

ANALYZING PACKET INTERARRIVAL TIMES DISTRIBUTION TO DETECT NETWORK BOTTLENECKS

Pál Varga

*Department of Telecommunications and Media Informatics,
Budapest University of Technology and Economics (BME-TMIT)
Address: Magyar tudósok körútja 2., Budapest, Hungary, H-1117.
Phone: +36-1-463-3424.*

pvara@tmit.bme.hu

Abstract This paper analyzes the properties of packet interarrival time (PIT) distribution functions of network segments including bottlenecks. In order to show the correlation between bottleneck behavior and packet interarrival time distribution, the alteration of probability distribution function (PDF) is observed through simulations including tighter and tighter bottleneck connections. The process of network bottleneck detection by passive monitoring requires effective metrics for distinguishing seriously congested links from normal or underutilized connections. The paper evaluates the third and fourth central moments (skewness and kurtosis, respectively) of PIT distribution as possible metrics for bottleneck detection. Simulation results as well as real measurement data analysis showed that PIT kurtosis can be a powerful measure of bottleneck behavior.

Keywords: passive monitoring, bottleneck detection, kurtosis

1. Introduction

Passive monitoring-based bottleneck detection requires a complex collecting, pre-processing and evaluation system. Although the collection and pre-processing of the large amount of data are demanding tasks by themselves, it is also challenging to find appropriate measures for evaluating whether a link is bottleneck or not. This paper introduces a novel metric, which seems to be powerful enough to distinguish traffic carried over a network bottleneck from "traffic flowing with its own pace".

There are various measures suggested by [Varga et al., 2003] and [Moldován et al., 2004], each derived from transport-level flow analysis. These studies investigate several metrics (loss-rate, speed-averages, variance of flow-level throughput, delay-factor) to be used in bottleneck detection by passive moni-

toring, and found that they work with different accuracy under different conditions. The two most promising ones are the "X-measure" [Moldován et al., 2004] (like peakedness [Molnár and Miklós, 1998], it is a *coefficient of variation*-type metric applied for the throughput) and the delay factor calculated from the inter-arrival times of flows (based on the M/G/R – PS model) [Riedl et al., 2000; Varga et al., 2003]. Nevertheless, a recent study [Varga et al., 2004] evaluating these metrics on live traffic showed that they are not yet accurate enough and need further fine-tuning.

Packet pair method [Keshav, 1991; Kang et al., 2004] is widely used for bottleneck bandwidth estimation. This technique includes, however, intrusion of active probe traffic to the network before analyzing the results [Carter and Crovella, 1996]. The current paper focuses on non-intrusive bottleneck detection methods only. The entropy-based clustering method (introduced in [Katabi and Blake, 2002]) seems to be powerful to find connections sharing a bottleneck; the referred paper, however, misses to provide a range for packet interarrival time entropy that suggests bottleneck behavior.

The current paper introduces a novel metric to be used in bottleneck detection. It is based on the analysis of the probability distribution function (PDF) of packet interarrival times, which exhibit different shapes (with some common patterns) as seen during studying network bottlenecks using transport-level flow interarrival time distributions.

2. Packet Interarrival Times Distribution

There is a good visual interpretation of packet interarrival times PDF patterns in [Katabi and Blake, 2002]. The authors describe PDFs computed in scenarios with no experienced queuing, significant queuing and queued traffic influenced by cross-traffic. Spikes and "spike-trains" in the PDF are found to suggest bottlenecks – or at least significant queuing in the analyzed path.

Nevertheless, there is an important difference in naming conventions used in [Katabi and Blake, 2002] and in this paper. The referred study considered connections as "bottlenecks" where the packets experienced "significant queuing". This is a very loose definition, even though it is hard to give a firm description of network bottleneck (other definitions would be connected with loss, throughput limits, high utilization, significant delay [Varga et al., 2003]).

