Traffic Equivalence and Substitution in a Multiplexer*

COSTAS COURCOUBETIS, ANTONIS DIMAKIS, AND GEORGE D. STAMOULIS

Institute of Computer Science (ICS)

Foundation for Research and Technology Hellas (FORTH)

and

Department of Computer Science

University of Crete

P.O. Box 1385 GR 711 10 Heraklion, Crete, Greece.

Email: {courcou,dimakis,gstamoul}@ics.forth.gr

Abstract

For a multiplexer fed with a large number of sources we derive conditions under which a single source can be substituted for a given subset of the sources while preserving the buffer overflow probability and the dominant time scales of buffer overflows. This equivalence is stronger than simple effective bandwidth equality and takes into account the context in which multiplexing takes place. This allows a substitution to be made for arbitrarily large proportions of the traffic without changing the operating point of the multiplexer as experienced by the rest of the traffic. It corresponds to defining a single source which is equivalent in a sense "local" to a given context, rather than equivalent in a sense which is "universal" to all contexts. The proposed methodology does not rely on traffic models and obtains the necessary information from the actual traffic traces. We study the case of fractional Brownian motion as a single source substitute and provide theoretical and experimental results that validate the approach.

1 Introduction

Modern broadband networks, with their large bandwidths, are able to carry simultaneously a very large number of connections, with diverse traffic characteristics. Statistical multiplexing is one of the factors

*This work was partly supported by the European Union under ACTS Project CA$hMAN (AC-039).

which accounts for the economies of scale that are foreseen in such large systems. Nevertheless, there are still many unresolved issues for the network operator who wishes to predict system performance and avoid congestion phenomena that lead to buffer overflows. One important reason is that simple traffic models traditionally used for performance analysis are not adequate to capture all essential aspects of the network traffic. For example, the encoding mechanisms used for multimedia traffic (e.g., MPEG) introduce important structure in the traffic stream, as well as burstiness at various discrete time scales and depending on the structure of the frame sequence. This structure is believed to greatly influence important performance parameters, such the buffer overflow probability, specially in the case of small buffers which is the case in networks carrying real time traffic.

A popular alternative to performance analysis is the simulation of reality. In this approach, traffic that looks as close to real as possible is generated and fed to the multiplexer, which is then monitored. The key obstacle in this approach is the generation of a realistic equivalent for the traffic resulting from a large number of sources (hundreds or thousands). An existing solution is the use of dedicated hardware that can be programmed to explicitly emulate the sources in real time. The popular approach is the use of Markov Modulated Poisson Process models (MMPPs) implemented as state machines. Besides the high cost of such dedicated hardware and doubts about the validity of MMPP models for multimedia traffic[1], state explosion soon becomes the bottleneck if multiple state models are used for improving accuracy (N sources, each represented by a k state model will in general result in k^N states).

In this paper we propose a new approach for creating a simple single source equivalent for a large number of such sources. This source can substitute for the given set of sources at the multiplexer without perturbing the dominant phenomena that cause buffer overflows in the particular multiplexing situation. Our methodology can use actual traffic measurements and source information given in terms of a number of representative actual traffic traces (no need of models).

The novelty in our approach lies in the fact that we construct a "local" model for a set of sources, instead of a "universal" one. It is by now well understood that not all aspects of a source contribute equally to the phenomena that cause the overflows at the multiplexer [3, 11]. Depending on the buffer size and the overall traffic mix, burstiness at different time scales becomes the major contributor to the overflows. Our motivation now becomes clear: we approximate the bursty behavior of the sources for the time scales that are only relevant. For doing that we use the Many Sources Asymptotic which provides us with the dominant

[1] due to heavy tails of state sojourn times instead of exponential, see [6]

term of the cell loss probability and abstracts the most important time scale in which source burstiness contributes to the overflows. This provides the formal framework for defining traffic equivalence: a single source can substitute for a set of sources if the substitution preserves some important properties of the above large deviation coefficient. We should mention that this equivalence is more subtle than simple effective bandwidth equality. This is because substituting for a significant part of the traffic by traffic with the same effective bandwidth as defined at the original operating point of the multiplexer, will in general perturb the above operating point. Hence, such a substitution can result in completely different dominant overflow phenomena and buffer overflow probabilities.

There are many advantages and prospective uses of the proposed methodology. We have already mentioned that the equivalent source can be constructed from information obtained from actual traces that span the different types of sources at the multiplexer, and from the information about the percentages of the different source types in the overall traffic mix. In this paper we have exploited the use of fractional Brownian motion (fBm), due to its analytical effective bandwidth representation and simple parameter fitting, for constructing the equivalent source. Using a single Gaussian process for traffic generation instead of hundreds of sources simplifies considerably the above task, and suggests that such "background" traffic can be generated in real time by software. This scenario is useful when one wants to conduct real experiments for assessing the performance of network control mechanisms, such as admission control, and resource allocation. An important issue was whether fBm can always be used to substitute for actual traffic. We have established by combining analytic properties of the large deviation coefficient with properties validated through experimentation that this is probably the case for MPEG traffic.

We conclude by mentioning a different application of the substitution approach, which can be used towards the fast calculation of the operating point parameters of the multiplexer, and hence for the on-line estimation of the buffer overflow probability and the effective bandwidth of the input traffic.

The paper is organized as follows. In Section 2 we provide the background information on the many sources asymptotic, in Section 3 we give the formal conditions for traffic equivalence, and we investigate theoretical issues including a detailed study of the use of fractional Brownian motion as the equivalent process; this reduces in the calculation of three parameters (the mean, standard deviation and Hurst parameter). In Section 4 we provide results from the experimental evaluation of the approach where we study substitution of a large number of MPEG sources in an ATM multiplexer. In Section 5 we discuss the use of the methodology for the fast calculation of the operating point of the multiplexer, and we conclude the paper with a discussion of the approach and with some directions for further research.

2 Background: the Many Sources Asymptotic

In this section we review some basic results from the theory of multiplexing. The model for the multiplexer is simple in order to illustrate the basic concepts; it can be extended to handle priorities as in [7].

The arrival process at a broadband link is the superposition of independent sources of J types. Let $N_j = N n_j$ be the number of sources of type j, and let $n = (n_1, \ldots, n_J)$ (note that the n_js are not necessarily integers). This system can be viewed as having N sources of the same type, where a single source consists of a proportion of the J source types and can be characterized by the vector n. The broadband link has a shared buffer of size $B = Nb$ and link capacity $C = Nc$. Parameter N is the scaling parameter (size of the system), and parameters b,c are the buffer and capacity per source, respectively.

