

An MMPP-Based Hierarchical Model of Internet Traffic

L.Muscariello, M.Mellia, M.Meo, M.Ajmone Marsan
Dipartimento di Elettronica, Politecnico di Torino, Italy

R.Lo Cigno
Dipartimento di Informatica e Telecomunicazioni, Università di Trento, Italy

Abstract— In this paper, we propose an MMPP (Markov Modulated Poisson Process) traffic model that accurately approximates the LRD (Long Range Dependence) characteristics of Internet traffic traces. Using the notion of sessions and flows, the proposed MMPP model mimics the real hierarchical behavior of the packet generation process by Internet users. Thanks to its hierarchical structure, the proposed model is both simple and intuitive: it allows the generation of traffic with desired characteristics by easily setting few input parameters which have a clear physical meaning. Results prove that the queuing behavior of the traffic generated by the MMPP model is coherent with the one produced by real traces collected at our institution edge router under different networking scenarios and loads. Due to its characteristics, the proposed MMPP traffic model can be used as a simple and manageable tool for IP network performance analysis, as well as for network planning and dimensioning.

I. INTRODUCTION

Identifying simple and accurate models of Internet traffic is not easy, as can be seen by scanning the vast literature on the subject (see Section II for a very short overview). A general consensus exists on the fact that Internet traffic is not Poisson at any level of aggregation, but researchers disagree on the approaches that should be used for the description of the traffic features. In this paper we propose an approach based on an MMPP (Markov Modulated Poisson Process) representation of traffic.

One of the main characteristics of Internet traffic, probably the one that impacts the most on network performance, planning, and dimensioning, is the Long Range Dependence (LRD) of the distribution of several parameters (e.g., packet inter-arrival time, amount of data transferred per time unit, etc.). A wide range of interpretations of the LRD of Internet traffic were proposed, and a large number of models for the representation of LRD traffic were devised, often quite different from one another, depending on the perspective of the work. Indeed, two main perspectives motivated the development of Internet traffic models: (i) trace fitting, and (ii) performance analysis. When the focus of the work is the fitting of real traffic traces, sophisticated stochastic processes are used to find the most elegant fitting, and to try to obtain a mathematical explanation of the observed LRD. On the contrary, when the focus of the work is on network performance analysis, with emphasis on queuing behavior and network planning and dimensioning, more “tractable” models are sought, that offer a phenomenological approximation of the traffic characteristics,

This work was supported by the Italian Ministry for Education, University and Research (MIUR) through the TANGO and GRID.IT projects.

not trying to explain *why* Internet traffic is LRD, but only to study *how* LRD impacts network performance. The work in this paper belongs to the second category. We present and discuss a model of Internet traffic based on an MMPP, that tries to mimic the hierarchical generation of data by Internet users, and is capable of effectively capturing the key aspects of the traffic measured on an edge router, hence representing an aggregation of the traffic generated by a number of sources. The proposed MMPP model is particularly appealing, because it is easy to understand, and it is controlled by a small number of parameters (just five), whose influence on the model output is predictable. The proposed MMPP traffic model was tuned on measured traffic traces, and its behavior was studied in queuing systems with finite buffers, representing multiple-node network topologies. A comparison of the queuing performance of the traffic produced by the MMPP Model and by traces measured on an edge router proves the accuracy of the traffic description.

II. RELATED WORKS

Most of the research on Internet traffic modeling stemmed from the seminal works [1], [2] showing that traffic traces captured on both LANs and WANs exhibit LRD properties. We mention here only works related to network performance evaluation, that share the same objective of our work. We thus ignore, for the sake of brevity, all the works focusing on the fitting of measured traces, without the goal of using the resulting models for network analysis and design.

One of the first attempts to describe LRD exploited *Fractional Brownian Motion* (FBM) models, whose Gaussian nature helps in the study of the queuing behavior [3]. However, FBM models present a restrictive correlation structure, that fails to capture the short-term correlation of real traffic and its rich scaling behavior.

