Bypassing Modelling: Further Investigations of Entropy as a Traffic Descriptor in the Fairisle ATM network

Simon Crosby and Ian Leslie

Cambridge University Computer Laboratory

Meriel Huggard, John T. Lewis, Fergal Toomey*, and Cormac Walsh

Dublin Institute for Advanced Studies.

January 1996

*Corresponding author. Dublin Institute for Advanced Studies, 10 Burlington Road, Dublin 4, Ireland. Phone: +353 1 6680748. Fax: +353 1 6680561. E-mail: toomey@stp.dias.ie.

Abstract

The thermodynamic entropy of an ATM traffic stream is a useful tool for predicting cell-loss ratios and cell-delay. In [1] it was proposed that the entropy of a cell stream be estimated directly at a multiplexor; from this data, predictions can be made rapidly using algorithms which are simple enough to be executed on-the-fly. Preliminary investigations using simulated data have suggested that the method is practicable. This paper describes the results of an investigation of the method using real ATM traffic on the Fairisle ATM network [2]. Fairisle is equipped with a wide range of ATM traffic sources, including compressed video, audio and LAN data. In addition traffic can be generated from a pre-recorded trace of traffic activity. One such source is the Star Wars movie, reconstructed from a trace of the output of a DCT-based VBR video codec at Bellcore [3]. In [4] we described experiments involving the estimation of the entropy of both single-source and multiplexed cell streams based on the measured cell inter-arrival times. Accurate estimates were obtained at high switch loads. In this paper we continue these investigations, introducing two new methods of estimating the entropy, one based on splitting the arrival stream into two processes, and the other on observing the empirical queue formed by the traffic. The first of these leads to significant improvements in accuracy at intermediate loads, while the second is shown to be accurate at all loads.

1 Introduction

ATM networks will rely on statistical multiplexing to maximise resource utilisation. This raises the question of how best to allocate resources between competing demands. Clearly, some knowledge of the queueing characteristics of different traffic streams will be needed in order to make good decisions in this regard. The problem is made difficult by the fact that ATM networks will carry traffic of several different types, many of them too complex to be described by simple stochastic models. Ideally we require a method which will work with traffic of any type, which contains information about the traffic's queueing properties in a form which is easy to extract, and which is not overly complex or difficult to apply in practice.

In recent years it has been realised by many researchers that, for a very wide range of stochastic traffic models, the queueing properties of the model can be deduced from its large deviation rate function. Specifically, if the traffic is passed through a single-server queue with constant service rate s, then the resulting steady-state queue length distribution has exponential asymptotics of the form

$$\mathrm{P}\{Q > q\} \sim \exp -\delta q , \tag{1}$$

with a decay parameter δ which can be calculated from rate-function $I(\cdot)$ of the model:

$$\delta = \min_{c>0} I(c+s)/c . \tag{2}$$

Similar results are known to hold for a variety of other simple queueing systems: see, for example, [5, 6]; a general introduction may be found in [7].

As these results hold for a very wide range of models of arbitrary complexity one is naturally led to enquire whether 'real' traffic, of the type likely to be carried on an operational ATM network, also exhibits asymptotically exponential queueing behaviour. The answer to this question is 'yes', as described in [4], and therefore we have proposed that traffic can be characterised by measuring its rate function. Such a scheme would provide the necessary information about the impact of the traffic on the network in a form which is easy to use, and has the advantage that it should work with traffic from any source.

We call $I(\cdot)$ the 'entropy' of the arriving traffic, to draw attention to an analogy with thermodynamics. The entropy function of a gas, which mathematically speaking is nothing other than a large deviation rate-function, contains just that information needed to calculate the bulk properties of the gas. It is not the practice of engineers to construct complex models of a gas in order to calculate its entropy — they measure it. A similar approach may prove fruitful in the study of statistical multiplexing.

In [4] we investigated whether the rate function of traffic can be measured in practical situations. In this paper we return to this question, again using video traffic from the Fairisle ATM network. Fairisle is an experimental ATM platform built at the University of Cambridge Computer Laboratory to permit research into performance, architectural, and management issues in ATM networks [2]. It is based on a network architecture composed of ATM switches which are interconnected to form a general topology ATM LAN. The Fairisle switch is based on a simple, 4×4 input buffered crossbar. Switches of degree greater than 4 are constructed by using multiple 4×4 crossbars in a 2-stage delta interconnection network.

