82×10^{-1}	2.82×10^{-1}	7.56×10^{-1}
4	2.50×10^{-1}	2.50×10^{-1}	2.50×10^{-1}	2.50×10^{-1}	—
5	1.46×10^{-1}	1.46×10^{-1}	1.46×10^{-1}	2.24×10^{-1}	—
6	5.84×10^{-2}	5.84×10^{-2}	5.84×10^{-2}	—	—
7	1.62×10^{-2}	1.62×10^{-2}	1.97×10^{-2}	—	—
8	3.09×10^{-3}	3.09×10^{-3}	—	—	—
9	3.86×10^{-4}	4.16×10^{-4}	—	—	—
10	2.86×10^{-5}	—	—	—	—
11	9.54×10^{-7}	—	—	—	—
A	0.32	0.39	0.50	0.70	1.17
l	0	1.21×10^{-5}	1.58×10^{-3}	4.07×10^{-2}	3.20×10^{-1}

Table 1: Total bit rate distributions in the bufferless case.

behavior of a couple of bursty sources can be described as a Markov-Modulated Rate Process (MMRP) with σ states with total rates r_s and individual contributions $r_s^{(i)}$ as well as transition rates $\lambda_{s \to s'}$. We apply standard fluid flow spectral analysis [**?**, **?**] with a special treatment of the indifferent states [**?**]. We assume the following scenario:

- A bottleneck of C Mbps shared by constant and variable bit rate traffic.
- One constant bit rate source with $R^{(i)} \equiv 1$ Mbps.
- 10 variable bit rate sources of exponential on-/off type each with peak rate $r_{on} = 1$ Mbps and $\lambda_{\text{on}\to\text{off}} = 3\lambda_{\text{off}\to\text{on}}$, which implies a bit rate distribution of $\Pr\big\{R^{(i)} = 0 \text{ Mbps}\big\} = 0.75$ and $\Pr\big\{R^{(i)} = 1 \text{ Mbps}\big\} = 0.25$, respectively.

¿From now on, units are omitted.

4.2 Total Output Bit Rate Distribution

Bufferless case. Table 1 shows total output bit rate distributions for different bottleneck capacities C, load values A and loss ratios $l = 1 - \mathbf{E}[R^o] / \mathbf{E}[R]$ in the bufferless case: The value of $C = 11$ implies that there is enough capacity at all times, i.e. there is no bottleneck. The distribution of the bit rate at the output is given by the convolution of the individual bit rate distributions at the input.

The other cases address bottlenecks of different capacities $C < 11$. They have in common that the probabilities $\Pr\{R^o = r\}$ for $r < C$ match the ones from the convolution, while the rest of the probability mass is concentrated at $r = C$. The bottleneck with $C = 9$ hardly shows up. The loss ratio is very low, and the $\Pr\{R^o = 9\}$ is merely 8 % higher than in the no-bottleneck case. The bottleneck with $C = 7$ is more noticeable with a loss ratio around 2 % and an increase of $\Pr\{R^o = 7\}$ by 21 %. The bottleneck with $C = 5$ produces already 4 % loss, while $\Pr\{R^o = 5\}$ has increased by 53 % compared to the no-bottleneck

r	$\Pr\{R^o = r\}$			
	$C = 11$	$C = 9$	$C = 7$	$C = 5$
1	5.63×10^{-2}	5.63×10^{-2}	5.63×10^{-2}	5.62×10^{-2}
2	1.88×10^{-1}	1.88×10^{-1}	1.88×10^{-1}	1.82×10^{-1}
3	2.82×10^{-1}	2.82×10^{-1}	2.82×10^{-1}	2.62×10^{-1}
4	2.50×10^{-1}	2.50×10^{-1}	2.50×10^{-1}	2.07×10^{-1}
5	1.46×10^{-1}	1.46×10^{-1}	1.45×10^{-1}	2.93×10^{-1}
6	5.84×10^{-2}	5.84×10^{-2}	5.65×10^{-2}	—
7	1.62×10^{-2}	1.62×10^{-2}	2.25×10^{-2}	—
8	3.09×10^{-3}	3.07×10^{-3}	—	—
9	3.86×10^{-4}	4.38×10^{-4}	—	—
10	2.86×10^{-5}	—	—	—
11	9.54×10^{-7}	—	—	—
A	0.32	0.39	0.50	0.70
$G(0)$	0	8.37×10^{-5}	9.66×10^{-3}	2.13×10^{-1}

Table 2: Total bit rate distributions in the unlimited buffer case.

case. Finally, the totally overloaded bottleneck with $C = 3$ shows a very high spike at $\Pr\{R^o = 3\}$ that is 2.68 times as large as before.