The current paper considers a link as "bottleneck" where packets experience continuous, severe queuing and even being dropped due to the finite queue-lengths. Let us take an other viewpoint: consider a user utilizing similar networked services on server a and on server b (the servers offer similar processing performance). The user can reach server a on route A , whereas he/she gets serviced by server b on route B . In case this user is satisfied with the network performance towards server a , but he/she can notice performance problems

towards server b , then route B contains bottleneck link(s) – at least more of them, than scenario A does.

Obviously this is not a precise definition of a bottleneck either. Finding an appropriate metric to distinguish bottleneck links from well-dimensioned ones could help clarifying the issues of having different definitions for the same underlying problem (which is ultimately reflected in user satisfaction of using networked services).

2.1 Higher Order Statistical Properties

The first and second central moment (mean and variance) of statistical distributions are widely used for briefly characterizing a distribution. Higher order statistical properties [Kenney and Keeping, 1951], as the third central moment (skewness) and the fourth central moment (kurtosis) are more rarely used in the engineering practice (although their applicability is wide-scale).

Skewness characterizes the degree of symmetry – or rather, the asymmetry – of a distribution around its mean. Positive skewness indicates a distribution with a probability-peak on the lower values and an extending tail towards the higher values. On the contrary, negative skewness indicates a distribution having a probability-peak on the higher values and an asymmetric tail extending towards the lower values.

Equation 1 shows the definition of skewness.

$$\gamma_1 = \frac{E[(\xi - E(\xi))^3]}{\sigma^3}, \quad (1)$$

where $E()$ stands for expectation, ξ is the statistical variable (hence $E(\xi)$ is the mean of ξ) and σ is the standard deviation. For measured data with finite number of measured entities, estimated skewness can be calculated as

$$\gamma_1 = \frac{n}{(n-1)(n-2)} \sum_{i=1}^n \left(\frac{x_i - x_{mean}}{s^*} \right)^3, \quad (2)$$

where n is the number of entities, x_i is the value of the actual item, x_{mean} is the mean of the measured values of x and s^* is the standard deviation.

Kurtosis characterizes the relative peakedness or flatness of a distribution compared to the normal distribution. The distribution is leptokurtic (or more peaked than the standard normal distribution) if the kurtosis excess (see Equation 3) is positive. Negative kurtosis excess indicates a platykurtic (or relatively flat) distribution. The term "kurtosis" was first used in [Pearson, 1905].

Equation 3 shows the definition of "kurtosis excess", which is widely used in the practice of mathematical statistics. The outcome of this type of kurtosis is normalized for easier comparison with the normal distribution. "Kurtosis

proper" is by definition the fourth central moment, and it misses the normalizing element of -3 .

$$\gamma_2 = \frac{E[(\xi - E(\xi))^4]}{\sigma^4} - 3. \quad (3)$$

The statistical estimate of kurtosis with finite number of measured entities comes from the formula

$$\gamma_2 = \frac{n(n-1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^n \left(\frac{x_i - x_{mean}}{s^*} \right)^4 - 3 \frac{(n-1)^2}{(n-2)(n-3)}. \quad (4)$$

The next chapter introduces the features of PDFs derived from packet interarrival times of computer networks. We shall see that deriving skewness and kurtosis helps distinguishing bottleneck scenarios from un-congested measurement setup.

2.2 Packet Interarrival Times PDF of Links with No Congestion

The probability distribution function (PDF) of packet interarrival times (PIT) should be extremely flat in a non-queued aggregated network-link. This is because several links with various capacities can carry packets to the observed aggregated link. Independent sources generating traffic in such topology causes absolutely random interarrival times.

In case a node aggregates numerous links with various capacities (and no queuing), the PIT PDF at that node appears to be flat, as shown in Figure 1.a (showing probability and PIT (in microseconds) on the axes).