Furthermore, let $X_j[0,t]$ be the total load produced by a source of type j in the time interval $(0,t]$, which feeds the above link. We assume that $X_j[0,t]$ has stationary increments. The *effective bandwidth* of a source of type j is defined as [7]

$$\alpha_j(s,t) = \frac{1}{st} \log \mathrm{E}\left[e^{sX_j[0,t]}\right] , \tag{1}$$

where s,t are system parameters which are defined by the context of the source, i.e., the characteristics of the multiplexed traffic, their QoS requirements, and the link resources (capacity and buffer). Specifically the *time* parameter t (measured in, e.g., milliseconds) corresponds to the most probable duration of the busy period of the buffer prior to overflow [4] (i.e., the time-to-fill the buffer) . The *space* parameter s (measured in, e.g., kb^{-1}) corresponds to the degree of multiplexing and depends, among others, on the size of the peak rate of the multiplexed sources relative to the link capacity. Effective bandwidths are increasing in s [7]. In particular, for link capacities much larger than the peak rate of the multiplexed sources, s tends to zero and $\alpha_j(s,t)$ approaches the mean rate of the source, while for link capacities not much larger than the peak rate of the sources, s is large as $\alpha_j(s,t)$ approaches the maximum value that the random variable $\frac{X_j[0,t]}{t}$ can attain.

Let $Q(Nc, Nb, Nn) = \mathrm{P}(\text{overflow})$ be the probability that in an infinite buffer which multiplexes $Nn = (Nn_1, \ldots, Nn_j)$ sources and is served at rate $C = Nc$, the queue length is above the threshold

$B = Nb$. The following holds for $Q(Nc, Nb, Nn)$:

$$\lim_{N\to\infty} \frac{1}{N} \log Q(Nc, Nb, Nn) = -\inf_t \sup_s \left[s(ct+b) - st \sum_{j=1}^{J} n_j \alpha_j(s,t) \right] = -I \; , \tag{2}$$

where I is called the asymptotic *rate function*. The last equation is referred to as the *many sources asymptotic* (or infsup formula) and has been proved for discrete time in [4], for continuous time in [2] and for a special case in [12]. Due to equation (2), the overflow probability can be written as P(overflow) $= e^{-NI+o(N)}$, which leads to the following approximation when N is large:

$$\log \mathrm{P(overflow)} \approx -\inf_t \sup_s \left[s(B+Ct) - st\alpha_R(s,t) \right] = -\Phi_R \; , \tag{3}$$

where $\alpha_R(s,t)$ is the effective bandwidth function of the aggregate input traffic (since the input sources are independent, this is the sum of the individual effective bandwidths). Φ_R can be thought as the scaled asymptotic rate function.

If a QoS requirement of the type P(overflow) $\leq e^{-\gamma}$ is imposed then the acceptance region (the points n where the above requirement is satisfied) for large N is approximated by an affine in N boundary [3]

$$\sum_{j=1}^{J} N_j \alpha_j(s^*, t^*) \leq C + \frac{1}{t^*}\left(B - \frac{\gamma}{s^*}\right) = C^* \; , \tag{4}$$

where (s^*, t^*) is an extremizing pair in equation (2), and C^* is the "effective capacity" of the system at the above operating point (s^*, t^*). Clearly, this point constitute a *saddle* point of the mapping $(s,t) \mapsto [s(B+Ct) - st\alpha_R(s,t)]$. That is, the derivatives of this function with respect to s and t both vanish at (s^*, t^*), while the corresponding Hessian matrix is non-definite (i.e. neither positive semi-definite nor negative semi-definite). The second order condition is equivalent to the determinant of the Hessian matrix being negative, because this is a 2×2 matrix. Henceforth, when referring to saddle points, it will be taken that these correspond to $\inf_t \sup_s$ as opposed to $\sup_t \inf_s$, unless otherwise specified.

The existence of effective bandwidths and the validity of the above asymptotic results assume only stationarity of sources. Illustrative examples discussed in [7] and [11] include periodic sources, fractional Brownian input, Markovian sources, policed and shaped sources. An important feature of equation (3) is that it can be solved based on trace information of the actual traffic sources. Given a trace of a source of type j, one can construct off-line the log-moment generating function on a grid over the (s,t) space; having done this for each j, one can then explicitly solve the optimization problem. A way to considerably speed-up the above procedure is described in Section 5.

The effective bandwidth $\alpha_j(s,t)$ provides a relative measure of resource usage *for a particular operating point* of the link, expressed through parameters s, t. For example, if a source of type j_1 has twice as much effective bandwidth as a source of type j_2, then, for this particular operating point of the link, one source of the first type can be replaced by two sources of the second type, while satisfying the QoS constraint. Thus effective bandwidths can be used as a basis for substitution between sources at a particular operating point. Unfortunately this argument fails if one wants to substitute a larger fraction of the total traffic, since this will in general lead to a different operating point (s,t). Indeed, it is logical to assume that the operating point (s,t) is dictated by the interaction of the overall traffic mix, and substituting a single source will have a negligible effect. On the other hand, changing the traffic composition can introduce new dominant time scales of buffer overflows, see [3].

3 Traffic Substitution

The key for providing the right equivalence condition with the appropriate substitution properties is in equation (2). The conditions we propose ensure that the term inside the brackets remains approximately unchanged in a neighborhood around the original operating point. This, unless a new local optimum is introduced with a smaller value, will guarantee that the buffer overflow probability will be the same and overflows will occur with the same mode (with respect to the dominant overflow time scales). One can easily extend the approach to include conditions on the first two moments of the processes in order to match other properties of interest such as the average queue length; this extension is outside the scope of this paper.

3.1 The Equivalence Condition

In this subsection, we define the condition to be satisfied by the modeled traffic in order to substitute for the actual one, and then we prove some important properties of the substitution.

We will first define some notation, which will be used throughout the rest of this paper. The multiplexer consists of buffer of size B and has service rate C. It is originally fed by traffic $R+V$, where R is the traffic to be substituted for, and V is the rest of the traffic. After substitution it is fed by traffic $M+V$, where is M is the "model" of the real traffic R. If the complete traffic is substituted, then by abusing notation $V=0$.

We define the function $G_X(s,t) = s(B+Ct) - st\alpha_X(s,t)$, where $\alpha_X(s,t)$ is the effective bandwidth of traffic $X \in \{R, V, M, R+V, M+V\}$, and use (s^*, t^*) to denote the operating point of the link under the actual traffic $R+V$, i.e., the point that solves $\inf_t \sup_s G_{R+V}(s,t)$. We will say that the function $\alpha(s,t)$ is differentiable at a given point if both its partial derivatives exist and are finite.

