DIAS-STP-94-07

THE ENTROPY OF AN ARRIVALS PROCESS: A TOOL FOR ESTIMATING QoS PARAMETERS OF ATM TRAFFIC

N.G.Duffield[1], J.T.Lewis[23], Neil O'Connell[2],
Raymond Russell[2] and F.Toomey[2]

ATM switching gives rise to new problems in queueing theory, problems which cannot be solved using the mathematical tools in common use in conventional queueing theory. The origin of the these new problems is to be found in the *bursty character* of the stream of data packets arriving at an ATM buffer. Classical queueing theory is applicable to telephone traffic at the call level: incoming calls to an exchange may be assumed to be independent so that the input stream can be modelled by a Poisson process; consequently, classical theory concerns itself with queues with Poissonian inputs. The queues which arise in the buffers of ATM multiplexers are queues of data packets waiting to be processed and the statistical characteristics of their input streams are very different from those of telephone traffic viewed at the call level. Since data packets are transmitted in bursts, their arrivals are strongly correlated and cannot be modelled well by a Poisson process.

We give a brief exposition of our strategy and then we describe some of the algorithms we have devised to implement it. Our aim is to estimate the Quality of Service (QoS) parameters of a buffer.

The QoS parameters we are interested in are: cell-loss frequency, cell-delay. They arise in the following way: cells arrive at an ATM switch in bursts and are stacked in a buffer of size b, the *buffer-size*, which is emptied on a first-come-first-served basis at a fixed rate s, the *service-rate*. The number of cells Q waiting to be served is the *queue-length*; if the queue-length exceeds the buffer-size, then the buffer overflows and cells are lost; the *cell-loss frequency* is the rate at which cells are lost because of buffer overflow. The number of clock-cycles a cell spends stacked in the buffer waiting to be served is the *cell-delay time*. It is important, for a variety of reasons, to be able to estimate the QoS parameters for a sample of traffic and see how they depend on buffer-size and service-rate: buffer-dimensioning, congestion management, traffic-shaping. Typical values being discussed as upper limits for the cell-loss ratio for an ATM switch range between 10^{-8} and 10^{-11}. Typical values mentioned for cell-delay in an ATM switch range between 100 μs and 1000 μs with an upper limit on the jitter of 100μs at a 10^{-10} quantile (meaning that the probability that the delay in the switch exceeds 100 μs is less than 10^{-10}). We shall concentrate here on the problem of estimating the cell-loss frequency. It is clear that very small probabilities are involved; dealing with these has become possible because of the rapid development in recent years of the theory of large deviations, providing exponential bounds for the probabilities of rare events. It is worth remembering that, although the probabilities we are dealing are so small that, in many applications of probability theory, they would be interpreted as meaning "the event never occurs", that is not the case here: at a transmission rate of one Gigabit per second, a cell-loss ratio of 10^{-8} means that, on average, we lose one cell per minute.

A Simple Model It is instructive to consider first a simple model, the *two-state Markov model,*

[1]School of Mathematical Sciences, Dublin City University, Dublin 9, Ireland
[2]Dublin Institute for Advanced Studies 10 Burlington Road, Dublin, Ireland
[3]Author presenting paper

which exhibits many of the features of bursty ATM traffic:

- the absence of a cell in a clock-cycle is represented by 0, the *dead* state, the presence of a cell in a clock-cycle is represented by 1, the *active* state;
- the model is described by two parameters: the probability a of switching from a dead state to an active state, and the probability d of switching from an active state to a dead state;
- there are L input lines, the traffic on each being independent of the traffic on the others but all having the same distribution;
- the service-rate, the maximum number of cells that can be processed in a clock-cycle, is denoted by s.

