

Modeling of Internet Traffic Data

Rika Taute *

As submitted for SATNAC, May 2006

Abstract--- This paper studies the progress made to date on the modeling of internet traffic. It more specifically looks at the repeatedly reported occurrence of extreme value behavior in aggregate internet traffic. A random variable that follows a heavy-tailed distribution can give rise to extremely large values with non-negligible probability. The use of an appropriate extreme value traffic model could be useful in avoiding overcrowding on network lines due to extreme behavior.

Index terms--- extreme values, internet traffic, long-range dependence, self-similar.

Introduction

A deep understanding of Internet traffic can contribute to network performance monitoring, equipment planning, quality of service, security, and the engineering of Internet communications technology. A lot of research has already been done towards understanding the Internet but little has been done to model observations in Internet data. [1]

This paper aims to give an effective analysis and summary of proposed traffic models as a much needed platform for progress in this area. Special attention is given to the modeling of extreme values or tail behavior of internet traffic; as network engineers need to plan particularly for extremely large numbers in aggregate internet traffic. The paper will proceed as follows: first the basics on data and internet is given, then the different models with their positive and negative features is discussed and the paper ends with the seemingly most effective model to date. Self-similarity is discussed in more detail as this seems to be a prevalent characteristic of internet traffic.

Data Collection

On the internet data is sent in manageable pieces called packets. Each packet has a TCP/IP header containing the following: the IP source and destination addresses, the protocol number (to identify TCP segments), the source and destination ports, the TCP flags, the data sequence number and the acknowledgement (ACK) number. These headers can be collected and analyzed to arrive at a model for internet traffic. [2] describes the data collection process in detail.

Internet Topology

The Internet is a network of routers, each belonging to a administrative authority or autonomous system (AS), connected by links. Both the router and the AS are referred to as nodes.

Let $P(k)$ be the connectivity distribution giving the probability that a node in the network is connected to k other nodes. Results on the World-Wide Web (WWW), the Internet and other large networks indicate that many systems belong to a class of inhomogeneous networks, called scale-free networks, for which $P(k)$ decays as a power-law, that is $P(k) \sim k^{-\alpha}$, free of a characteristic scale. [3]

Various mechanisms have been proposed as the origins of power law phenomena, including reflected or drifted multiplicative processes, exponentially stopped geometric Brownian motion, self-organized criticality, preferential growth, and highly optimized tolerance among others. [4]

[1] sheds more light on the topology and operation of the Internet than can be covered here.

Previously proposed network traffic models: their success and failure

Poisson (Classical) Model

First studies on data traffic indicated that the data traffic sources in communication networks were bursty in nature, i.e. relatively short sequence of source activities are followed by long idle periods. According to [5], studies suggested the following assumptions for data traffic are reasonable:

1. The interarrival times of messages generated by an external data source are exponentially distributed, i.e., each external data source behaves as a Poisson process. Let $G(i)$; $i = 1; 2; \dots; N$, be a random variable denoting interarrival times of messages generated from the i th data source.
2. The length of messages generated by an external data source are exponentially distributed. Let $H(i)$; $i = 1; 2; \dots; N$ be a random variable denoting length of messages from the i th data source.
3. Processes described by random variables $G(i)$ and $H(i)$ are stationary and independent.

Failure of Poisson model:

[6] reports that for both local-area and wide area network traffic, the distribution of packet inter-arrivals differs from exponential, that Poisson processes are valid only for modeling the arrival of user sessions; that they fail as accurate models for other wide-area network(WAN) arrival processes; that WAN packet arrival processes appear better modeled using self-similar processes; and that in some cases commonly used Poisson models underestimate the burstiness of TCP traffic over a wide range of time scales. [7] established that simple Poisson models were not appropriate for network traffic.

*Assisted by Prof. D. J. de Waal, UFS
rika_taute@yahoo.com

Packet Train Model

This model was a result of the discovery of long-range dependence and self-similarity in traffic data. [7] have modeled traffic between a single source and destination as an on/off or packet train process [8]. The packet train model assumes that a group of packets travel together. The packet train model is a source model. It applies only when we look at the packets coming or going to a single node. In order to allow analytical modeling with a simplified form of train model, usage of a two-state Markov model is suggested. The source can be either in generation state or idle state. The transitions between these states are memoryless (Markovian). The duration of the two states is exponentially distributed.

This models is a good application to the heavy tails observed in data by [9].