The next definition defines traffic equivalence at some arbitrary point (s_0, t_0).

Definition 1 *Let M and R be traffic with effective bandwidth functions $\alpha_M(s,t)$ and $\alpha_R(s,t)$ respectively that are differentiable at (s_0, t_0). Traffic M is said to be equivalent with traffic R at some point (s_0, t_0) iff*

$$\begin{cases} \alpha_R(s_0,t_0) = \alpha_M(s_0,t_0) \\ \\ \frac{\partial \alpha_R}{\partial s}(s,t_0)\Big|_{s=s_0} = \frac{\partial \alpha_M}{\partial s}(s,t_0)\Big|_{s=s_0} \\ \\ \frac{\partial \alpha_R}{\partial t}(s_0,t)\Big|_{t=t_0} = \frac{\partial \alpha_M}{\partial t}(s_0,t)\Big|_{t=t_0} \end{cases}. \tag{5}$$

It is easy to see that the above conditions define a reflective, symmetric, and transitive relation, i.e., an equivalence relation defined locally with respect to the point (s_0, t_0). An important observation is that in order to simultaneously satisfy the above conditions between a real traffic R and the traffic generated by a model M, any such candidate model must have at least three free parameters whose value will be defined by the above system of three equations.

The important property of traffic equivalence is that when defined on a saddle point it preserves the above saddle point after substitution. The problem is that substitution might introduce more saddle points which might in turn produce a new operating point. The above property is formally stated as follows.

Theorem 1 *Assume that $G_{R+V}(s,t)$ has a saddle point at (s_0,t_0), in a neighborhood D of which $\alpha_R, \alpha_M, \alpha_V$ are twice differentiable. If at (s_0,t_0) M is equivalent with R, and G_{M+V} has a non-definite Hessian, then G_{M+V} has also a saddle point at (s_0,t_0), and furthermore $G_{M+V}(s_0,t_0) = G_{R+V}(s_0,t_0)$.*

Proof: The last two conditions in (5) imply the equality of the partial derivatives of G_{R+V} and G_{M+V} at (s_0,t_0), and since (s_0,t_0) is a saddle point for G_{R+V}, these derivatives must be zero. Now, since the Hessian of G_{M+V} is non-definite it follows that the above point is also a saddle point for G_{M+V}. It only remains to show that this corresponds to $\inf_t \sup_s$ as opposed to $\sup_t \inf_s$. But this follows from the fact that $G_{M+V}(s,t)$ is concave w.r.t. s for every $t > 0$, due to convexity of moment generating functions [1].

Indeed, by concavity of $G_{M+V}(s,t)$ w.r.t s, the extreme w.r.t s can only be a local $\sup_s$. Moreover by the first condition in (5), it follows that $G_{M+V}(s_0,t_0) = G_{R+V}(s_0,t_0)$. ∎

The main consequence of Theorem 1 is as follows: If traffic $R+V$ feeds the link and has (s^*,t^*) as the operating point, then, if M substitutes for R, (s^*,t^*) will still be a saddle point of $G_{M+V}(s,t)$ and thus a *potential* operating point; it will not be the new operating point if more saddle points are introduced that achieve a smaller value for G_{M+V}.

A physical interpretation is that substitution of R by the model M may introduce new most probable time scale(s) of overflow. Thus eventually the two systems would have different overflow probabilities occurring under different modes of buffer overflow.

Summarizing the above, an essential requirement for our approach of traffic substitution to be effective is that traffic model M *preserves* the operating point at which it substituted for the actual traffic R. If traffic M satisfies this requirement then it preserves both the type and the frequency of the dominant phenomena that determine overflows. Below we show that the above always holds in the case of substitution with fractional Brownian motion, in the absence of background traffic V.

Regarding the assumption of a non-definite Hessian matrix, since this involves real traffic, it can be checked only through experiments (it can be validated analytically in special cases where all the traffic is generated by models). We have conducted extensive experiments involving (see also Section 4) combinations of modeled (M) and MPEG traffic (V); in all cases, the Hessian matrix was indeed non-definite.

Deriving substitution conditions from traffic traces: an example

As already explained in the Introduction, network designers would like to be able to assess performance under real load scenarios. Using our method, this can be done by substituting large traffic aggregations by an equivalent source having a simple model (with at least three free parameters as we discussed previously). In order to do this, one should evaluate numerically for the actual system the value of the operating point (s^*,t^*) and the values of the LHSs of the equations (5). We illustrate in terms of an example the necessary steps involved.

Suppose that we want to assess the QoS of a particular video (say James Bond) in a link with high degree of multiplexing, by using our method for substituting the background traffic. To make things concrete, consider a link with capacity 155 Mbps and buffer size of 2000 cells which is fed by the aggregation of 40 independent James Bond movies and 400 independent Star Wars movies (both encoded in MPEG), where we are interested to observe visually the performance of the James Bond movies only. In this case we need to substitute all Star Wars movies by a single modeled source. To apply the proposed approach, we first calculate the effective bandwidth functions $\alpha_{SW}(s,t)$, $\alpha_{JB}(s,t)$ of Star Wars and James Bond movies respectively using only a single trace for each[2], and then calculate the operating point of the system. The latter is computed by solving numerically $\inf_t \sup_s \left[s(C+Bt) - st(40\alpha_{JB}(s,t) + 400\alpha_{SW}(s,t))\right]$ involving the effective bandwidth function of both movies. It should also be noted that if we add say one (or only a few) more video (e.g. Simpsons cartoon), then the new operating point will be a small perturbation of that already calculated, and for most practical purposes can be considered to be the same. This remark can be particularly helpful in cases where the application traffic is small relative to the overall traffic mix and is not known prior to substitution; e.g., a videoconferencing application.

3.2 Substitution with Gaussian Sources

Gaussian sources are a good candidate for substitution due to their simple analytic form of the effective bandwidth, and the rather straightforward way to implement traffic generation.