The hardware design of the Fairisle network components permits a high degree of experimental flexibility owing to the use of field programmable gate arrays which implement a number of features to support experimental work, including high resolution clocks and hardware time stamping of cells. Fairisle can be configured as either an input or output buffered switch. Each port is equipped with 2K cell buffers, 4 MB of RAM and an ARM RISC processor, which implements all policy-specific traffic control functions, using a flexible control plane which exports RPC-style interfaces [8, 9] for signalling and management. This has enabled us to implement a wide range of measurement facilities to accurately monitor network performance. In the experiments reported in this paper, Fairisle was configured as an output-buffered switch with FIFO queueing.

2 Measuring the rate function

Rather than estimating the rate function itself it is more convenient to estimate a transform of $I(\cdot)$ called the *scaled cumulant generating function* (CGF for short), defined by

$$\lambda(\theta) = \lim_{n\to\infty} \frac{1}{n} \ln \mathrm{E} \exp(\theta A_n), \tag{3}$$

where A_n denotes the number of cells arriving in an interval of length n. The CGF $\lambda(\cdot)$ is related to the rate function by the Legendre transform:

$$I(x) = \max_{\theta}\{x\theta - \lambda(\theta)\}, \tag{4}$$

where the maximisation is over all values of θ. The decay parameter δ can be calculated directly from $\lambda(\cdot)$ using the equation

$$\delta = \max\{\theta > 0 : \lambda(\theta) < s\theta\}. \tag{5}$$

Thus δ is the ordinate of the intercept of the graph of $\lambda(\cdot)$ and a line of slope s through the origin.

How might we estimate $\lambda(\cdot)$? Of course, the CGF of a traffic set can always be measured indirectly by passing the traffic through a queue served at constant rate s (implemented in software), and measuring the aymptotic decay parameter of the resulting queue length distribution. By repeating this experiment for various values of s, and using the formula

$$\lambda(\delta) = s\delta, \tag{6}$$

one may build up a picture of $\lambda(\cdot)$ for various values of θ. The results will naturally depend on the size of the traffic sample. If the traffic comes from a stationary, mixing source then one may associate a 'true' CGF with the source, and as the size of the traffic sample increases the estimates of $\lambda(\theta)$ obtained from the empirical queue should converge to the 'true' value. This estimation method may not be the most efficient, as it effectively involves simulating a single-server queue, and its indirectness is an additional disadvantage: values of $\lambda(\theta)$ for arbitrary θ must be found by interpolating between measured values of $s\delta$. Nevertheless we have found this method, which we refer to as the 'empirical queue' method, to be a useful tool for analysing the queueing properties of traffic, and it also provides a handy benchmark test for other estimation methods.

The CGF may be estimated directly by taking its definition, equation (3), as a starting point. It can be shown that for a very large class of arrivals processes it is possible to find a block-length T such that the aggregated arrivals A_T are approximately independent and identically distributed; in this case the CGF is given approximately by

$$\lambda(\theta) \approx \lambda_T(\theta) = \frac{1}{T} \ln \mathrm{E}\ \exp(\theta A_T). \tag{7}$$

The most straightforward approach to estimating $\lambda(\cdot)$ therefore is to use the CGF of the empirical distribution of the block arrivals:

$$\hat{\lambda}_T(\theta) = \frac{1}{T} \ln \frac{1}{N} \sum_{i=1}^{N} \exp(\theta A_T^{(i)}). \tag{8}$$

The advantages of this estimator, which we will refer to the 'empirical measure' estimator, are its simplicity and the fact that it makes minimal assumptions about the traffic. To use it in practice, we must make a choice concerning the block length T. This should be sufficiently large so that successive observations of A_T are approximately independent, but not so large that short bursts of activity, which make an important contribution to the queue-length, are effectively smoothed out. Observations of A_T can be taken not only from disjoint blocks, but from all blocks of length T. This naturally increases the correlations between blocks, but also increases the number of samples; the net result is an improvement in the properties of the estimator.