Infinite buffer. Table 2 shows total output bit rate distributions for different bottleneck capacities C, load values A and probabilities of buffering $G(0) = \sum_{s \in \mathcal{S}^o} G_s(0)$. Due to reasons of stability, the case $C = 3$ has to skipped.

In case of $C = 9$, the buffer is hardly used, but its effect is already visible: $\Pr\{R^o = C\}$ is now 13 % higher than if there was no bottleneck, compared to 8 % in the bufferless case. This effect becomes even clearer for $C = 7$ (38 % instead of 21 %) and $C = 5$ (100 % instead of 53 %). Another difference to the bufferless case is seen in $\Pr\{R^o < C\}$ that is in general smaller than in the no-bottleneck case.

Consequently, one may assume that bottlenecks might be found from comparisons of the bit rate statistics of the total output stream with the convolution of the bit rate statistics, given that the bit rates of the input streams are uncorrelated. Buffered bottlenecks are likely to be visible since they show a larger peak close to the capacity. The bit rates smaller than the capacity appear with smaller probabilities than in the convolution case.

4.3 Individual Output Bit Rate Distributions

Table 3 and Figure 4 compare the individual bit rate statistics of the constant bit rate stream with those of *one* variable bit rate stream. They show input and output bit rate statistics of a bufferless bottleneck with $C = 5$ (i.e. a load of 70 %).

Behind the bottleneck, there are also lower speeds visible in the individual bit rate statistics. These stem from the overload states in which the capacity is shared between the competing streams. As it shares the resource with the variable bit rate streams in the same fair manner, the constant bit rate stream experiences the same kind of "slow-downs".

Table 4 and Figure 5 show the corresponding individual bit rate statistics for a system with infinite buffer. As in the bufferless case, we see "slow-downs" stemming from the overload states. As opposite to the bufferless case, there are also "speed-ups" exceeding the original (peak) bit rates at the input. These have their origins in situations of underload

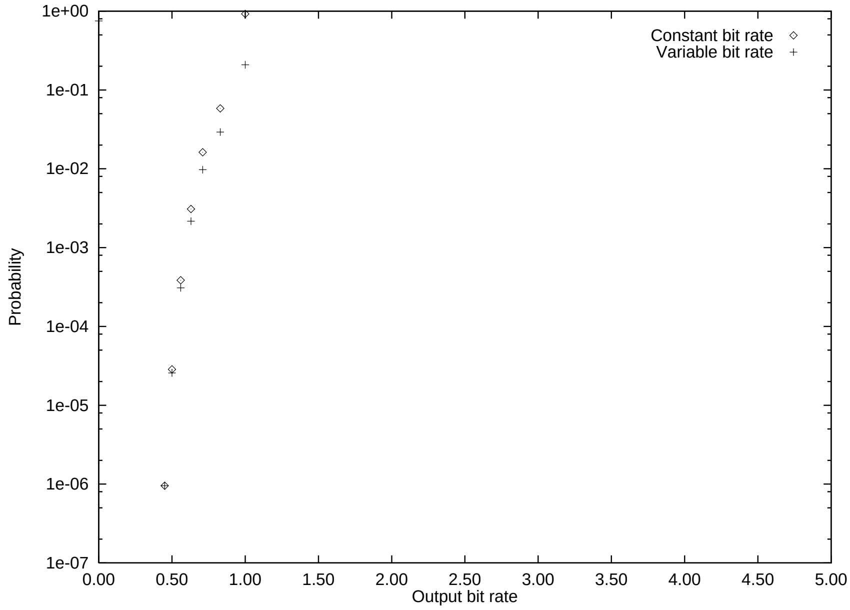

Figure 4: Rate distributions at output for individual streams, bufferless case.

r	Variable bit rate		Constant bit rate	
	$\Pr\{R^{(i)} = r\}$	$\Pr\{R^{(i),\mathrm{o}} = r\}$	$\Pr\{R^{(i)} = r\}$	$\Pr\{R^{(i),\mathrm{o}} = r\}$
0.00	7.50×10^{-1}	7.50×10^{-1}	0	0
0.45	0	9.54×10^{-7}	0	9.54×10^{-7}
0.50	0	2.57×10^{-5}	0	2.86×10^{-5}
0.56	0	3.09×10^{-4}	0	3.86×10^{-4}
0.63	0	2.16×10^{-3}	0	3.09×10^{-3}
0.71	0	9.73×10^{-2}	0	1.62×10^{-2}
0.83	0	2.92×10^{-2}	0	5.84×10^{-2}
1.00	2.50×10^{-1}	2.09×10^{-1}	1.00×10^{0}	9.22×10^{-1}

Table 3: Comparison of individual bit rate distribution of a constant bit rate stream with those of *one* variable bit rate stream, bufferless case.