[10] reports that if the ON (OFF, or both) periods are generated according to heavy-tailed distributions, and the number of multiplexed flows is large, then the resulting aggregate traffic exhibits asymptotic self-similar properties with LRD behavior. An M/G/N queue with infinite variance service time exhibits long-range dependence (LRD) properties in the number of active servers. Since a heavy-tailed distribution of file sizes was measured on storage devices, it can be mapped onto the service time of a file transfer. [10]

Self-similarity

Leland et. al. (1995) argues in favor of self-similar models following findings that traffic exhibits ‘burstiness’ as opposed to a characteristic burst length which would tend to be smoothed by averaging over a long enough time scale. Self-similarity is the property we associate with fractals -the object appears the same regardless of the scale at which it is viewed. Since a self-similar process has observable bursts on all time scales, it exhibits long-range dependence; values at any instant are typically correlated with values at all future instants. (Crovella and Bestavros, 1994)(L. Muscariello, M. Mellia, M. Meo, M. Ajmone Marsan, R. Lo.)

According to Willinger; Taquq; Leland and Wilson, 1995, self-similar processes were introduced by Kolmogorov (1941). In terms of stochastic modeling, self-similar processes or their increment processes are almost exclusively used in situations where the modeler tries to account for the presence of long-term correlations in a parsimonious manner.(Willinger et al., 1995)

Willinger et al. reports that the data of Ethernet traffic measurements is clearly different than predictions by stochastic models considered up to 1995. This includes the Markov-modulated Poisson process, fluid-flow models, ARMA models, TES processes and packet-train models.

Definition of self-similar processes

Continuous self-similar process: A continuous time process $Y(t)$ is said to be exactly self-similar with self-similarity parameter H (the Hurst parameter) if it satisfies the following condition

$$Y(t) \stackrel{d}{=} a^{-H} Y(at), \forall t \geq 0, \forall a > 0, 0 < H < 1$$

where the equality is in the sense of finite dimensional distributions.

Self-similar time series: In the network traffic context, we normally deal with a time series rather than a continuous process. The definition of self-similarity in that context goes as follows. Let

$X = \{X(i), i \geq 1\}$ be a stationary sequence. Let

$$X^{(m)}(k) = \frac{1}{m} \sum_{i=(k-1)m+1}^{km} X(i), k = 1, 2, \dots$$

be the corresponding aggregated sequence with level of aggregation m obtained by averaging over non-overlapping blocks of size m . Then if X is self-similar, it has the same autocorrelation function, $r(k) = E[(X_t - \mu)(X_{t+k} - \mu)]$ as the series $X^{(m)}$ for all m . Note that this means that the series is distributionally self-similar: the distribution of the aggregated series is the same (except for changes in scale) as that of the original.

Second-order self-similar: A stationary sequence is said to be second-order self-similar if $m^{1-H} X^{(m)}$ has the same variance and auto-correlation as X for all m .

Asymptotically second-order self-similar: A stationary sequence is said to be asymptotically second-order self-similar if $m^{1-H} X^{(m)}$ has the same variance and auto-correlation as X as $m \rightarrow \infty$.

Asymptotically second-order self-similar processes are also called long range dependent (LRD) processes.

Gong, Liu, Misra, Towsley (2005) stresses that this is the property that network traffic exhibits, and ‘‘self-similar’’ is often used interchangeably with LRD processes which is misleading.

An equivalent way to describe asymptotic second-order self-similarity is in terms of it’s power spectral density ψ

$$\varphi_X(f) \sim k|f|^{1-2H}, f \rightarrow 0$$

These processes are also equivalently called 1/f processes. Gong, Liu, Misra, Towsley (2005) further explains that the Hurst parameter has been used as a parsimonious measure to describe the correlation structure of analyzed data, which is misleading since it must assume that the underlying process is self-similar. Thompson, Miller and Wilder (1997) has shown that internet traffic exhibits strong daily and weekly patterns. Thus, internet traffic is not truly self-similar or LRD, as for long enough timescales the LRD-like behavior disappears.

Properties of Self-Similar Process

(Willinger, Taquq, Leland and Wilson, 1995)

1. Long-range dependence and the Hurst-effect. Thus, processes with long-range dependence are characterized by an autocorrelation function that decays hyperbolically as the lag increases.
2. Slowly decaying variances.

3. Parsimonious modeling.

There are a number of approaches for mathematical modeling and analysis of self-similar processes. These include fractional Brownian motion, wavelet transform based models, physics based models and mathematical framework of scale stationarity.