Another approach to describe LRD consists in looking at packet traffic as a superposition of source-destination traffic flows, assuming that each single flow has an ON/OFF behavior [4], [5]. If the ON, OFF (or both) period durations are generated according to heavy-tailed distributions, then the resulting aggregate traffic exhibits asymptotic self-similar properties with LRD behavior.

An $M/G/\infty$ queue with service time with infinite variance is used in [6], [7] to model video sources. Moreover, since the heavy-tailed distribution of file sizes was measured on storage

devices [8], the $M/G/\infty$ with high variability services model is a popular model to generate traffic with LRD properties.

Wavelet analysis (see [9] for a review) was shown to be one of the most powerful methods for the description of stochastic properties of traffic. Several authors ([10], [11], [12]) used a wavelet approach to generate synthetic traffic, possibly taking explicitly into account a multi-fractal model as in [10]. The main advantage of these models, specially when multifractality is taken into account, is their rich scale-invariance property. These models are computationally very efficient, but they suffer for the lack of a simple mapping between the traffic parameters and the model coefficients.

All these traffic models deviate considerably from classical Markovian models, which, however, continue to be widely used for performance evaluation purposes [13], [14], [15], [16]. In these works, the Markov Modulated Poisson Process (MMPP) is considered as the best Markov process to emulate LRD [14] and scale invariance [13] (multifractality in particular), though in [15], [16] it is correctly pointed out that any MMPP cannot exhibit LRD in a mathematically proper way, i.e., it is always possible to find a time lag above which an MMPP correlation decays exponentially. The authors of these works thus define a “local Hurst parameter” using an approximate LRD definition, valid on a limited range of time scales only. In this work we will also apply such an approach. In [17], the authors show that the long-term correlation of traffic beyond a certain threshold does not influence the performance of a system: this result paves the way to the use of models where correlation is limited (such as MMPP models) to derive credible results in network performance studies. These papers, and [13], [14] in particular, are those that most influenced our work, though the approach we used to obtain the proposed MMPP model, and its final structure, differ considerably from the approaches in those papers, since the parameters of our model have a direct mapping on measurable traffic characteristics, as explained in Sect. III

III. THE TRAFFIC GENERATION MODEL

The main objectives in the model development were: i) simplicity, and ii) possibility of mapping the model parameters onto easily understandable traffic measures.

In order to achieve those goals, we decided to replicate the key aspects of the packet generation mechanisms activated by Internet users and applications. Packets are generated with a nested process, where *sessions* are started by an external process, *flows* are generated within sessions and *packets* are generated within flows.

Fig. 1 depicts this situation. The hierarchy in the packet generation process introduces memory in the system, regardless of the nature of the generation processes of sessions, of flows within sessions, and of packets during flows. However, the stochastic properties of the overall packet arrival process clearly depend on the underlying generation processes.

The model we propose mimics the hierarchical structure of flow and packet generation, but models neither the complex stochastic characteristics of the generation processes of flows and packets, nor the complex distributions of the number of flows within sessions and of the number of packets within




Fig. 1. Hierarchical representation of the packet generation process in the Internet.




Fig. 2. CTMC of the proposed model.

flows. The model assumes that flows within sessions arrive following a Poisson process, and also that packets within flows arrive following a Poisson process; in addition, the number of flows within sessions, and the number of packets within flows, is assumed to be geometrically distributed, so that the global model can be described by a Continuous-Time Markov Chain (CTMC) (see Fig. 2) and the generated traffic can be computed as a reward measure on the CTMC. Authors in [18] followed a similar hierarchical approach, but tried to fit the measured distributions for flows and packets arrivals, so that the overall model is not Markovian and is much more complex. The difference between our model and the model proposed in [14] lies instead in the hierarchical structure: The MMPP process in [14] does not have a pre-defined hierarchical structure, and LRD is obtained through a direct parameter fitting, which results in a very large number of parameters within the model.