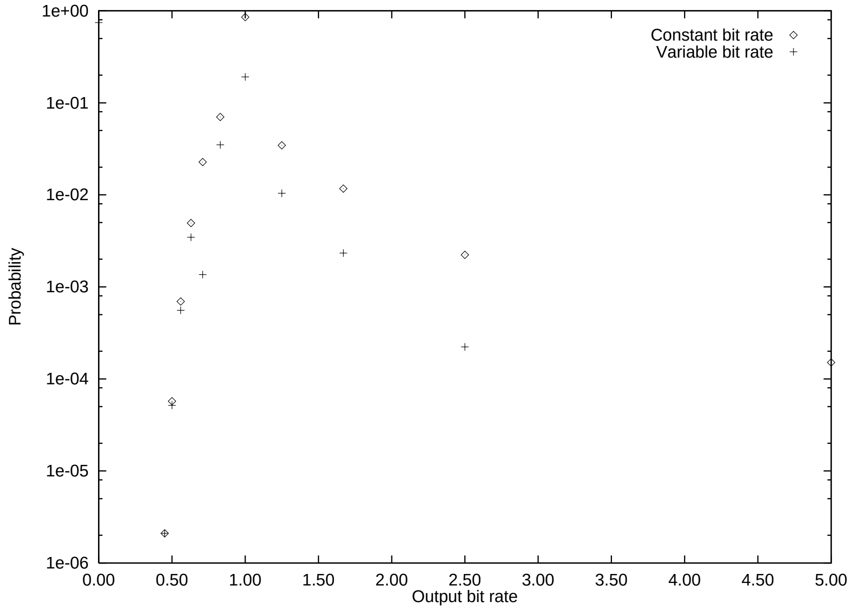

Figure 5: Rate distributions at output for individual streams, unlimited buffer case.

r	Variable bit rate		Constant bit rate	
	$\Pr\{R^{(i)} = r\}$	$\Pr\{R^{(i),\mathrm{o}} = r\}$	$\Pr\{R^{(i)} = r\}$	$\Pr\{R^{(i),\mathrm{o}} = r\}$
0.00	7.50×10^{-1}	7.43×10^{-1}	0	0
0.45	0	2.10×10^{-6}	0	2.10×10^{-6}
0.50	0	5.15×10^{-5}	0	5.72×10^{-5}
0.56	0	5.56×10^{-4}	0	6.95×10^{-4}
0.63	0	3.46×10^{-3}	0	4.94×10^{-3}
0.71	0	1.36×10^{-2}	0	2.27×10^{-2}
0.83	0	3.50×10^{-2}	0	7.01×10^{-2}
1.00	2.50×10^{-1}	1.91×10^{-1}	1.00×10^{0}	8.53×10^{-1}
1.25	0	1.04×10^{-2}	0	3.45×10^{-2}
1.67	0	2.33×10^{-3}	0	1.17×10^{-2}
2.50	0	2.23×10^{-4}	0	2.23×10^{-3}
5.00	0	0	0	1.51×10^{-4}

Table 4: Comparison of individual bit rate distribution of a constant bit rate stream with those of *one* variable bit rate stream, unlimited buffer case.

combined with a non-empty buffer. In this case, a stream may be served at a speed that even can reach bottleneck speed if it is the only stream in progress, cf. the last row of Table 4. Such "speed-ups" are typical for buffered bottlenecks. Again, the shapes of the distributions are the same for constant and variable bit rate traffic. In other word, the constant bit rate streams inherits the characteristics introduced by the variable bit rate streams.

Last but not least, the individual bit rate statistics can be used in an end-to-end manner to reveal the overall bottleneck performance of a network path between two peers.

5 Related Work

6 Conclusions and Open Issues

- Comparison of the total output statistics with the convolution of the individual input statistics may give hints on capacity and buffering capabilities. Typically used per link.
- Individual densities reveal the type of bottleneck (speed-ups = buffered) and may even reveal the capacity of the bottleneck. Can be used per link, but also end-to-end.
- Constant bit rate traffic seems to be best suited for bottleneck observations.

While the total bit rate statistics is basically shaped by the capacity of the bottleneck, the individual bit rate statistics at the output of the bottleneck reveal the interaction of traffic streams with different characteristics when passing the bottleneck. For instance, the bit rate pattern of streams with variable bit rate that are causing overload in the bottleneck is also visible in the pattern streams with constant bit rate sharing the same link. Moreover, the individual output bit rate distributions in particular contain clear indications about whether the bottleneck is equipped with buffering capabilities or not.

Further work will include simulations to validate our analytical results and to imitate real traffic conditions. Another important piece of future work will cover measurements of bit rate statistics and bottleneck identification and classification in real environments based on the qualitative results obtained in this paper.

W.r.t. fluid flow approach:

- Finite buffer size
- Simulation

W.r.t. "reality":

- Discretization effects
- Simulations
- Measurements