We have found that estimator (8) works well with certain types of simulated traffic (see [1]). In [4] we applied it to video traffic from the Fairisle network and found that, while results remain good at high loads, at lower loads the estimator's performance deteriorates: it typically over-estimates the decay parameter of the queue length distribution. This video traffic was found to contain both strong correlations and short bursts of very high activity. We hypothesised that this combination caused the poor performance of the estimator. The correlations necessitate a large block size T; as a result, the estimator cannot 'see' the short bursts, which are smoothed out of the data during aggregation.

We have therefore experimented with a modification of (8) which involves splitting the traffic into high activity and low activity observations, and estimating the CGF of each separately. For the low activity observations a large block size may be used to ensure that successive blocks are independent; for the high activity observations one may use a smaller block size so that short bursts are not smoothed out. Making the assumption that the two sets of observations are approximately independent one may then recombine the estimates using the prescription,

$$\lambda(\theta) = \lambda_H(\theta) + \lambda_L(\theta), \tag{9}$$

where $\lambda_H(\cdot)$ and $\lambda_L(\cdot)$ represent, respectively, the CGFs associated with the high and low activity observations. Clearly this method of splitting the traffic is a crude one, and more sophisticated techniques could undoubtedly be applied. Nevertheless this scheme, referred to as the 'split streams' technique, provides a first test of the hypothesis that splitting the data may improve performance.

3 Experiments with simulated traffic

Before applying an estimation method to 'real' ATM traffic it is sensible to explore its properties using traffic from a simple simulation model. The advantage of using a model is that the statistical properties of the estimator may be examined in a context where stationarity is assured. Moreover, the correct value of the decay parameter may readily be calculated for many models; this allows the accuracy of the estimation scheme to be assesed. (For a discussion of how to calculate δ for Markovian traffic models, see [10]).

With these aims in mind we have performed a number of experiments using simulated traffic from a four state discrete time Markov model. This model was chosen in order to test our hypothesis mentioned above concerning the performance of the empirical measure estimator. Its parameters were adjusted to produce traffic containing short bursts of high activity combined with strong positive correlations.

Figure 2 contains histograms of δ estimates made with simulated traffic using the empircal measure and the empirical queue estimators. One hundred estimates were taken from each method at four different service rates corresponding to loads of 65%, 75%, 85%, and 95%. In the figure the horizontal axes of the plots have been rescaled to show the ratio of the estimates to the true value of δ. The estimates from the empirical queue are seen to be clustered around the theoretical value, rarely straying from it by more than about 10% of its value (however, since δ increases as the load is decreased, the absolute spread of the estimates therefore also increases).

The estimates from the empirical measure, on the other hand, exhibit bias, being typically larger than the theoretical value of δ. The bias increases at lower loads, reproducing the effect seen with video traffic from Fairisle. This confirms our hypothesis that strong correlations and short bursts lead to poor performance with this estimator.

4 The Starwars video traffic

We turn our attention in this section to results obtained in an experiment with 'real' ATM traffic, namely the 'Starwars' video data set. By now well known in the ATM traffic research community, the Starwars data set was produced by Garrett and Vetterli at Bellcore. It has been studied in some detail, for example in [3], and comprises the information content in bytes per frame for 171,000 frames (approximately 2 hours) of the film, as transmitted by a Discrete Cosine Transform (DCT) based codec, with a compression algorithm similar to that of JPEG. The byte count per frame is broken down further to the 'slice' level – each slice is 16 video lines. With 30 slices per frame and a frame rate of 24 per second (i.e. the original film rate), one slice represents the information transmitted in about 1.4 milliseconds.

Our data was constructed from the Starwars traces by encoding each slice of video as a single AAL5 PDU; each PDU was transmitted in such a fashion that its cells were evenly spaced throughout the slice time. This smoothing effect over the slice time serves to reduce the peak rate of the source from the line rate (100 Mb/s, if all cells in a slice are transmitted back-to-back) to about 24.2 Mb/s for the worst slice, and is typical of ATM video sources (see, for example [9]). The mean rate of the source is about 5.3 Mb/s.

Our experiment was based on the transmission of a single Starwars video source across an output buffered Fairisle switch, as described previously. The absolute arrival time of each cell in the film was recorded, and the recorded data was passed offline for analysis. The task of making these measurements was found to be well within the capacity of the switch port controller, which was operating at loads between 70% and 80%. The question of exactly how heavily the switch can be loaded while making measurements is a topic for future investigation.