Exactly self-similar models such as fractional Gaussian noise or asymptotically self-similar models such as fractional ARIMA processes can be used to fit hour-long traces of Ethernet traffic very well (Willinger et al 1995). Their successful application to packet traffic modeling relies on:

I. Having readily available statistical methods for real-time parameter estimation for these processes.

II. Being able to generate quickly long traces of synthetic observations from these processes.

III. Demonstrating network-related implications of self-similarity which, when not accounted for, lead to mediocre or unacceptable network performance.

[9] also reports a statistical analysis showing the burstiness and self-similar nature of data. In their paper they have shown that traffic due to World Wide Web transfers can be self-similar when demand is high. They have traced the genesis of Web traffic self-similarity along two threads: first, they have shown that transmission times are heavy tailed, primarily due to the very heavy-tailed distribution of Web files. Second, they have shown that silent times are heavy-tailed, primarily due to the influence of user "think time". In addition, they have shown the distribution of user requests is lighter-tailed than the set of available files; but that the action of caching serves to make the distribution of actual files transferred similar to the more heavy-tailed distribution of available files.

Long-Range Dependent Network Traffic

Let X_t be a stationary process with autocovariances $r(k) = \text{cov}(X_t, X_{t+k})$. Then X_t is said to be long-range dependent if, as $|k| \rightarrow \infty$,

$$r(k) \sim L_1(k) |k|^{2H-2}, H \in \left(\frac{1}{2}, 1\right)$$

Where $L_1(k)$ is a slowly varying function as $|k| \rightarrow \infty$, that is $L_1(ta)/L_1(t) \rightarrow 1$ as $t \rightarrow \infty$ for any $a > 0$. This implies that the correlations are not summable and the spectral density has a pole at zero. Under suitable conditions on $L_1(\cdot)$, the spectral density

$$f(x) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} r(k) e^{ikx} \sim L_2(x) |x|^{1-2H}$$

as $|x| \rightarrow 0$ for some $L_2(\cdot)$ slowly varying at the origin. The best models with above described properties are fractional Gaussian noise and fractional ARIMA. [11]

Among the numerous generic LRD models proposed in the literature, **Fractional Brownian Motion (FBM)** received a lot of attention, since its Gaussian nature helps in the study of the queuing behavior. Fractional Brownian motion is an example of a self-similar process.

The fractional Brownian motion is a stochastic process $V_H(t)$ in continuous time with the property that for any t_1 ,

$$\text{var}[V_H(t_2) - V_H(t_1)] \propto |t_2 - t_1|^{2H}$$

where $0 < H < 1$. Thus the increment process is normally distributed with mean zero and variance t^{2H} . H is the self-similarity parameter.

When $H = 1/2$, we get the standard Brownian motion: The increment $V_{1/2}(t_2) - V_{1/2}(t_1)$ has a variance proportional to the time difference $t_2 - t_1$. The fractional Brownian motion

$V_H(t)$ is a non-stationary process, but its increments form a stationary process. We see that the usual Central Limit Theorem does not apply for the sample mean of increments of $V_H(t)$ unless $H = 1/2$, in which case the increments are a (continuous time) white noise process.

The "derivative" of $V_H(t)$ (which can be heuristically defined, even though $V_H(t)$ is not actually differentiable anywhere), is called **fractional Gaussian noise**. The fractional Gaussian noise is stationary, and is essentially a continuous-time version of the Fractional ARIMA (0, d , 0).

The process $V_H(t)$ is said to be **statistically self-similar**, because if we change the time scale we find that, apart from a multiplicative constant, the stochastic behavior remains unchanged.

While this model has been successful in developing insights to the source of self-similarity in network traffic, it does not have predictive and recursive capability [12] and the model presents a restrictive correlation structure, that fails to capture the short-term correlation of real traffic and its rich scaling behavior. [10]

Fractional Gaussian noise: Fractional Gaussian noise (FGN) is the increment process of fBm.

The FGN processes represent the traditional LRD traffic model. Due to the huge link capacity of high-speed networks, many network applications are likely to be served by one multiplexer. By applying the central limit theorem, the aggregated process can be characterized as a Gaussian process. By the nature of Gaussian processes, the FGN is an exactly self-similar process. [13]

MWM (Multi-fractal Wavelet Model)

A relatively new approach to describing the dynamics of network traffic is based on multi-fractals. The investigation into the fractal nature of network traffic is based on using wavelet based analysis. [13]

The primary goal of the MWM is to model a class of non-negative discrete long-range dependent processes. The flexibility and accuracy of the model and fitting procedure results in a close fit to real data statistics. Multifractal models is an improvement on FBM due to their rich scale-invariance properties.