Let $X[0,t]$ denote the amount of work generated by a source in the interval $[0,t]$. Suppose that this is given by

$$X[0,t] = \lambda t + Z(t)\,, \tag{6}$$

where $Z(t)$ is a zero mean Normally distributed random variable. The effective bandwidth is then [7]

$$\alpha_M(s,t) = \lambda + \frac{s}{2t}\mathrm{Var}Z(t)\,. \tag{7}$$

[2]the functions are encoded by their values as estimated from the trace, on a finite grid over the $s-t$ positive orthant

The system of equations (5) in this case evaluated at (s_0, t_0) leads to

$$\begin{cases} \lambda = \alpha_R(s_0,t_0) - s_0 \, \frac{\partial \alpha_R}{\partial s}(s,t_0)\Big|_{s=s_0} \\ \\ \mathrm{Var}Z(t_0) = 2t_0 \, \frac{\partial \alpha_R}{\partial s}(s,t_0)\Big|_{s=s_0} \\ \\ \frac{s_0}{2}\frac{d}{dt}\left(\frac{\mathrm{Var}Z(t)}{t}\right)\Big|_{t=t_0} = \frac{\partial \alpha_R}{\partial t}(s_0,t)\Big|_{t=t_0} \end{cases} . \tag{8}$$

Therefore, for substitution of a given source by a Gaussian one, the only information on the function Var$Z(t)$ that is necessary is the value of the function and of its derivative at $t = t_0$.

3.3 Heavy Multiplexing

When the system has a large degree of multiplexing, the space parameter s of the system approaches zero. Large links with transmission rate 622 Mbps or 1.2 Gbps are expected to offer such degree of multiplexing. If $X[0,t]$ denotes the amount of work generated by the source R in the interval $[0,t]$, then for small s the effective bandwidth is given by [7]

$$\alpha_R(s,t) = \frac{\mathrm{E}X[0,t]}{t} + s\frac{\mathrm{Var}X[0,t]}{2t} + \mathrm{o}(s)\,; \tag{9}$$

that is, we only omit terms of second, or higher, order. If a Gaussian source G substitutes for a source R at the operating point (s^*, t^*) with s^* near zero, then system of equations (5) gives

$$\begin{cases} \lambda = \frac{\mathrm{E}X[0,t^*]}{t^*} + \mathrm{o}(s^*) \\ \\ \mathrm{Var}Z(t^*) = \mathrm{Var}X[0,t^*] + \mathrm{o}(s^*) \\ \\ \frac{d\mathrm{Var}X[0,t]}{dt}\Big|_{t=t^*} = \frac{d\mathrm{Var}Z(t)}{dt}\Big|_{t=t^*} + \mathrm{o}(s^*) \end{cases} . \tag{10}$$

So the parameters of the model depend on the first two moments of the input process at the time scale where overflows occur. In large systems, modeling should be performed using only the first and second order characteristics of the input process (diffusion approximation). This shows that our method agrees with the established ones for the cases of heavy multiplexing (see [9, 5, 6]).

3.4 The Case of $B = 0$

We consider bufferless links where we substitute real traffic using Gaussian sources. The overflow probability of the link depends only on steady-state characteristics of the input process. The time parameter of bufferless systems is always (independently of the input traffic) near 0. So we need to model the behavior of the sources at an operating point $(s^*, 0)$. The third equation of the system (5) should not be considered because in the case of no buffer, $t^* = 0$. Thus every model that satisfies the two first equations of (5) will have eventually the same asymptotic decay of overflow probability as given by the many sources asymptotic. Consequently two-parameter models suffice for substitution at the operating point $(s^*, 0)$. Brownian motion is such a model with $\mathrm{Var}Z(t) = \sigma^2 t$ and it commonly arises from heavy traffic models. Using, (7), the effective bandwidth is then

$$\alpha_M(s,t) = \lambda + \frac{s\sigma^2}{2}\,. \tag{11}$$

Since Brownian motion is a Gaussian source, the first two equations of (8) lead to

$$\lambda = \lim_{t\to 0}\left[\alpha_R(s^*,t) - s^*\,\frac{\partial \alpha_R}{\partial s}(s,t)\bigg|_{s=s^*}\right] \tag{12}$$

and

$$\sigma^2 = 2\lim_{t\to 0}\frac{\partial \alpha_R}{\partial s}(s,t)\bigg|_{s=s^*}\,. \tag{13}$$

The accuracy of substitution is depicted in Figure 1. Actual traffic is superposition of independent Star Wars MPEG-1 streams. The MPEG-1 compressed video sequence, made available[3] by O. Rose [10], has the following frame structure: IBBPBBPBBPBB (12 frames) and rate 25 fps (frames per second). The total duration is approximately 30 minutes and the mean rate 0.26 Mbps. The constituent traces are independent in the sense that the MPEG-1 streams are superimposed with each one starting at a random time instant which is independent of the starting time of the other streams. The accuracy is measured against that of a modeling approach using Brownian motion that matches only the first two moments of the marginal distribution of the input traffic. It is evident that for overflow probabilities smaller that 10^{-3} (which is the area of interest) our substitution approach is considerably better than the above approach, which is based

[3] Available at http://ftp-info3.informatik.uni-wuerzburg.de/pub/MPEG/

on the Central Limit Theorem approximation. This corroborates the fact that Chernoff bounds are more accurate that the Central Limit Theorem for tail probabilities.

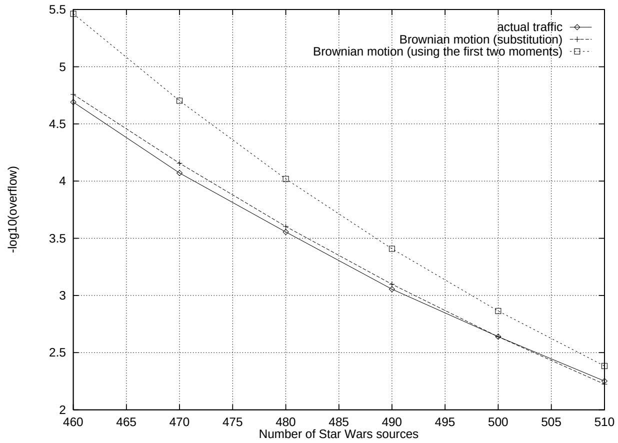

Figure 1: Overflow probability for a 155Mbps link with $B = 0$ as the number of superimposed Star Wars MPEG-1 encoded streams is increased.

3.5 Fractional Brownian Motion

Fractional Brownian motion (fBm) is a Gaussian source with Var$Z(t) = \sigma^2 t^{2H}$, where σ and $H \in (0, 1)$ are constants. Using also (7), the effective bandwidth of fBm is

$$\alpha_M(s, t) = \lambda + \frac{s}{2}\sigma^2 t^{2H-1} \ . \tag{14}$$

H is the Hurst parameter and depending on whether $H > \frac{1}{2}$, $H \leq \frac{1}{2}$ the source is long- or short-range dependent respectively. Long-range dependent processes have non-summable autocorellation function while short range dependent sources have summable ones. A remarkable feature of fBm is that it is completely characterized by three parameters, that is, as many parameters as equations in (5).