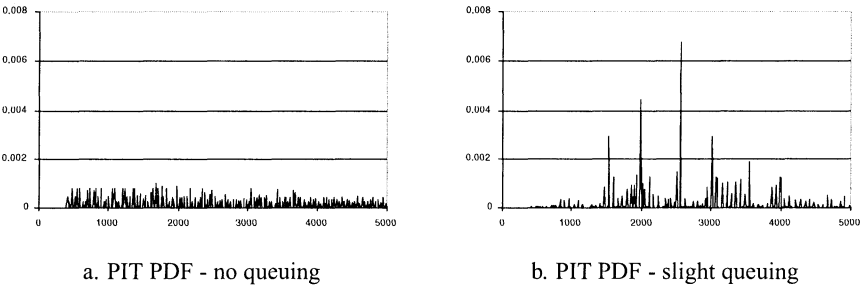


Figure 1. Packet interarrival times PDF on an aggregated link with no queuing (left) and on an aggregated link, with eventual queuing (right)

Once the link gets more busy, the relevant network node must queue some packets, and place them on the line right after the previous packet (back-to-back). The more of this queuing is applied, the less "flat" the PDF becomes: spikes starting to appear at the interarrival times where queued packets has followed each other back-to-back. Figure 1.b depicts a scenario where several lower capacity links – having different peak rates – connect to an aggregation node. The packets have experienced some – not severe – queuing before arriving to the aggregated link. Skewness of such PIT distributions are close to zero, or negative, whereas their kurtosis is negative, emphasizing that these PDFs are relatively flat.

2.3 Effects of the Bottleneck Behavior on Interarrival Times Distribution

Theoretically what one should expect during the observation of a link getting congested is the following. As the observed link starts showing bottleneck behavior, most of the eventual spikes of the PDF gets less noticeable and the spike around the lowest possible interarrival times starts dominating the PDF. Under severe congestion this spike fully dominates the PDF, as packets arrive back-to-back during the whole measurement period. Both skewness and kurtosis exhibit more positive values. The more dominant the spike around the lower PIT values gets, the more skewed the PDF becomes. Similarly, as kurtosis is, as the definition suggests the "peakedness" or "spikeness" of the distribution [Darlington, 1970], PIT kurtosis assumes higher values as the spike dominates the PDF. Since eventual queuing already makes the PDF skewed, one can expect that kurtosis should be more robust metric of bottleneck behavior than skewness. The following sections provide simulation results and analysis of real-life experiments to detail and support the above theory.

3. Bottleneck Detection: PIT PDF, Skewness and Kurtosis

3.1 Simulation Environment

The network topology used during the simulation is shown by Figure 2. OPNET was chosen as a simulation tool, as previous work on bottleneck detection has proven its applicability [Varga et al., 2003]. During each simulation period, traffic was generated to traverse the network for 15 minute long measurements.

From the simulation's point of view the traffic sources were the ISP's (aggregation nodes of Internet Service Providers) The characteristics of the traffic matches service types such as e-mail, ftp, http and database-access. There is also some asymmetry in the link capacities, since one of the ISP's is connected to the backbone through a bottleneck link (*link_2* between routers R2 and R5 in Figure 2).

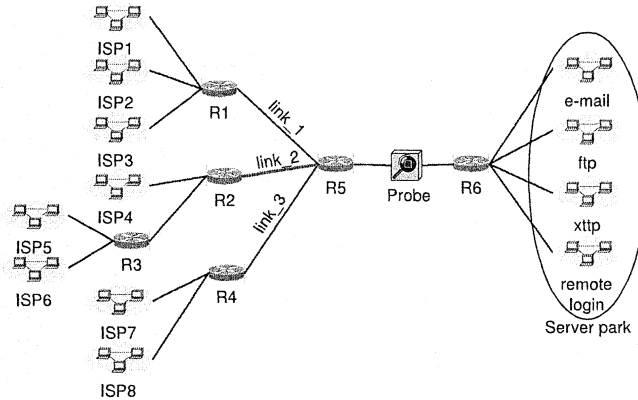


Figure 2. Network topology used during the simulations

The applied data collection and evaluation steps are as follows.

- the Probe captures traffic on an aggregated link (between R5 and R6 in Figure 2),
- traffic flowing from the same directions are distinguished by IP-address ranges belonging to the source ISP's of a given direction,
- after separating the captured traffic by directions, the PIT-characteristics of each direction is analyzed.