The number of active lines at each clock-cycle is called the *arrivals* process; this process is *stationary* — its statistical characteristics do not change with time. The *activity*, the fraction of time spent in the active state, is $a/(a+d)$. We take $1-(a+d)$ to be a measure of the *burstiness* — zero burstiness corresponds to independent cell-arrivals. The Duffield–Buffet formula [2] gives a safe bound for $f(L,b,s)$, the frequency of cell-loss when L lines input a buffer of size b emptied at service-rate s:

$$f(L,b,s) \leq \exp(-L\mu(s/L)\,)\exp(-b\,\delta(s/L)\,)\,, \tag{1}$$

where the coefficients μ and δ can be got from the parameters a and d by solving numerically two equations. This formula has some nice features:

- it tells us that the graph of $\log f$, the log-frequency of cell-loss versus the buffer-size b always lies below the straight line with slope $-\delta$ and intercept $-L\mu$;
- the coefficients μ and δ do not depend on b and depend on s and L only through the ratio s/L, the *effective band-width*;
- the pre-factor $\exp(-L\mu(s/L)\,)$ exhibits the economy of scale to be got from statistical multiplexing: if we keep the effective band-width s/L fixed, the frequency of cell-loss f decreases exponentially fast with the number L of input-lines (it is more profitable to multiplex $2L$ lines into an optical fibre of capacity $2s$ than to have two optical fibres, each of capacity s, serving L lines each);
- the slope $-\delta$ of the straight-line bound is, in fact, the slope of the graph of $\log f$ for large b: the slope of the bound is asymptotically correct.

The practical importance of the Duffield–Buffet formula is:

> *a straight-line bound can be used to give safe bounds for the QoS parameters; a bound for the frequency of cell-loss can be got directly from the straight-line and a bound for the queueing-delay jitter is easily calculated.*

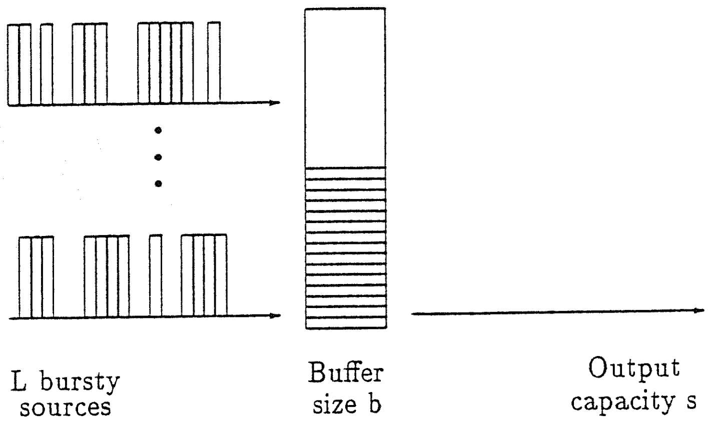

Figure 1: Simple Model of an ATM Multiplexer

Simulations We can see the Duffield–Buffet formula in action, using simulated data. Getting sufficient data can be a problem: simulations of input streams can take a long time to run since we are interested in events which occur with very small frequencies. Nevertheless, we have succeeded in doing this using two programmes: the *simulator* and the *analyser*.

In the simulator programme, the user-specified parameters are the *mean activity* and the *burstiness* of the two-state Markov model, together with the *total run-length* (in clock-cycles). The programme computes iteratively the number of active lines in each clock-cycle; we call this output the *arrivals* process.

The user-specified parameter of the analyser programme is the *service rate*, an integer-multiple of clock-cycles. The programme takes an arrivals process and feeds it into a water-marked virtual buffer which is emptied at the specified service-rate. The output is the log of the fraction of the total run-length for which each water-mark level was exceeded; this is interpreted as the log-frequency of cell-loss as a function of buffer-size.