Our model is completely described by five parameters:

- λ_s : the arrival rate of new sessions;
- λ_f : the flow arrival rate per active session;
- λ_p : the packet arrival rate per active flow;
- N_f : the average number of flows per session;
- N_p : the average number of packets per flow.

$\beta = 1 - 1/N_f$ is the probability that a flow is not the last of a session; $\mu_f = \lambda_p/(N_p - 1)$ is the average duration of a flow.

TABLE I
SUMMARY OF THE ANALYZED TRACES

Name	date	start time	stop time	packets (millions)	flows (thousands)
Peak'01	2 Feb 01	10:52	13:52	11	540
Night'01	2 Feb 01	04:52	07:52	0.43	30
Peak'00	13 Apr 00	08:10	14:10	12	564
Night'00	13 Apr 00	02:10	08:10	0.92	79

In all states, the global packet generation rate (the reward) is $n_f \lambda_p$, while the flow generation rate is $n_s \lambda_f + \lambda_s$.

The CTMC that modulates the packet and flow arrival processes is defined by the state variable $\bar{s} = (n_f, n_s)$, where n_f and n_s denote the number of active flows and the number of active sessions, respectively. The state-transition diagram of the CTMC model of the modulating process is reported in Fig. 2. Transitions can be classified in four groups:

- From state (i, j) to state $(i + 1, j + 1)$: arrival of a new flow which is the *first* flow of a session; the rate is λ_s ;
- From state (i, j) to state $(i + 1, j)$: arrival of a new flow which is neither the first nor the last flow of a session; the rate $(1 - \beta)\lambda_s + j\beta\lambda_f$;
- From state (i, j) to state $(i + 1, j - 1)$: arrival of a flow which is the *last* of a session; the rate is $j(1 - \beta)\lambda_f$
- From state (i, j) to state $(i - 1, j)$: termination of a flow which is *not* the last of a session; its rate is $i\mu_f$.

A. Tuning the Model

In order to tune and validate the model, we used traces collected at the access link of our institution (see [19]). We selected four traces with different characteristics, whose key parameters are summarized in Table I. The last two columns report the number of samples in each trace. We considered incoming streams of data, taking into account the fact that our campus LAN behaves mainly as a “client.” The outgoing streams are less interesting due to the very low link utilization.

The statistical analysis of the traces gives immediately three of the model parameters, N_p , λ_p , and λ_f , that are directly mapped on the measured quantities. λ_p is obtained as the ratio between the average number of packets per flow and the average flow duration. On the contrary, the parameters λ_s and N_f cannot be mapped directly on measurable quantities, because it is not easy to identify sessions within the data flow. We use these parameters to fit the correlation properties of the traffic.

The correlation analysis is based on the estimation of the Hurst parameter H through the wavelet decomposition of processes embedded in the traces [9]. Within the traces we identify the processes of event inter-arrival times $I(k)$, and event counting $N_T(n)$, taking both the flows and the packets as events. The counting process is obtained by counting the number of arrivals in a time interval $[nT, (n + 1)T)$. We use three values of $T \in \{1, 0.1, 0.01\}$ s, and name the three counting processes N_1 , N_2 and N_3 , respectively. This means that, within each trace, eight Hurst parameters are defined, four for packet-level processes and four for flow-level processes.




Fig. 3. Impact of N_f on the Hurst parameters generated by the model, for different values of the ratio $C = \lambda_s / \lambda_f$; (a) on the flow process; (b) on the packet process.

N_f is a measure of the memory of the system (the length of sessions), while the role of λ_s is less intuitive. Fig. 3, reports a sensitivity analysis for the Hurst parameter of the interarrival process for both flows and packets, for different values of the ratio $C = \lambda_s / \lambda_f$. As expected, increasing N_f monotonically increases the Hurst parameter at both the flow (H_f) and packet (H_p) levels¹. Notice that $N_f = 1 \rightsquigarrow H_f = 0.5$ by definition, while H_p is already larger than 0.5, due to the memory introduced by the flow length. On the other hand, C influences more H_f than H_p .