Solving the system (8) in this case we have

$$\begin{cases} \lambda = \alpha_R(s_0,t_0) - s_0 \, \frac{\partial \alpha_R}{\partial s}(s,t_0) \mid_{s=s_0} \\ \\ 2H-1 = \frac{t_0 \, \frac{\partial \alpha_R}{\partial t}(s_0,t)\Big|_{t=t_0}}{s_0 \, \frac{\partial \alpha_R}{\partial s}(s,t_0)\Big|_{s=s_0}} \\ \\ \sigma^2 = \frac{2 \, \frac{\partial \alpha_R}{\partial s}(s,t_0)\Big|_{s=s_0}}{t_0{}^{2H-1}} \end{cases} . \tag{15}$$

The fact that the number of parameters equals the number of equations makes one wonder whether there is enough freedom to guarantee that the fBm with parameters derived from the above equations is always a *valid* one, i.e. that $\lambda \in (0, C)$, $H \in (0, 1)$ and $\sigma^2 > 0$. Note that the condition on σ^2 is always guaranteed, since effective bandwidths are increasing in s [7]. On the other hand, there is no algebraic proof that λ, as given by the first of the above equations, satisfies $\lambda \in (0, C)$ for an arbitrary point (s_0, t_0). However, there is important evidence that this holds if (s_0, t_0) is the operating point (s^*, t^*), as discussed at the end of this subsection. Henceforth, we focus on studying equivalence of R and an fBm at the operating point, and we assume that the condition guaranteeing $\lambda \in (0, C)$ does hold, namely that

$$0 < \alpha_R(s^*,t^*) - s^* \, \frac{\partial \alpha_R}{\partial s}(s,t^*) \mid_{s=s^*} < C \, . \tag{16}$$

The following lemma establishes a condition equivalent to the righthand inequality of (16) expressed in terms of the scaled asymptotic rate function Φ_R defined in (3); this condition will be used in the proof of the main result of this subsection given by Theorem 2.

Lemma 1 *The condition* $\alpha_R(s^*,t^*) - s^* \, \frac{\partial \alpha_R}{\partial s}(s,t^*) \mid_{s=s^*} < C$ *is equivalent to*

$$2\frac{\Phi_R}{B} > \frac{\partial \Phi_R}{\partial B} \, . \tag{17}$$

Proof: Since (s^*, t^*) is a saddle point of $G_R(s^*, t^*)$, we have

$$\frac{\partial G_R(s,t)}{\partial s}|_{s=s^*} = B + Ct^* - t^*\alpha_R(s^*,t^*) - s^*t^*\frac{\partial \alpha_R(s^*,t^*)}{\partial s}|_{s=s^*} = 0 \, ,$$

which together with (3) implies that

$$s^* \frac{\partial \alpha_R(s^*, t^*)}{\partial s}|_{s=s^*} = C - \alpha_R(s^*, t^*) + \frac{B}{t^*} = \frac{\Phi_R}{s^* t^*} . \tag{18}$$

This now implies that the inequality $\alpha_R(s^*, t^*) - s^* \frac{\partial \alpha_R}{\partial s}(s, t^*) |_{s=s^*} < C$ is equivalent to $\frac{\Phi_R}{s^* t^*} > \alpha(s^*, t*) - C$. Using (3) again, this inequality is equivalent to $\frac{\Phi_R}{s^* t^*} > \frac{B}{t^*} - \frac{\Phi_R}{s^* t^*}$ and hence to $2\frac{\Phi_R}{B} > s^*$. The lemma follows from the property that $s^* = \frac{\partial \Phi_R}{\partial B}$, which can be established by taking derivatives of Φ_R w.r.t. B and using the fact that (s^*, t^*) is a saddle point. ■

The following theorem establishes that if condition (16) applies to a case of traffic equivalence with $V = 0$, then H indeed belongs to $(0, 1)$, and thus the fBm model derived is valid. Moreover, the theorem states that the operating point is *always preserved* after substitution.

Theorem 2 *Consider the case of traffic R with operating point (s^*, t^*). If (16) holds, then there exists a valid fBm M that is equivalent to R at (s^*, t^*). Furthermore, if M substitutes for M, then the same operating point is preserved.*

Proof: First it can be shown similarly to (18) that

$$t^* \frac{\partial \alpha_R(s^*, t^*)}{\partial t}|_{t=t^*} = C - \alpha_R(s^*, t^*) .$$

Combining this with (18) and the second equation of system (15) we obtain

$$2H - 1 = \frac{C - \alpha_R(s^*, t^*)}{C - \alpha_R(s^*, t^*) + \frac{B}{t^*}} = 1 - \frac{\frac{B}{t^*}}{C - \alpha_R(s^*, t^*) + \frac{B}{t^*}} = 1 - \frac{B}{\Phi_R / s^*} , \tag{19}$$

where we have also used the definition of Φ_R in (3). Since Φ_R is positive it follows that $H < 1$. Furthermore, it is straightforward to check that the condition $H > 0$ is equivalent to $2\frac{\Phi}{B} > s^*$ which was established in Lemma 1.

Next, we show that the operating point (s^*, t^*) is preserved after substitution of R by M. In particular, we show that $G_M(s, t)$ has a unique saddle point at (s^*, t^*). For this purpose, we explicitly compute the saddle point of $G_M(s, t)$. First note that by (14), we have $G_M(s, t) = s(B + (C - \lambda)t) - \frac{s^2}{2} t^{2H}$. It is straightforward to see that, for any fixed $t > 0$, $\sup_s G_M(s,t)$ is attained at $\hat{s}_t = \frac{B+(C-\lambda)t}{\sigma^2 t^{2H}}$. Furthermore,

$$\inf_t \sup_s G_M(s,t) = \inf_t \frac{(B+(C-\lambda)t)^2}{2\sigma^2 t^{2H}} \, .$$