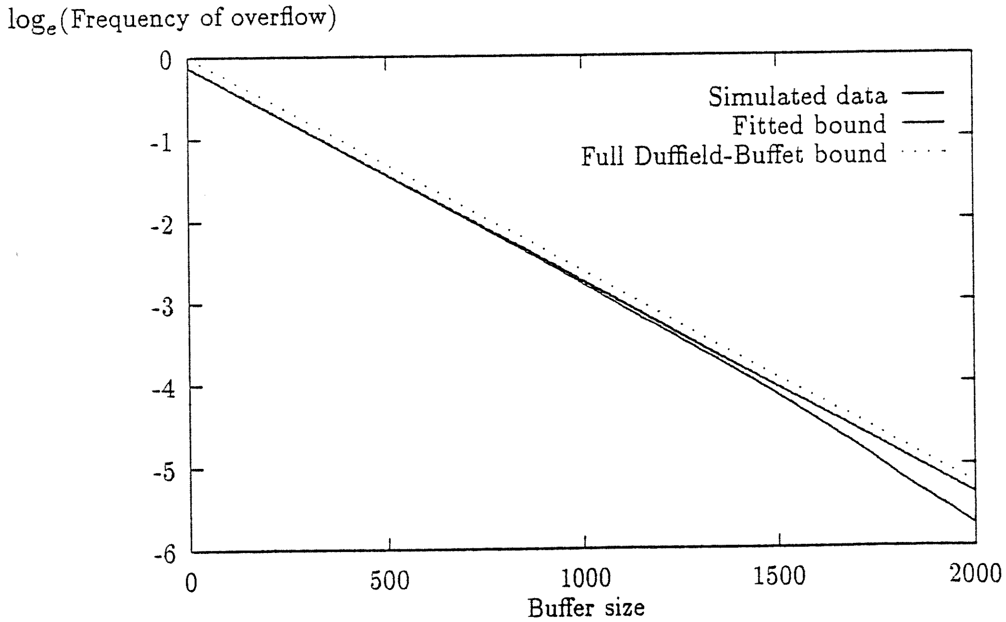

Figure 2: Comparison of simulation
with Duffield–Buffet bound

Figure 2 shows the results from a simulation as well as the Duffield–Buffet bound. In addition, we show an empirical bound got by fitting a regression line to the linear part of the data and shifting it upwards to the first non-zero data point, making it mimic the behaviour of the theoretical bound. Note that the two bounds are almost parallel. These results are typical of all the simulation results, curving initially at small buffer sizes, but quickly reaching the asymptotically linear regime. They eventually fall away when the data becomes too sparse, but we have not shown that region in this plot.

Using Real Data So far, we have no ATM data to investigate; we have made do with Ethernet data, processed to simulate ATM traffic. We have used two sources: Trinity College, Dublin, Computer Science Network; Bellcore Morristown Laboratory Network. In the case of TCD, the data was collected using the *etherfind* command on the SUN operating system. The Bellcore data, made available to us by Dr.W.Willinger, was collected using a specially designed high-speed device. The user-specified parameter in the *traffic processor* programme is a *clock cycle*, the basic time-slot. The raw data taken from the TCD network is a record, for each packet, of: time of arrival, size (in bytes), origin and destination. The programme uses the origin and destination to determine the *line* (virtual circuit) to which each packet belongs. Next, it uses the time-of-arrival and the packet-size to output an estimate of the number of active lines in each clock-cycle. We call this output the *arrivals* process from the Ethernet network; it is processed by the analyser, just as is done in the case of simulated traffic.