4.1 Statistical properties of the Starwars traffic

Figure 1 shows the measured activity on the network caused by a single Starwars source in the absence of other traffic. Clearly the activity of the source is highly variable over time. This results from variation in the compression ratio acheived by the codec. Some video frames have a large number of high frequency components and do not compress well. In addition the activity observations are strongly correlated over timescales of the order of a few seconds. This is to be expected, as frames within the same scene contain similar information and hence generate similar numbers of cells. Long term fluctuations in the source activity can also be produced by, for example, transitions from 'action' sequences to stationary shots, and perhaps even by changes in the background scenery. In view of the nature of the source, it does not seem appropriate to attempt to model the whole of the Starwars traffic set using a stationary stochatic model. A model fitted to a few seconds of data, for example, would remain valid only over timescales of the order of seconds. At longer timescales, further fluctuations in the traffic would render the model invalid.

As mentioned in section 2 the traffic is found to contain a preponderance of short bursts of high activity. From figure 3 one can clearly see the heavy-tailed nature of the activity distribution. Although infrequent, these spikes have important consequences for the queueing properties of the traffic as they make the dominant contribution to the queue.

4.2 Queueing properties of the Starwars traffic

Empirical queue length distributions made with the Starwars traffic are found to exhibit approximately exponential asymptotic behaviour, in accordance with what one might expect from large deviation theory. We have performed a number of experiments to investigate how the decay parameter δ varies over time as the whole of the Starwars traffic is passed through a queue, and also how δ varies between different subsets of the traffic. These experiments also indicate how the three estimation schemes described in section 2 differ in performance on the Starwars traffic.

Figure 4 contains graphs of δ estimates against sample size. For each sample size, ranging from 1700 to 170,000 frames, the service rate was chosen so as to keep the load over the sample constant. The experiment was performed using load values of 75%, 85%, and 95%, and taking estimates of δ from each of the three estimation schemes. The estimates made from the empirical queue show that the queueing properties of the traffic are very stable over time, especially at lower loads. This is remarkable given that the traffic contains fluctuations at time scales of the order of the size of the entire sample. It appears that segments which cause large variations in the queueing properties of the traffic are very rare, especially at low loads. This experiment was repeated using a fixed service rate for all sample sizes. The resulting δ estimates contained even less variation, being in fact constant over much of the film at low loads. At these loads it was found that only a few short segments of the traffic made any contribution to the queue.

The values of δ obtained using the empirical measure estimator are typically larger at low loads than those made from the empirical queue. However, splitting the arrivals into high and low activity streams leads to a considerable improvement in performance. The 'split stream' estimates are more accurate than those made from the empirical measure, and at higher loads reproduce much of the variation in the estimates from the empirical queue. Thus they provide a reasonably good impression of the actual queueing behaviour of the traffic. Results obtained while keeping the service rate fixed were very similar.

Our second experiment consisted of splitting the traffic into 100 disjoint segments of equal length and estimating δ for each. Again the service rate was chosen to keep the load constant from sample to sample, and the experiment was repeated with different values for the load. The results are depicted in figure 5, which contains histograms of estimates obtained from each of the three estimation methods. The bulk of the estimates fall within a small range of values: again, a remarkable finding given the nature of the traffic. Each of the distributions also contains a fair number of outliers at larger values, which come from occasional segments with more favourable queueing properties than the bulk of the traffic.

From the histograms, it appears that the two CGF estimation methods are able to pick out the main hump of the distribution of δ, and also reproduce many of the outliers. The 'split stream' method again gives better performance in this respect, providing a useful picture of the distribution even at low loads.

5 Discussion

In this paper we have discussed three methods of estimating the rate function of a traffic stream, and examined their performance with both simulated traffic and a measured traffic stream taken from the Fairisle ATM network. The first method, based on measuring the decay parameter of the empirical queue length distribution, was found to work very reliably with the simulation model, producing δ estimates close to the known theoretical value with an acceptably small spread. We are therefore confident that if the rate function of a traffic stream exists at all in any meaningful sense then it may be measured using this method.