For $t > 0$, the derivative of the function $\frac{(B+(C-\lambda)t)^2}{2\sigma^2 t^{2H}}$ w.r.t. to t equals

$$\frac{(B+(C-\lambda)t)(C-\lambda)(1-H)}{\sigma^2 t^{2H+1}}(t-\hat{t}) \, , \tag{20}$$

where

$$\hat{t} = \frac{HB}{(1-H)(C-\lambda)} \, , \tag{21}$$

which is positive due to assumption (16) and its consequence that $H \in (0,1)$. Since all terms multiplying $(t-\hat{t})$ are positive, and the above derivative only vanishes at $\hat{t}$, it follows that $\inf_t \sup_s G_M(s,t)$ is attained at $\hat{t}$, while $\sup_s G_M(s,t)$ has no other local optima w.r.t. t. Thus, $G_M(s,t)$ has a unique saddle point at $(\hat{s}_{\hat{t}}, \hat{t})$.Since, by construction, the derivatives of $G_M(s,t)$ at (s^*, t^*) are equal to zero, it must be the case that $(\hat{s}_{\hat{t}}, \hat{t}) \equiv (s^*, t^*)$. ■

Regarding the validity of (17), there are interesting indications that in many cases a stronger condition holds, namely

$$\frac{\Phi_R}{B} > \frac{\partial \Phi_R}{\partial B} \, . \tag{22}$$

This is equivalent with the concavity of the scaled asymptotic rate function Φ_R w.r.t. B. Indeed, for the superposition of ON/OFF Markov fluids and small B it was shown in [13] that $\Phi_R \approx C_1 + C_2\sqrt{B}$. Also, for large B, it is shown in [4] that Φ_R converges to a linear function in B; similar indications can be found in [5]. On the other hand, concavity is not always the case, as for the sub-bursty Markovian sources considered in [2].

Since we could not hope for a general result validating (17) and (22), we had to resort to experimentation. Thus, we have done extensive experiments to validate (17) and (22) for actual traffic of interest (MPEG). In all cases, both these conditions were verified. See Figure 2 where Φ_R is plotted as a function of the buffer size B for several cases where many independent yet identical video movies feed a link with capacity 155 Mbps; in all experiments, the utilization of the link was 88%.

An interesting observation is that if (22) holds, then (19) implies that $H > 1/2$, and the corresponding fBm is long-range dependent. It is somewhat surprising that in certain cases a long-range dependent fBm can be equivalent to a short-range dependent traffic.

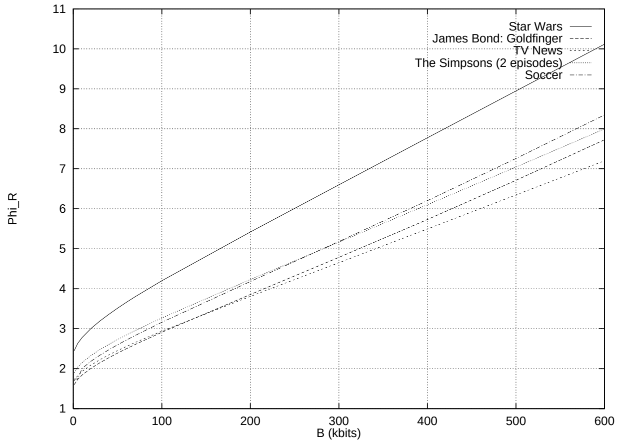

Figure 2: Φ as a function of the buffer size B for various MPEG-1 encoded streams. In each case many yet identical independent video movies feed a link with capacity 155 Mbps; in all experiments, the utilization of the link was 88%.

4 Experimental Results

In the present section, we validate our approach by means of experiments. Three types of experiments have been conducted:

1. Study of the accuracy of substitution, by comparing buffer overflow probabilities (BOP) of the actual traffic against the BOP achieved by the substituting traffic. Such experiments have been carried out for complete substitution of traffic, as well as for substitution of a subset of the actual traffic. To further validate our approach, we have also compared the actual BOP to that attained by a source model derived by statistical parameter matching in the case of complete substitution.

2. Validation that indeed preserving t^* results in preserving the dominant timescales at which overflows occur. This has to do with the accuracy and the interpretation of the many sources asymptotic.

3. Comparison of the actual BOP to that attained after complete substitution according to our approach as well as to the BOP attained by a naive approach matching only effective bandwidths of the actual traffic and the traffic model, without taking into account the first order conditions. This comparison motivates the need for taking such conditions into account.

In all types of experiments, the overflow probability for every input traffic is measured by brute force simulation.

In the first set of experiments, for performance assessment, the real traffic we consider is a superposition of independent MPEG-1 encoded streams of the Star Wars movie. The streams are multiplexed into a 34 Mbps link with 500 cells buffer (about 6 msec maximum delay) and into a 155 Mbps link with 2000 cells buffer (about 5.5 msec maximum delay). We examine the accuracy of two models used for substitution: a fBm source and a superposition of i.i.d. ON/OFF Markov Modulated Fluids. For the latter, the unknown parameters determined by the equations (5) is the rate ν of transitions from state OFF to ON, the rate μ of transitions from state ON to OFF, the rate h of traffic produced at state ON, and the number K of superimposed ON/OFF sources. The effective bandwidth for such a source is known analytically [7]. In order to fully determine the parameters (ν, μ, h and K) by solving the system (5) one more relation is needed: we have set that the mean rate of MMF equal the actual source mean rate.

To demonstrate the effectiveness of the substitution method against straightforward modeling approaches, we calculate the BOP of a system fed by a superposition of i.i.d. ON/OFF MMF constructed according to a statistical matching method in which we match the mean rate, peak rate, rate variance and the average time at which the rate is below a threshold corresponding to 10% of the mean. (We established that the 10% threshold led to better results than other threshold levels.) Figure 3 depicts the accuracy of the aforementioned models; note that "stat.MMF" stands for the above straightforward method. The figure implies that the BOP attained by traffic substitution is very close to the actual one, for both MMF and fBm models, while stat.MMF method is considerably inaccurate, particularly when employed for the 155 Mbps system.

Figure 5 depicts BOP comparison for substitution of subsets of 100 Star Wars sources, multiplexed into a 34 Mbps link with 500 cells buffer. There is excellent agreement of the BOP attained by substitution with fBm to the actual value for 25%, 50%, 75% and 100% of sources substituted; the actual value corresponds to 0% of sources substituted.

In the second set of experiments, we measure the time-to-fill the buffer t_f for the actual and the substitution traffic, which in both cases should be close to the value of t^*. This indeed applies to the cases studied, as depicted in Figure 4, in which fBm and MMF substitute for aggregations of Star Wars sources that are again multiplexed at a 34 Mbps link with 500 cells buffer; t^* is given by the curve marked as "theoretical". Agreement is very good, particularly for the fBm model, thus confirming that the dominant timescales for overflow are indeed preserved.