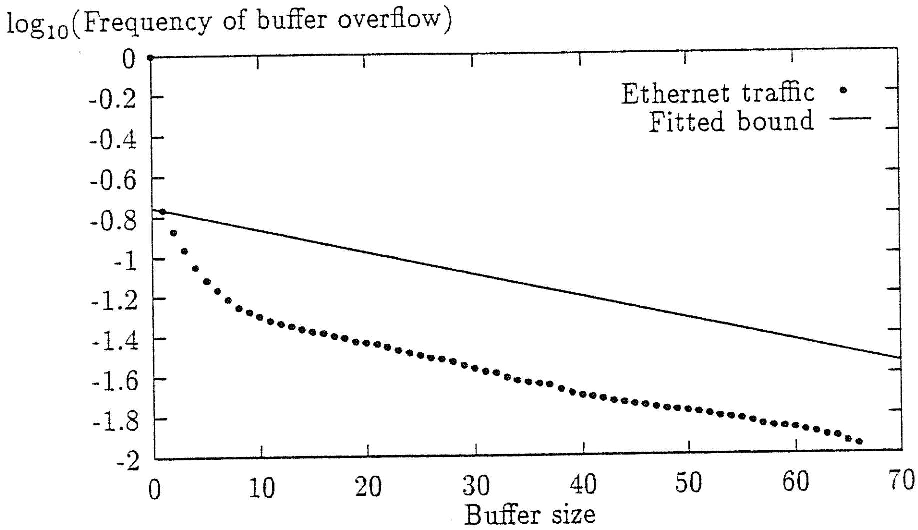

Empirical Estimation of QoS Parameters

We have seen that it is possible to take an arrivals stream and produce a plot of log-frequency of cell-loss versus buffer-size; from the plot, we can determine the asymptotic slope of the graph by fitting a regression line to the linear part of the plot. As we explained in Section 2, the straight-line bound can be used to compute bounds on the QoS parameters. The draw-back with this procedure is that it is difficult to get a good determination of the slope of the regression line with real data. Nevertheless, we have found the log-frequency versus buffer-size plot got from watermarking to be very useful in our work; it provides a visual test of any proposed scheme for determining a straight-line bound. For example, we tried out the method used in the MINOS algorithm (Monitor to Infer Network Overflow Statistics), devised by Courcourbatis et al [3] to enable an ATM switch to to predict its spare capacity. To estimate the frequency of overflow in a buffer of size B, the MINOS algorithm keeps track of the occupancy of three much smaller buffers of size B0, B1 and B2 and fits a three-parameter curve through their overflow frequencies. For simulated data got using a two-state Markov model, the results are reasonably good, as judged against the plot got from watermarking using our analyser programme. But, on real Ethernet traffic, we found the MINOS algorithm to be extremely sensitive to the choice of the virtual buffer-sizes B0, B1 and B2 and we know of no way of optimizing this.

Using the Rate-Function There is a better way of estimating the asymptotic slope. To explain it, we have to go back to the idea of queue-length. The current queue-length Q is the number of cells which have arrived less the number which have been served; to be precise, we have to make a couple of definitions:

- the number A_n of cells which have arrived in the preceding n clock-cycles is called the *arrivals* process;
- the *work-load* process W_n is A_n less the number of cells which *could* have been served in the preceding n clock-cycles:

$$W_n \ := \ A_n - ns \ ; \tag{2}$$

then a remarkable formula of queueing theory, see [1], gives

$$Q \ = \ \max_{n \geq 0} W_n \ . \tag{3}$$

Under very general conditions on the arrivals process, we know that $\log p(n, w)$, the log-probability that W_n/n exceeds the level w, is asymptotically linear in n:

$$\log p(n, w) \ \sim \ -n \, I(w). \tag{4}$$

(In the theory of large deviations, developed in recent years to estimate the probability of rare events, the coefficient I is called the *rate-function.*) Denote by $f(b, s)$ the probability that the current queue-length Q exceeds b for given service-rate s; the theory of large deviations tells us [4] that

> $\log f(b, s)$ *is asymptotically linear in* b *whenever* $\log p(n, w)$ *is asymptotically linear in* n*:*
>
> $\log p(n, w) \ \sim \ -n \, I(w)$ *implies* $\log f(b, s) \ \sim \ -b\delta$ *;*
>
> *moreover,* δ *is determined by the rate-function* I*:*

$$\delta \ = \ \min_{w>0} I(w)/w \ . \tag{5}$$

It follows that if we can determine I, then we can compute $-\delta$, the asymptotic slope of the graph of log-frequency of cell-loss versus buffer-size and so estimate the QoS parameters.

The traditional way to proceed in such a situation is to choose a model of the arrivals process (such as the two-state Markov model), use the data to estimate the adjustable parameters of the model (a and d in the two-state Markov model) and then compute the rate-function I. The difficulty with this procedure is that it is very difficult to find a satisfactory model for real traffic; indeed, it has even been suggested that Ethernet data is too complex to model at all. Certainly, it is unlikely that bursty traffic can be described by a model with a small number of parameters. But a full characterization of the traffic is not required in order to compute $-\delta$; we saw that all we need is the rate-function of the arrivals process, so why not estimate the rate-function directly from the data, by-passing the modeling? There is a good reason for believing that this might work.

The Analogy with Statistical Thermodynamics There is a simple statistical model of an ideal gas; using it, we can compute the thermodynamic entropy function from which all the bulk properties of the gas can be calculated. In the case of a real gas, we could take a much more elaborate statistical model, one with several adjustable parameters; we could adjust the parameters to fit the model to experimental data and then use the model to compute the thermodynamic entropy function. Since the bulk properties of a gas can all be calculated from the thermodynamic entropy function, chemical engineers by-pass the modeling procedure and measure the thermodynamic entropy function directly.

The relevance of this analogy to the problem of characterizing bursty traffic is that, from a mathematical point of view, the entropy function is just he negative of a rate-function: the problem of making sense of Boltzmann's celebrated formula equating the thermodynamic entropy of a macroscopic state with the logarithm of the volume of the corresponding set in phase-space is the same as the problem of making precise the relation

$$\log p(n, w) \sim -n\, I(w). \tag{6}$$

Our claim is that, for the purpose of estimating QoS parameters, it is enough to know the rate-function of the arrivals process; the modeling procedure can be by-passed if we can estimate the rate-function directly.