There is no direct mapping between λ_s , N_f and the Hurst parameters; thus, after setting N_p , λ_p and λ_f , we resort to the following iterative procedure to tune λ_s and N_f .

- 1) Let $C = \lambda_s / \lambda_f$ and set $N_f = 1$ and $C = 1$;
- 2) Generate a synthetic sequence with the same number of samples as the real trace;
- 3) Estimate the Hurst parameter of the synthetic trace at both packet and flow level, and compare them with those of the real trace;
- 4) If the fitting is good, exit, else assign new values to N_f and C and go to 2

The new values to be assigned to N_f and C at step 4 of the procedure are chosen following the empirical evidence that the larger N_f is, the larger H_p and H_f are; and also, the larger C is, the larger H_f , while C has little influence at packet level.

The selection of λ_f and N_f by means of the above fitting procedure requires only few iterations to provide accurate results (approximately 10 in our tests, to obtain a fit within 10% in the worst case). Table II reports the results obtained by fitting the four traces we selected;

Notice that all processes of all traces show similar values for the Hurst parameter ranging from 0.71 to 0.88. Indeed, \hat{H}_f and \hat{H}_p are almost independent from the considered trace or process, even if the operating conditions in the four selected traces are quite different (different link speeds, different loads,

¹The Hurst parameters of the traces are indicated as \hat{H}_f, \hat{H}_p respectively.

TABLE II

HURST PARAMETERS OF THE FOUR CONSIDERED TRACES (\hat{H}) AND OF THE FITTED SYNTHETIC TRACES (H)

Flow level analysis								
	Peak'01		Night'01		Peak'00		Night'00	
	\hat{H}_f	H_f	\hat{H}_f	H_f	\hat{H}_f	H_f	\hat{H}_f	H_f
I	0.74	0.74	0.86	0.84	0.76	0.75	0.74	0.73
N_1	0.76	0.71	0.76	0.82	0.75	0.75	0.78	0.77
N_2	0.75	0.71	0.73	0.83	0.74	0.72	0.76	0.76
N_3	0.74	0.78	0.80	0.79	0.75	0.73	0.78	0.79
Packet level analysis								
	Peak'01		Night'01		Peak'00		Night'00	
	\hat{H}_p	H_p	\hat{H}_p	H_p	\hat{H}_p	H_p	\hat{H}_p	H_p
I	0.87	0.84	0.71	0.82	0.84	0.85	0.84	0.83
N_1	0.88	0.87	0.73	0.87	0.86	0.86	0.83	0.83
N_2	0.88	0.82	0.72	0.79	0.87	0.86	0.83	0.82
N_3	0.88	0.84	0.76	0.86	0.88	0.85	0.87	0.88






Fig. 4. The third topology T3.

different patterns between Peak and Night). This can be taken as a strong indication that LRD is an intrinsic characteristic of the Internet traffic and is not induced by network conditions.

IV. PERFORMANCE EVALUATION

For the evaluation of the queuing performance of the traffic generated by our MMPP model, we consider three different network topologies. Topology T1 comprises a single link and a single source; topology T2 aggregates four different sources on a single link; topology T3, shown in Fig.4 is a tandem network with three links and four traffic relations, arranged so as to aggregate two traffic relations on every link, with relation Φ_1 crossing all three links, and all other relations interfering with it on a different link.

We test the queuing behavior by simulation using “*ns-2*” [20], comparing the behavior of the synthetic traces to that of the original measured traces. To avoid useless repetitions, we report the results for the Peak'01 trace only, since all other results are equivalent. When the measured traces are used, the service times at nodes depend on the measured packet length. The MMPP traffic model, instead, generates only arrival instants, and the packet length, determining the service time at the nodes, is randomly associated with packets. The packet length value is drawn from a distribution fitted on the measured one, which exhibits the well-known multi-mode behavior, with peaks for very short packets and for the different MTUs (Maximum Transfer Units) in the network. The peak of the Ethernet frame at 1,500 bytes dominates the lot. Notice that this procedure destroys all possible correlation between the arrival process and the packet length distribution.