The results obtained with the Starwars traffic were somewhat surprising. The traffic contains fluctuations at many different timescales, and would be difficult or impossible to model using a stationary stochastic model. Nevertheless, estimates of δ from the empirical queue were remarkably stable both over time and between different subsets of the data. This makes us hopeful that measurements of δ and the CGF will prove to be useful tools for the purpose of analysing and predicting the queueing behaviour of traffic.

The experiments with simulated traffic using the CGF of the empirical arrivals distribution confirm that this estimator has difficulties with traffic containing a combination of strong correlations and short bursts of high activity. These features are both present in the Starwars traffic, and for this reason the estimator's performance at low loads with Starwars is not good. However splitting the arrivals into high and low activity streams leads to a considerable improvement. The 'split streams' estimates are accurate at high loads and give a reasonable qualitative picture of the traffic's queueing behaviour at lower loads. Although these estimates are not yet as good as those from the empirical queue, they show that improvements can be made by adapting the estimator to the traffic.

Acknowledgements: We thank Amir Dembo for suggestions which led us to use the simple estimator (8).

References

[1] N. G. Duffield, J. T. Lewis, N. O'Connell, R. Russell, and F. Toomey. Entropy of ATM traffic streams. *IEEE Journal on Selected Areas in Communications*, 13(6), August 1995.

[2] R. Black, I. Leslie, and D. McAuley. Experiences of building an ATM switch for the Local Area. In *SIGCOMM*, volume 24(4), September 1994.

[3] M. W. Garrett and M. Willinger. Analysis, modelling, and generation of self-similar vbr video traffic. In *Proceedings, 1994 SIGCOMM Conference*, pages 269–280, London, 1994.

[4] S. Crosby, I. Leslie, J. T. Lewis, N. O. Connell, R. Russell, and F. Toomey. Bypassing modelling: an investigation of entropy as a traffic descriptor on the fairisle ATM network. In *Proceedings of the 12th UK Teletraffic Symposium*, Windsor, March 1995.

[5] N. G. Duffield and N. O. Connell. Large deviations and overflow probabilities for the general single server queue, with applications. To appear in Proc. Cam. Phil. Soc.; available as: ftp:stp01.stp.dias.ie/DAPG/dapg9330.ps.

[6] N. O'Connell. Large deviations for queue lengths at a multibuffered resource. Submitted to Appl. Prob.; available as: ftp:stp01.stp.dias.ie/DAPG/dapg9434.ps.

[7] A. Schwartz and A. Weiss. *Large Deviations for Perfomance Analysis*. Chapman and Hall, 1995.

[8] R. Black and S. Crosby. Experience and results from the implementation of an ATM socket family. In *Proceedings USENIX Winter Technical Conference*, San Francisco, January 1994.

[9] U. of Cambridge Systems Research Group. ATM document collection. Technical note, University of Cambridge Computer Laboratory, Pembroke Street, Cambridge, CB2 3QG, UK, Mar. 1994. available as: http://www.cl.cam.ac.uk/Research/SRG/bluebook.html.

[10] N. G. Duffield. Exponential bounds for queues with markovian arrivals. To appear in Queueing Systems; available as: ftp:stp01.stp.dias.ie/DAPG/dapg9301.ps.