Estimating the Rate-Function It turns out that it is better to estimate a transform of the rate-function, rather than the rate-function itself. The *scaled cgf* (cumulant generating function) λ of the arrivals process is defined by

$$\lambda(\theta) := \lim_{n\to\infty} \frac{1}{n} \log E \exp(\theta \sum_{i=1}^{n} X_i) . \tag{7}$$

Provided this limit function exists (and satisfies some technical conditions), the rate function $I(w)$ exists and is determined by it; the slope $-\delta$ can be calculated directly from λ using the formula

$$\delta = \max\{\theta : \lambda(\theta) \le s\theta\} . \tag{8}$$

The scaled cgf exists for a wide class of stationary processes. If we are to estimate λ empirically, it is important that the arrivals process be approximately stationary; provided this is the case, and the process is *mixing* in the sense that there exists a *block-size* k such that the block-sums

$$\tilde{X}_1 := \sum_{i=1}^{k} X_i, \quad \tilde{X}_2 := \sum_{i=k+1}^{2k} X_i, \ ... \tag{9}$$

are approximately independent and identically distributed, we can estimate λ using

$$\widehat{\lambda}(\theta) := \frac{1}{k} \log \frac{1}{m} \sum_{i=1}^{m} e^{\theta \tilde{X}_i} \tag{10}$$

and δ using

$$\widehat{\delta} := \max\{\theta : \widehat{\lambda}(\theta) \le s\theta\} . \tag{11}$$

We have written a programme which takes the arrivals stream, estimates the scaled cgf λ and then computes δ. This is a much more efficient procedure the ones described earlier and offers the prospect of on-line estimation of QoS parameters, a topic we are at present investigating.

Notes Several proofs of formula (5) have appeared in the literature: a heuristic argument can be found in Kesidis, Walrand and Chang [6] and proofs under very general conditions in Glynn and Whitt [5] and in Duffield and O'Connell [4]; further bibliographical details can be found in de Veciana, Courcoubetis and Walrand [11].

The connection between the theory of large deviations and thermodynamics was pointed out by Ruelle [9] and Lanford [7] over twenty-five years ago but only now is it beginning to be exploited; it is explained in a recent paper by Lewis and Pfister[8]. An account of related ideas addressed to a general audience has been given by Ruelle [10].

Acknowledgements This work was supported by grants from EOLAS and Mentec Computer Systems Ltd. under the Higher Education –Industry Cooperation Scheme. We are grateful to Walt Willinger for making the Bellcore data available to us.

References

[1] A.A.Borovkov, Stochastic Processes in Queueing Theory, Springer, Berlin (1976)

[2] E.Buffet, N.G.Duffield, Exponential Upper Bounds via Martingales for Multiplexers with Markovian Arrivals, to appear *J.Appl.Prob.* (1994)

[3] C.Courcoubetis, G.Kesidis, A.Ridder, J.Walrand and R.Weber, Admission control and routing in ATM networks using inferences from measured buffer occupancy, to appear in *IEEE Trans. Comm.*

[4] N.G.Duffield, N.O'Connell, Large Deviations and Overflow Probabilities for the General Single-Server Queue, with Applications DIAS Research Report DIAS-STP-93-30 (1993)

[5] P.W.Glynn, W.Whitt, Logarithmic Asymptotics for Steady-State Tail Probabilities in a Single -Server Queue, to appear *J.Appl.Prob.* (1994)

[6] G.Kesidis, J.Walrand, C.S.Chang Effective Band-widths for Multiclass Markov Fluids and other ATM Sources, preprint (1993)

[7] O.E.Lanford, Entropy and equilibrium states in classical mechanics, in Statistical Mechanics and Mathematical Problems, A. Lenard, ed., Lecture Notes in Physics 20, Springer, Berlin (1973)

[8] J.T.Lewis, C.-E.Pfister, Thermodynamic Probability Theory: some aspects of large deviations, preprint DIAS-93-33 (1993),to appear in *Theor.Prob.Appl.*

[9] D.Ruelle, Correlation functionals, *J.Math.Phys.* 6: 201 (1965)