Fig. 5. Buffer occupancy distribution, topology T1 with infinite buffer: (a) load $\rho = 0.9$ and (b) load $\rho = 0.8$.Fig. 6. Buffer occupancy distribution, $B = 32, 64, 128, 256, 512$ packets: (a) topology T1 and (b) topology T2.

To simplify reading the plots (unless otherwise stated), we normalize the link load, changing the link speed appropriately. Results at different load levels or with variable loads on different links in topology T3 yield similar results. Fig.5 reports the queue length distribution for infinite buffer in topology T1. The thin dashed line refers to the measured trace, the solid line to the MMPP model. The thick dashed line refers to a plain Poisson arrival process (the packet length distribution is still the measured one) and is reported for reference. The left plot refers to load 0.9; the right plot to load 0.8; results show that the load does not influence the MMPP model behavior. The buffering behavior of the model is satisfactory, and the reference Poisson curve shows that the packet length distribution has a marginal impact on the queuing behavior.

Queueing behavior is much more interesting (and realistic) in the finite buffer case. We now consider the three topologies with five different values of buffer size: $B = 32, 64, 128, 256,$ and 512 packets. Fig. 6 refers to the T1 and T2 topologies. The matching is satisfactory, though the trace-driven simulations always show a higher probability of full buffer, that is reflected in a slightly higher loss rate. The Poisson driven simulations refer only to $B = 32$, and show how inaccurate is a simple




Fig. 7. Buffer occupancy distribution, topology T3, $B = 32, 64, 128, 256, 512$ packets; the load is normalized on each link; (a) first queue, (b) second queue, (c) third queue.

Poisson model, even for a very short buffer size. All other curves for Poisson arrivals are practically coincident with this one, and are not reported to avoid cluttering the graph. Multiplexing traffic does not change the situation, and actually the accuracy of the results obtained with the MMPP traffic model is even better. Fig. 7, finally, refers to the multi-node T3 topology, and reports the behavior of all three buffers. Once again, the MMPP model behavior is very close to that of traces, and we can observe that traffic crossing several bottleneck (as the traffic relation Φ_1 does) does not alter results significantly.

V. CONCLUSIONS

In this paper we presented a simple hierarchical MMPP traffic model capable of generating traffic that accurately emulates the aggregate Internet traffic measured at an edge router. The model is based on a layered structure of sessions, that generate flows, that finally generate packets. The properties of the model, with respect to the emulation of the LRD traffic characteristics, and to the queuing behavior in several networking scenarios, were analyzed in detail, comparing synthetic traffic with real measured traces. The comparison shows that the MMPP traffic model emulates real traffic quite accurately. The simplicity of the MMPP traffic model makes it an ideal traffic generator to drive simulations (the code for *ns-2* is available for download at [21]) aimed at planning and dimensioning links and buffers. In addition, the Markovian properties of the MMPP traffic model make it analytically tractable, so that analytic solutions for simplified networking scenarios are possible (e.g., deterministic or exponential service times). Our work proceeds in this direction.