While the models mentioned so far give good approximations of the LRD properties of Internet traffic, they are difficult to manage, due to their analytical complexity.

Wavelet decomposition has been widely used as a natural approach to study scale invariance, but only recently they were introduced in the field of data networks. There are many examples of measurement-based traffic models, which try to fit the LRD properties of real traffic.

These models are computationally very efficient, but they are complex and difficult to tune, due to the lack of a mapping between the traffic parameters and the model coefficients.

FARIMA models are widely used in video trace modeling, and can be used to generate LRD sequences. These models are derived by filtering white Gaussian noise, and capture both the short and the long period correlations of traffic. However, the models are quite complex, and their structure makes it very hard to understand the relationship among the filter coefficients and the real traffic data. [10]

Markovian models for LRD traffic

[14] showed that the long-term correlation of traffic beyond a certain threshold does not influence the performance of a system, so that simple models where correlation is limited can be successfully employed.

[15] propose the use of discrete time Markov chains (DTMC) to modulate the packet arrival process. The key of the modeling process is the use of nearly decomposable chains, that have transients behaviors that are local to the initial state and evolves with a slower time constant between different groups of states, in other words a model that has some implicit form of memory.

The model is proved to emulate well both the correlation structure (Hurst parameter) and the queueing behavior of measured data.

Unfortunately, the parameters in this model is not easily related to traffic conditions. [10]

[4] present a model, that is Markovian in its components, and yet exhibits LRD-like correlation and spectral behavior. This follows after observed that there is little evidence that the tail impacts the design of various algorithms and infrastructure on the Internet. They hold that engineers should focus on the behavior of a “waist”, that refers to the portion for which there exist enough data to summarize the distributional information, not its “tail”. They further argue that the generation mechanism of traffic is inherently multiple timescale, and hence sustained correlations are inevitable, coupled with the fact that TCP generates sustained correlations on its own.

[10] proposed a simple **Markov Modulated Poisson Process (MMPP)** Internet traffic model that is capable of approximating well their traffic data characteristics.

The model is based on a layered structure of sessions that generate flows that finally generate packets. The characteristics of the synthetic traffic generated with the model match the LRD characteristics observed in the measured traces over the time scales of interest. One of the interesting features of the MMPP model is that it requires five parameters only. Three of these parameters can be

directly mapped onto average traffic parameters, such as the average flow arrival rate, the average number of packets per flow, and the average arrival rate of packets within flows. The other two parameters define the notion of session, and are used to control the Hurst parameter of the synthetic traffic on the considered scaling range.

The key features of the proposed MMPP model are its simplicity and its intuitive structure.

[16] compared the customized batch Markovian arrival process (BMAP) with the MMPP and the Poisson process by means of visual inspection of sample paths over four different time scales, by presenting important statistical properties, by formal analysis of traffic burstiness using R/S statistics, and by queuing system analysis. The challenge for employing BMAPs to model IP traffic constitutes the proper parameter estimation for this arrival process from the given trace data. [16] introduced an efficient and numerical stable method for estimating the parameters of a BMAP.

A timescale decomposition approach

[17] proposes a timescale decomposition approach to real time traffic prediction. The raw traffic data is first decomposed into multiple timescales using the *a trous* Haar wavelet transform. The wavelet coefficients and the scaling coefficients at each scale are predicted independently using the ARIMA model. The predicted wavelet coefficients and scaling coefficient are then combined to give the predicted traffic value.

This timescale decomposition approach can better capture the correlation structure of the traffic caused by different network mechanisms, which may not be obvious when examining the raw data directly. The proposed prediction algorithm is applied to real network traffic. It is shown that the proposed algorithm outperforms traffic prediction algorithms namely the fractional autoregressive integrated moving average (FARIMA), the neural network approach and method based on the α -stable model and gives more accurate results.

Traffic prediction based on the FARIMA model relies on the accurate estimation of the Hurst parameter. Despite a number of estimators reported in the literature, the accurate estimation of the Hurst parameter remains a difficult problem even in off-line conditions. The presence of non-stationarity and complex scaling behavior in network traffic makes the situation even worse. Therefore, traffic prediction based on the FARIMA model is not suitable for real applications.