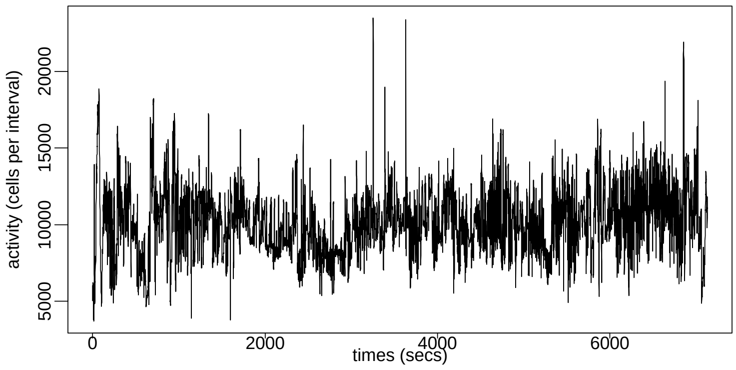

Figure 1: The Starwars video traffic

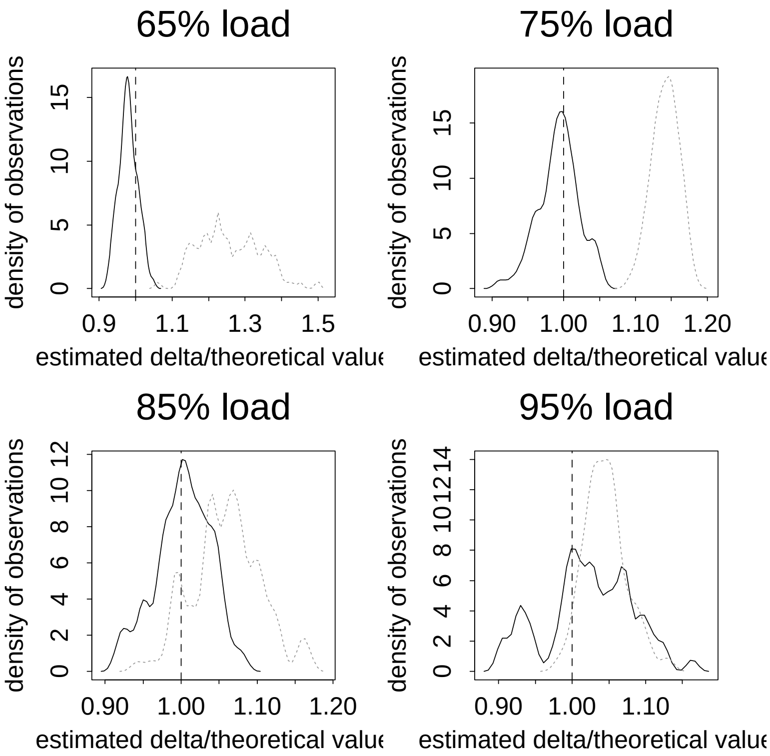

Figure 2: Distribution of delta estimates from simulated traffic. The solid line represents estimates from the empirical queue, the dotted line estimates from the empirical measure.