REFERENCES

- [1] W.E. Leland, M.S. Taqqu, W. Willinger, V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended version)," *IEEE/ACM Trans. on Networking*, Vol. 2, N. 1, pp. 1–15, Jan. 1994.
- [2] V. Paxson, S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Trans. on Networking*, Vol. 3, N. 3, pp. 226–244, Jun. 1995.
- [3] I. Norros, "On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks," *IEEE JSAC*, Vol. 13, N. 6, pp. 953–962, Aug. 1995.
- [4] M.E. Crovella, A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Trans. on Networking*, Vol.5, N.6, pp. 835–846, Dec. 1997.
- [5] W. Willinger, M.S. Taqqu, R. Sherman, D.V. Wilson, "Self-Similarity Through High Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *IEEE/ACM Trans. on Networking*, Vol.5, N.1, pp. 71–86, Jan. 1997.
- [6] M. Krunz, A. Makowski, "A Source Model for VBR Video Traffic Based on M/G/infinity Input Processes," *Proc. of IEEE Infocom'98*, San Francisco, CA, USA, pp. 1441–1448, Apr. 1998.
- [7] M. Zukerman, T. Neame, R.G. Addie, "Internet Traffic Modeling and Future Technology Implications," *Proc. of IEEE Infocom'03*, S. Francisco, CA, USA, Apr. 1–4, 2003.
- [8] M. E. Crovella, M. S. Taqqu, A. Bestavros, "Heavy-Tailed Probability Distributions in the World Wide Web," In *A Practical Guide To Heavy Tails*, Chapman & Hall, New York, pp. 3–26, 1998.
- [9] P. Abry, P. Flandrin, M.S. Taqqu, D. Veitch, "Self-Similarity and Long-Range Dependence Through the Wavelet Lens," in *Theory and Application of Long-Range Dependence*, (Doukhan, et al. editors), Birkhuser, Boston, USA, 2002, ISBN 0-8176-4168-9.
- [10] H. Riedi, M.S. Crouse, V.J. Ribeiro, R.G. Baraniuk, "A multifractal Wavelet Model with application to data traffic," *IEEE Trans. on Inf. Theory*, Vol. 45, No. 4, pp. 992–1018, Apr. 1999.
- [11] S. Ma, C. Ji, "Modeling Heterogeneous Network Traffic in Wavelet Domain", *IEEE/ACM Trans. on Networking*, Vol. 9, No. 5, pp. 634–649, Oct. 2001.
- [12] A. De Vendiectis, A. Baiocchi, "Wavelet Based Synthetic Generation of Internet Packet Delays," *Proc. of ITC17*, Dec. 2–7 2001.
- [13] A. Horváth, M. Telek, "A Markovian Point Process Exhibiting Multifractal Behavior and its Application to Traffic Modeling," *Proc. of 4-th Int. Conf. on Matrix-Analytic Methods in Stochastic Models (MAM4)*, Adelaide, AU, Jul. 14–18, 2002.
- [14] A. T. Andersen, B. F. Nielsen, "A Markovian Approach for Modeling Packet Traffic with Long-Range Dependence," *IEEE JSAC*, Vol.16, No.5, pp. 719–732, 1998.
- [15] S. Robert and J.Y. Le Boudec, "On a Markov Modulated Chain Exhibiting Self-Similarities Over Finite Timescale," *Performance Evaluation*, Elsevier, Vol. 27–28, pp. 159–173, Oct. 1996.
- [16] S. Robert, J.Y. Le Boudec, "New Models for Pseudo Self-Similar Traffic," *Performance Evaluation*, Vol. 30, No. 1–2, pp. 57–68, Jul. 1997.
- [17] M. Grosslauser, J. Bolot, "On the Relevance of Long-Range Dependencies in Network Traffic," *IEEE/ACM Trans. on Networking*, Vol. 7, N. 5, pp. 629–640, Oct. 1999.
- [18] A. Reyes Lecuona, E. González Parada, E. Casilari, J.C. Casasola and A. Díaz Estrella, *A Page-Oriented WWW Traffic Model for Wireless System Simulations*, ITC16, Edinburgh, Jun. 1999
- [19] M. Mellia, A. Carpani, R. Lo Cigno, "Measuring TCP and IP Behaviour on a Edge Node," *Proc of IEEE Globecom 2002*, Taipei, Nov. 2002.
- [20] "ns - network simulator". <http://www.isi.edu/nsnam/ns/>
- [21] "mmpp traffic generator for ns". http://www.tlc-networks.polito.it/muscariello/mmpp_tg/