[10] D.Ruelle, Chance and Chaos, Princeton U.P., Princeton (1991)

[11] G.de Veciana, C.Courcoubetis and J.Walrand, Decoupling bandwidths for networks: a decomposition approach to resource management, Memorandum # UCL/ERL M93/50, University of California

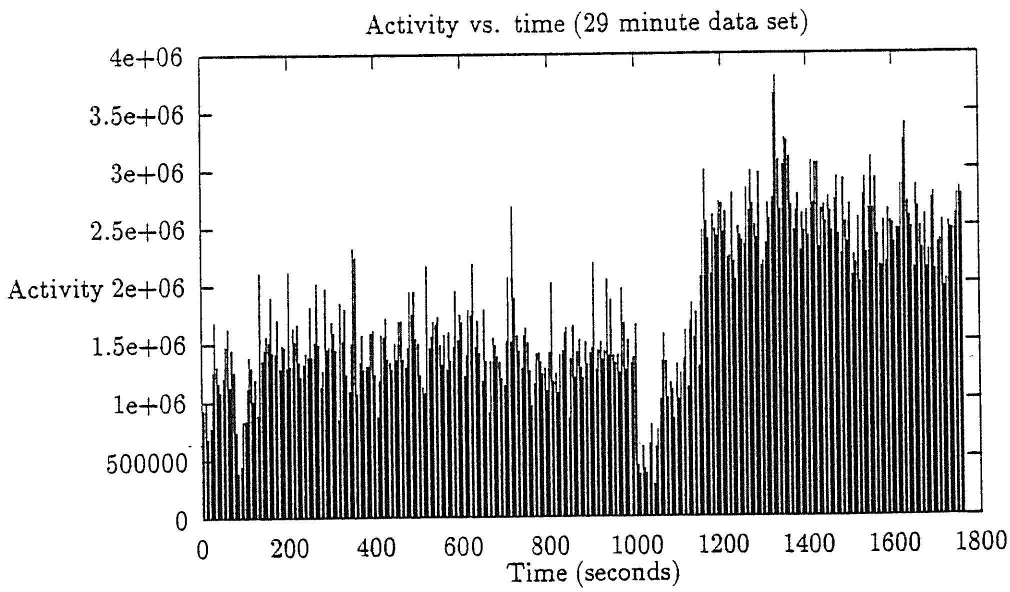

Figure 4: Aggregated activity of Bellcore data.

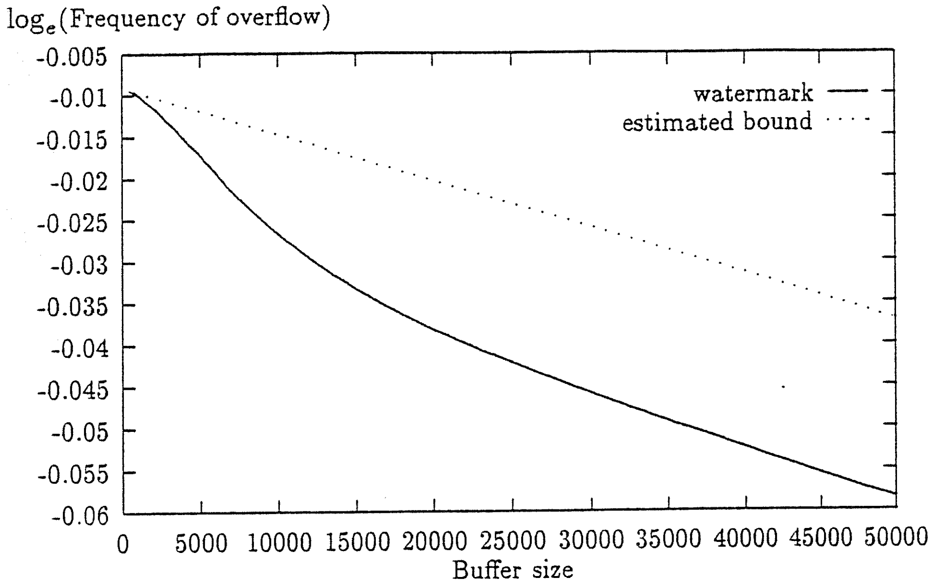

Figure 5: Watermark plot for a stationary section of the Bellcore data: the slope of the bound was obtained from an empirical estimate of the scaled cumulant generating function.