To evaluate kurtosis and skewness as possible metrics for bottleneck detection, the bottleneck in our simulation environment has been set tighter in several steps. This way one can evaluate how much difference in utilization and also in available bandwidth (ABW) would result in a loose or tight bottleneck. The bottleneck was created in the simulation topology by defining *link_2* having relatively low bandwidth, whereas *link_1* and *link_3* were equal, higher capacity connections.


The bottleneck was set tighter and tighter by increasing the maximum ABW of *link_1* and *link_3* (up to 1000 Mbps) and simultaneously decreasing the capacity of *link_2* (down to 300 Mbps). The tightest bottleneck was set to be able to handle merely 30% of the traffic of the other links (loaded by the same amount of traffic). In the following this scenario will be referred to as "1000/300", suggesting the "higher/lower"-bandwidth values applied for "*link_1, 3/link_2*".

3.2 Common patterns of PIT PDFs


In order to validate the theory about PIT PDF structures (see Section 2.3) an exhaustive study was carried out using different source traffic characteristics,

topologies and utilizations (the latter has ranged from scenario 675/625 to scenario 1000/300). Typical results for the two extreme scenarios are plotted in Figures 3.a-f.


The somewhat normal condition can be described as "packets experience queuing, but no loss" (scenario 700/600). The other extreme is a real bottleneck condition (scenario 1000/300).




a. *link_1*, ABW = 700




d. *link_1*, ABW = 1000




b. *link_2*, ABW = 600



e. *link_2*, ABW = 300



c. *link_3*, ABW = 700



f. *link_3*, ABW = 1000

Figure 3. Packet interarrival time PDFs in simulation scenarios 700/600 (left) and 1000/300 (right)

Comparing the resulting PDFs, there are numerous observations that can be made:

- the PDF keeps its general shape (number and place of the spikes) in a given link, as long as the same type of traffic flows through it,

- great spikes appearing in the PDF do not necessarily hide bottleneck condition,
- if the highest spike is not positioned at the lower PIT values, the packets are not following each other back-to-back (hence there is no severe congestion, see Figures 3.c,f),
- the PDF loses some of its peakedness if the available bandwidth increases (compare Figures 3.a and d),
- as expected, the PDF gets extremely peaked and skewed under bottleneck conditions (see Figure 3.e),
- as expected, the spike at the lower PIT values are dominant under bottleneck conditions (see Figure 3.e).

Unfortunately, drawing conclusions from the comparison of plotted diagrams is not a feasible network maintenance practice. Even in a medium-sized network the high number of network connections makes such comparisons impossible. A more decent metric is needed that is able to distinguish PDFs of congested connections from other cases without significant queuing. As anticipated, skewness and kurtosis of PITs should be able to suggest whether a severe bottleneck exists or not. The following section evaluates the usability of these higher order statistical properties in bottleneck detection.

3.3 Performance Evaluation Based on Simulations

The statistical properties skewness and kurtosis of PITs were calculated using Equation 2 and Equation 3 respectively. During the simulations the same network topology (see Figure 2.) were loaded by similar source traffic. The only major difference between the simulation scenarios is the ABW assigned to links 1, 2 and 3 in Figure 2.

Table 1. provides skewness values of the three links for the different bottleneck scenarios. The following observations can be made by evaluating the results:

- as expected, skewness of a bottleneck link is positive,
- the tighter the bottleneck is, the more positive skewness is observed,
- skewness assumes positive values also for links with eventual queuing, hence it is difficult to distinguish these from severely congested links,
- connections with less traffic and almost no queuing can be characterized with negative or close to zero skewness value.

Table 1. PIT skewness values in simulated environment

Name	675/625	700/600	800/500	900/400	1000/300
<i>link_1</i>	0.75527	0.80733	0.66632	0.57561	0.41202
<i>link_2</i>	0.95506	1.05409	1.10767	1.48124	1.65051
<i>link_3</i>	-0.09488	-0.09007	-0.11049	-0.04702	-0.03974

To conclude, skewness of PITs noticeably acquires positive values in bottleneck scenarios, although it may do so for links with eventual queuing.