Traffic prediction using the neural network approach can be quite complicated to implement. The accuracy and applicability of the neural network approach to traffic prediction is limited.

Finally, traffic prediction based on the α -stable model has the same problem as traffic prediction based on the FARIMA model, which relies on the accurate estimation of the Hurst parameter. Moreover, the α -stable model is based on a generalized central limit theorem and its application is limited by that. It might achieve a good performance in

heavy traffic or when there is a high level of traffic aggregations. However when conditions deviate from that, the performance may be poor. Furthermore, the α -stable model is a parsimonious model, which may not be able to capture the complex scaling behavior of the traffic.

Wavelet transform has been widely used in traffic analysis and modeling. In their paper, [17] use a special kind of redundant wavelet transform, i.e. the α -trous Haar wavelet transform, which is particularly suited for traffic prediction, and the ARIMA model for traffic prediction.

Kalman filtering and multivariate self-similarity

According to [12] in most of the studies on self-similarity, the self-similarity parameter is considered as the major characteristic of self-similarity, and its reliable estimation from the network measurements is an important task for both characterization and simulation of the traffic.

They proceed to define multivariate self-similarity as joint self-similarity, in which the self-similarity is governed by a matrix valued parameter H . This generalization could suit the nature of Multi-Input Multi-Output (MIMO) systems, since each channel is likely to be governed by a different self-similarity parameter. The system and measurement models for their, [12], proposed Kalman filter are defined as $tx(t) = A^t x(t) + B^t n(t)$ and $y(t) = Cx(t) + Dv(t)$, respectively. Here, the derivative operator $tx(t)$ indicates that the memory of the process is stored in time scales, unlike the memory stored in time shifts for stationary processes.

[18] also studied the formulation and analysis of multivariate self-similarity. He models self-similarity by using the long range dependent ordinary stationary processes and the correlation structure, but that does not provide a state space representation which can further be used in Kalman filtering formulation.

Heavy Tails

[9], [10] among others reported the observation of heavy tails in their data. If the asymptotic shape of the distribution is hyperbolic regardless of the behavior of the small values of the variable, it is heavy-tailed. A random variable that follows a heavy-tailed distribution can give rise to extremely large values with non-negligible probability. [9]

This should cause great concern for network engineers and could be planned for if the correct extreme value models are used although [4] claim that the tail has little impact on Internet performance.

In my opinion most of network problems experienced specifically in South Africa could be solved by application of extreme value theory. [19] also argued that a common explanation for observed LRD and self-similarity of network traffic, is heavy tailed transmission times and that empirical and theoretical evidence supports this explanation of self-similarity.

Classically these heavy tails has been modeled by the Pareto distribution. Log-normal distributions have also

been suggested and contrary to previous convictions, these are not contradictory to LRD in aggregated traffic [20].

Last mentioned authors discovered what they call "wobbles" of a type not present in the tails of classical Pareto and log-normal distributions. The wobbling effect is due to a mixture of distributions which will be investigated further through a simulation* study. These could be contributed to variability in the tail, but with repeating the analysis with different datasets, they conclude that the wobbles are important underlying distributional phenomena. They proceed to find the parameters for these mixed distributions. An interesting further study would be to confirm these "wobbles" with more recent data on the Internet and to check their parameter estimates.

Conclusion

In this article I have summarized progress made on the modeling and characterizing of internet traffic. Main developments included the repeated observation of self-similarity, long-range dependence and heavy tails. Models aimed at simulating these characteristics had to balance accuracy and tractability. Most fail in this regard. I aimed to stress the importance of considering extreme value modeling or tail behavior of aggregate traffic in the planning and expanding of networks. Interesting development in this area involves the modeling of file flows with mixed Pareto and log-normal models. I proposed further investigation in the application of this promising model.

References:

- [1] Cleveland, W.S. and Sun, D.X. "Internet Traffic Data". (2000). *Journal of the American Statistical Association*, Vol. 95, No. 451, pp. 979 – 995.
- [2] Smith, F.D. Campos, F.H. Jeffay, K. and Ott, D. "What TCP/IP Headers Can Tell us About the Web". (2001). *Proceeding of ACM SIGMETRICS 2001/Performance 2001, Cambridge, MA*, pp. 245-256.
- [3] Albert, R. Jeong, H. & Barabási, A.-L. "Error and attack tolerance of complex networks". *Department of Physics, 225 Nieuwland Science Hall, University of Notre Dame, Notre Dame, Indiana 46556, USA*
- [4] Gong, W.B. Liu, Y. Misra, V. Towsley, D. "Self-similarity and long range dependence on the internet: a second look at the evidence, origins and implications". (2005) *Computer Networks* 48, 377–399.
- [5] Babic, G., Vandalore, B., Jain R. "Analysis and Modeling of Traffic in Modern Data Communication Networks". Ohio State University Technical Report, February 1998.
- [6] Paxson, V. and Floyd, S. "Wide Area Traffic: The Failure of Poisson Modeling". (1995) *IEE/ACM TRANSACTIONS ON NETWORKING*, VOL.3, NO.3, pp. 226-244.
- [7] Willinger, W. Taqqu, M.S. Leland, W.E. Wilson, D.V. "Self-Similarity in High-Speed Packet Traffic:

* The results will be reported in the author's final thesis.

- Analysis and Modeling of Ethernet Traffic Measurements". (1995) *Statistical Science*, Vol. 10, No. 1 (Feb., 1995), 67-85.
- [8] Heath, D., Resnick, S., Samorodnitsky, G. "Patterns of Buffer Overflow in a Class of Queues with Long Memory in the Input Stream". (1997). *The Annals of Applied Probability*, Vol. 7, No. 4, pp. 1021-1057.
- [9] Crovella, M. E. and Bestavros, A. "Explaining World Wide Web Traffic Self-Similarity". (1995). Technical Report TR-95-015. Revised {crovella,best}@cs.bu.edu
- [10] Muscariello, L. Mellia, M. Meo, M. Ajmone Marsan, M. and Cigno, R. Lo. "Markov models of internet traffic and a new hierarchical MMPP model". (2004). *Computer Communications* 28 (2005) 1835-1851.
- [11] Beran, J., Terrin, N. "Testing for a change of the long-memory parameter ". (1996), *Biometrika*, 83, 3 pp. 627-638.
- [12] Yazıcı, B. Izzetoglu, M. Onaral, B. and Bilgutay, N. "Kalman filtering for self-similar processes". (2006). *Signal Processing* 86, 760-775
- [13] Nguyen, C. T. Sugawara, S. and Tetsuya. "A New Approach to Cell Loss Analysis for Long-Range Dependent Network Traffic". (2005), *MIKI IEICE TRANS. COMMUN.*, VOL.E88-B, NO.9
- [14] Grosslauser, M. and Bolot, J. "On the relevance of long-range dependencies in network traffic". (1999). *IEEE/ACM Transaction on Networking* 7 (5), pp. 629-640.
- [15] Robert, S. and Le Boudec, J.Y. "On a Markov modulated chain exhibiting self-similarities over finite timescale". *Performance Evaluation, Elsevier* 27-28 (1996) 159-173
- [16] Klemm, A. Lindemann, C. Lohmann, M. "Modeling IP traffic using the batch Markovian arrival process". (2003). *Performance Evaluation* 54 (2003) 149-173
- [17] Mao, G. "A Timescale Decomposition Approach to Network Traffic Prediction". (2005). *IEICE Trans. Commun.* Vol E88-B, No. 10, 3974-3981.
- [18] Bladt, M. "Multivariate self-similar processes: second-order theory". *J. Appl. Prob.* 31 (1994) 139-147.
- [19] Mikosch, T. Resnick, S. Hootzen, H. and Stegeman, A. (2002). "Is Network Traffic Approximated by Stable Levy Motion or Fractional Brownian Motion?" *The Annals of Applied Probability* Vol. 12, No. 1, 23-68.
- [20] Hernández-Campos, F. Marron, J.S. Samorodnitsky, G. and Smith, F.D. "Variable Heavy Traffic Durations in Internet Traffic". *To appear in the Journal of Performance Evaluation*.
- [21] Kolmogorov, A. N. "The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers". (1941) *C. R. Acad. Sci. URSS (N. S.)* 30 301-305. [Translation in *Turbulence* (S.K. Friedlander and L. Topper, eds.) 151-155 (1961). Interscience, New York.]
- [22] Thompson, K., Miller, G., Wilder, R. "Wide-area Internet traffic patterns and characteristics". (1997). *IEEE Network* 11 (6) 10-23.



Rika Taute was born and matriculated in Bloemfontein, Free State. She completed her B.Comm degree and honours in Actuarial Science at the University of the Free State. During her studies in actuarial science she developed a keen interest in extreme value theory and is currently employing this in further M-studies.