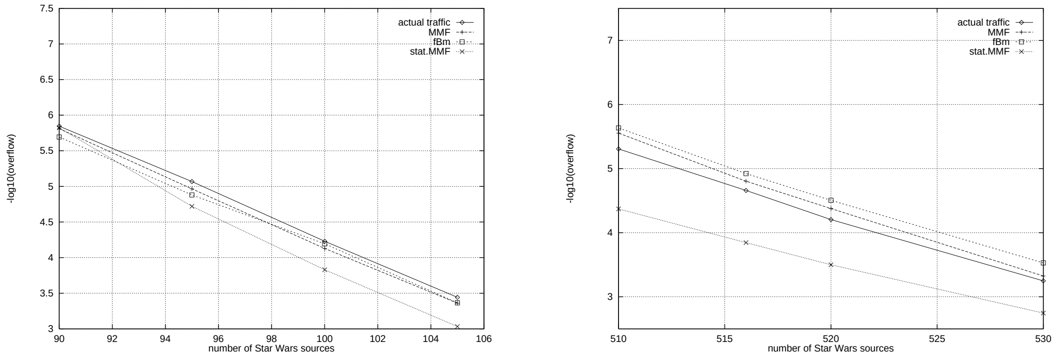

(a) (b)

Figure 3: BOP for various models against the actual BOP for a 34 Mbps link with 500 cells buffer (a) and for a 155 Mbps link with 2000 cells buffer (b) as the number of superimposed Star Wars MPEG-1 encoded streams is increased

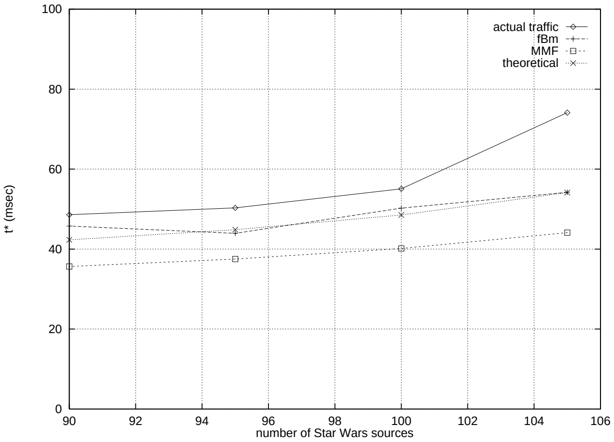

Figure 4: The time-to-fill the buffer is depicted for the actual traffic, fBm and MMF models used to substitute for the former at a 34 Mbps link with 500 cells buffer.

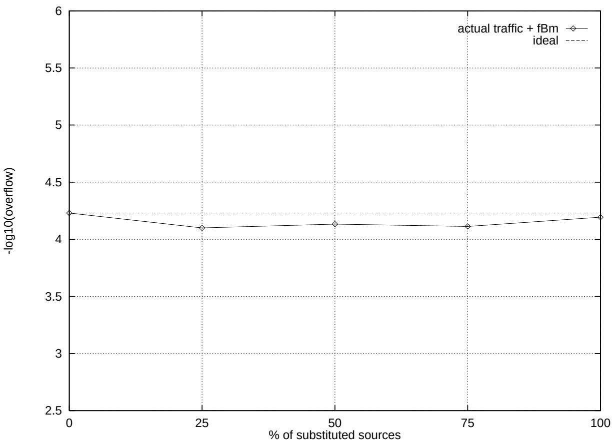

Figure 5: BOP achieved by substituting subsets of the input sources by fBm. The BOP value for the actual traffic corresponds to 0% of traffic substituted

Finally, to further motivate our approach, we study in the third set of experiments the performance attained when traffic substitution with MMF is solely based on equality of effective bandwidths, without taking into account the first order conditions of traffic substitution. Again it is assumed that multiple Star Wars sources are multiplexed in a 34 Mbps link with buffer capacity of 500 cells. Equality of effective bandwidths can be attained by different combinations of values of the parameters involved. As depicted in Figure 6, this can lead to considerable mismatch of the performance for certain such combinations. On the other hand, substitution according to our approach matches the actual BOP almost perfectly. This comparison clearly demonstrates the importance of the first order conditions.

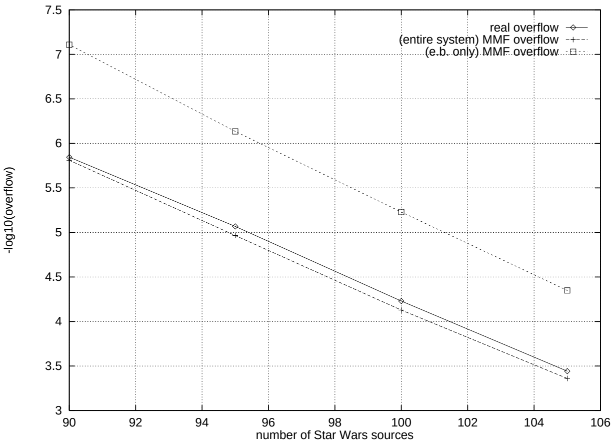

Figure 6: BOP comparison when two MMF models substitute for the actual traffic. The first model substitutes for the traffic according to our notion of equivalence, while the other model is based only on equality of the effective bandwidths and ignores the first order conditions.

5 Fast Calculation of Operating Point

The application of our substitution approach relies on knowledge of the operating point (s^*, t^*). We have already mentioned that this can be computed off-line by solving the infsup formula in equation (2). This solution is a saddle point of the $G_R(s,t)$. There is no general method for solving such problems especially when $\alpha_R(s,t)$ is computed from real traffic traces. A common way to solve this problem is by taking the maximum w.r.t. s for each t, and then taking the minimum over t. This method is computationally demanding, because two optimization problems need to be solved sequentially[4]. Having in mind that a single evaluation (on a particular point (s,t)) of the effective bandwidth function for a real trace may take

[4]We can exploit the concavity of $s(B+Ct) - st\alpha_R(s,t)$ in terms of the parameter s in order to speed-up the solutions of the $\sup_s$ problems.

seconds, this method becomes prohibitively expensive to be used for the on-line evaluation of the above operating point.

We have devised an iterative method for calculating the operating point using the idea of traffic substitution. The key idea is to construct a mapping having the value of the operating point as a fixed point. The method starts with an initial guess of the operating point, say (s_0, t_0). Then we substitute an fBm model for the actual input traffic *at that point*, and using system (15) we compute the parameters of the fBm. Let's call fBm_0 the obtained fBm having parameters λ_0, σ_0, H_0. Next, we compute the new value (s_1, t_1) to be the operating point of the link fed by fBm_0. This can be computed explicitly by the equations

$$t_1 = \frac{H_0 B}{(1 - H_0)(C - \lambda_0)} \tag{23}$$

$$s_1 = \frac{B + (C - \lambda_0)t_1}{\sigma_0^2 t_1^{2H_0}}\,. \tag{24}$$

The method then iterates by using (s_1, t_1) as the current guess of the operating point.