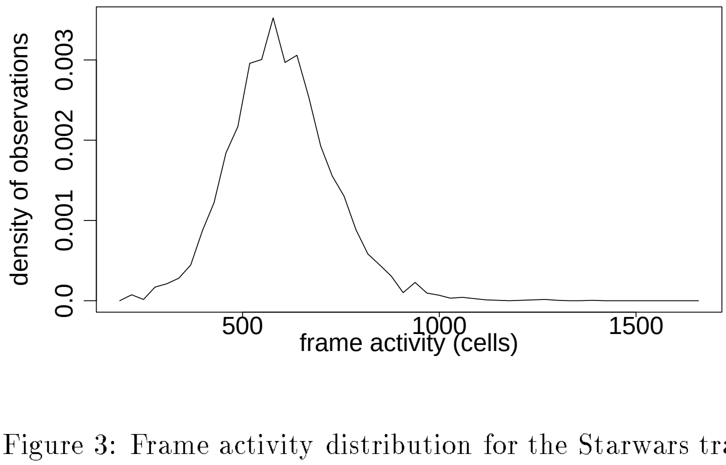

Figure 3: Frame activity distribution for the Starwars traffic

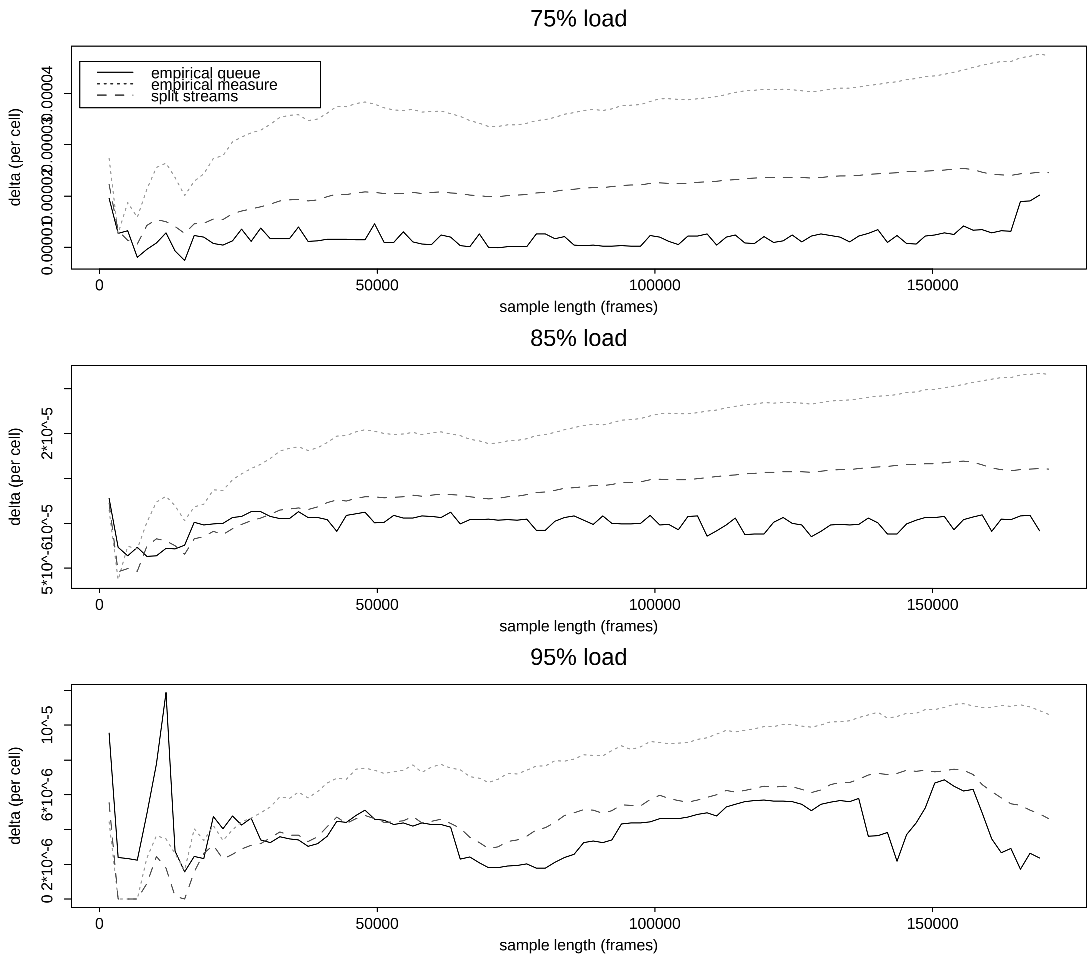

Figure 4: Delta estimates vs. sample length for the Starwars traffic

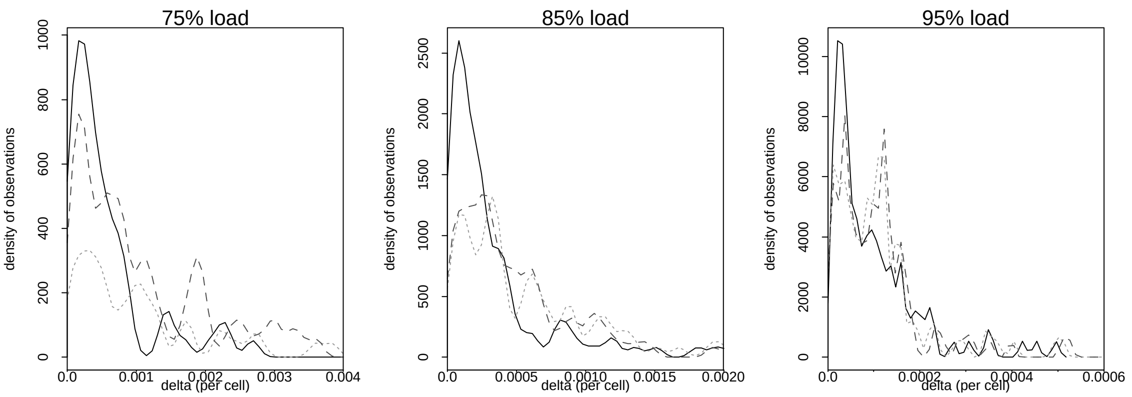

Figure 5: Distribution of delta estimates from disjoint segments of the Starwars traffic. The solid line represents estimates from the empirical queue, the dotted line estimates from the CGF, and the dashed line estimates obtained from the CGF using the 'split streams' technique.