It is easy to see that the operating point (s^*, t^*) is a fixed point in the above algorithm, since having (s^*, t^*) as the initial guess will produce a fBm implying the same operating point (see Subsection 3.5). Unfortunately we have not been yet able to prove convergence properties for the above procedure, although extensive experimentation has provided extremely positive results: mostly less than six steps needed, no divergent case observed. Figure 7 depicts the series of steps for two starting points.

We must add that the actual traffic (MPEG) used in the experiments produced a number of saddle points, but there was a substantial difference in the value of the optimum between the dominant one and the rest of the points, which were also rather flat compared to the dominant one. One could construct cases where more than one time scales for buffer overflows are likely, in which case the initial point for starting the algorithm could lead to different values for the operating point. Although theoretically possible, such a case has a small chance to occur in practice.

6 Conclusions

In this paper we have proposed a new notion of traffic equivalence based on preserving important properties of the large deviation coefficient. It can be used for constructing simple traffic generators in software which can emulate an arbitrarily large number of real sources from their actual traces possibly in real time. The accuracy of the approach increases with the size of the system and experimentation has shown very good results for the case of broadband ATM links, but nevertheless the approach is not particular to ATM.

Although the use of fBm for constructing traffic substitutes is promising, using processes with more parameters should be further investigated. This is important for extending the approach to match other measures of the queuing process, such as the mean queue length. Since this particular direction is well understood by now (e.g. see [9, 8]), combining these approaches should be feasible.

An important application of the fast calculation procedure of the previous section is the on-line calculation of the buffer overflow probability and the effective bandwidths of the input streams. The idea is to solve system (15) on-line. This can be done by estimating the value of the effective bandwidths and their derivatives in parallel on a grid of size ϵ around (s, t). For real-time traffic (hence small buffers) the value of t is small of the order of 10-100ms, and hence each such estimate can be done fairly accurately in 1-2s by measuring the actual traffic. Then using the above fast calculation procedure, the desired quantities can be available in 5-10s, which is of significant practical interest.

Acknowledgment: The authors are grateful to Frank Kelly and Vasilios Siris for useful discussions on the subject of the paper.

References

[1] P. Billingsley. *Probabilty and Measure*. Wiley, New York, 2nd edition, 1986.

[2] D. D. Botvitch and N. Duffield. Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. *Queueing Systems*, 1995.

[3] C. Courcoubetis, V. A. Siris, and G. D. Stamoulis. Application and evaluation of large deviation techniques for traffic engineering in broadband networks. In *ACM SIGMETRICS '98/ PERFORMANCE '98 Joint International Conference on Measurement and Modeling of Computer Systems*, Madison, Wisconsin, June 1998.

[4] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a switch handling many traffic sources. *Journal of Applied Probability*, 1996.

[5] A. Elwalid, D. Heyman, T. V. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13(6), 1995.

[6] P. Jelenkovic, A. Lazar, and N. Semret. The effect of multiple time scales and subexponentiality of MPEG video streams on queueing behavior. *IEEE Journal on Selected Areas in Communications*, 15(6):1052–1071, August 1997.

[7] F. P. Kelly. Notes on effective bandwidths. In F. P. Kelly, S. Zachary, and I. Ziedins, editors, *Stochastic Networks: Theory and Applications*, pages 141–168. Oxford University Press, 1996.

[8] L. A. Kulkarni and S.-q. Li. Measurement-based traffic modeling: Capturing important statistics. submitted to Journal of Stochastic Models.

[9] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J. D. Robbins. Performance models of statistical multiplexing in packet video communications. *IEEE Trans. on Comm.*, 36(7):834–844, July 1988.

[10] O. Rose. Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems. Technical Report No. 101, University of Wuerzburg. Institute of Computer Science Research Report Series, February 1995.

[11] B. K. Ryu and A. Elwalid. The importance of the long-range dependence of VBR video traffic in ATM traffic engineering: Myths and realities. In *Proc. of ACM SIGCOMM '96*, pages 3–14, August 1996.

[12] A. Simonian and J. Guibert. Large deviations approximations for fluid queues fed by a large number of on/off sources. *IEEE JSAC*, 13(7):1017–1027, August 1995.

[13] A. Weiss. A new technique for analyzing large traffic systems. *Adv. Appl. Prob.*, 18:506–532, 1986.

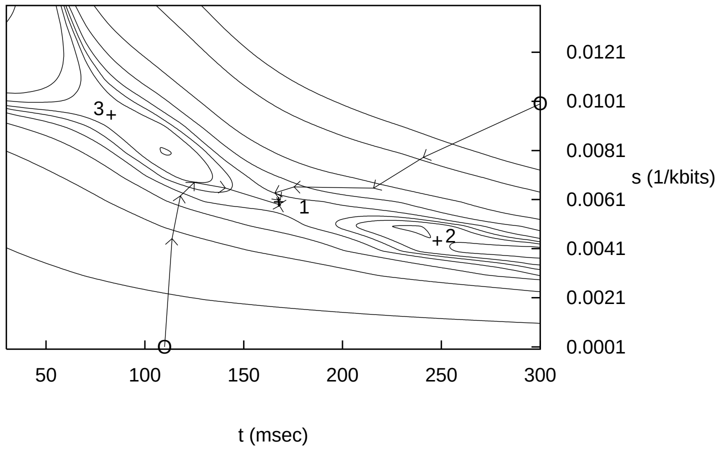

Figure 7: The steps taken to converge by the method for fast calculation of the operating point. 510 Star Wars sources are multiplexed in a ($C =$) 155 Mbps link with ($B =$) 2000 cells buffer size. The contour of $G_R(s, t)$ is depicted for the aggregate actual traffic. Two starting points are considered, one at $s = 0.01, t = 300$ and the other at $s = 0.0001, t = 110$. It can be seen that convergence is fast for both starting points. The operating point of the system is $s = 0.006$, $t = 168$ (marked as '1' at the figure), while there exist two more fixed (saddle) points at $s = 0.0042$, $t = 248$ and at $s = 0.0097, t = 82$ (marked as '2' and '3' respectively). For both starting points the method converged to the desired operating point.