Table 2. summarizes PIT kurtosis values observed in different simulation scenarios. There are merely two positive values appearing among the results, both observed on the bottleneck link in serious bottleneck scenarios (900/400 and 1000/300). It is also noticeable that the PIT distribution gets more platykurtic (flat) as ABW increases (*link_1*) and changes its platykurtic character into leptokurtic (peaked) as ABW decreases as far as causing bottleneck behavior (*link_2*). There is a less practical, but interesting feature of PIT kurtosis, namely, that it does not assume noticeably lower (more negative) values if the capacity of an underutilized link increases (*link_3*). To summarize,

Table 2. PIT kurtosis values in simulated environment

Name	675/625	700/600	800/500	900/400	1000/300
<i>link_1</i>	-0.96502	-0.87263	-1.04085	-1.23243	-1.45336
<i>link_2</i>	-0.59994	-0.35450	-0.28027	0.71819	1.31911
<i>link_3</i>	-1.62645	-1.63841	-1.63470	-1.63417	-1.61983

PIT kurtosis seems to be an appropriate metric for detecting bottlenecks. This is further validated in the following section.

3.4 Using PIT Kurtosis to Detect Bottlenecks in Real, Operational Networks

In order to validate metrics for passive bottleneck detection, a series of measurements have been taken place at sites of a major Hungarian network operator. During the measurements a passive network monitoring tool (Network Associates' Sniffer and Snifferbook Ultra) have been used, capturing continuous data traversed on Gigabit Ethernet interfaces. The measurements covered normal and busy hour conditions, connections to the operator's Internet Data Center and to routers/switches at the edge of the core network.

The measurements captured traffic during severe bottleneck conditions also. The topology of the network segment that contained a bottleneck link during the measurements can be studied in Figure 4. Traffic is flowing from the Dig-

ital Subscriber Line Access Multiplexers (DSLAM) towards the core network (and back), going through internal ATM (STM-1) links. The network monitoring tool has been connected to the aggregated link, and captured all packet headers of the measurement period. As part of the data-processing, packet-headers were sorted by direction (ATM sub-network). The moments of packet interarrival time distribution was calculated for each direction. Most of the measurements have been carried out under normal network conditions. During one of the monitoring sessions, an ATM link loading the *OpR_1 - OpSW* segment got overloaded. (this fact has been indicated by other network analysis tools also).




Figure 4. Measurement topology at the operator's site

Skewness analysis resulted positive values for all cases. The bottleneck link provided higher PIT skewness value comparing to the others, but still the difference was not significant, leaving skewness to be unable to detect bottleneck by itself.

The results of PIT kurtosis calculations are demonstrated in Table 3. These kurtosis values clearly meet the expectations set up in Section 3.3. Traffic measured under normal network conditions (underutilized links) provide negative values; although *OpR_1 - OpSW* shows kurtosis close to zero, suggesting some uncertainty. During the congestion on this route (remember that not the monitored link, but an ATM link, being two hops behind the monitoring unit has been congested), kurtosis reached the positive domain, suggesting severe bottleneck condition. In this period noticeable packet loss and over 90% link utilization was observed also (using monitoring tools covering the targeted, lower capacity links).

The PIT PDFs calculated for the three connected links under normal and high traffic conditions are shown in Figure 5.a-c and Figure 5.d-f respectively. After analyzing these PDFs based on the assumptions of Section 3.2, the following observations can be made. Figures 5.a, c and d suggest healthy traffic

Table 3. PIT kurtosis values of measurements at the operator’s site

Link name	normal conditions	OpR_1 overloaded
OpR_0 - OpSW	-0.46420	-0.45571
OpR_1 - OpSW	-0.06267	1.15119
OpR_2 - OpSW	-0.11180	-0.13289

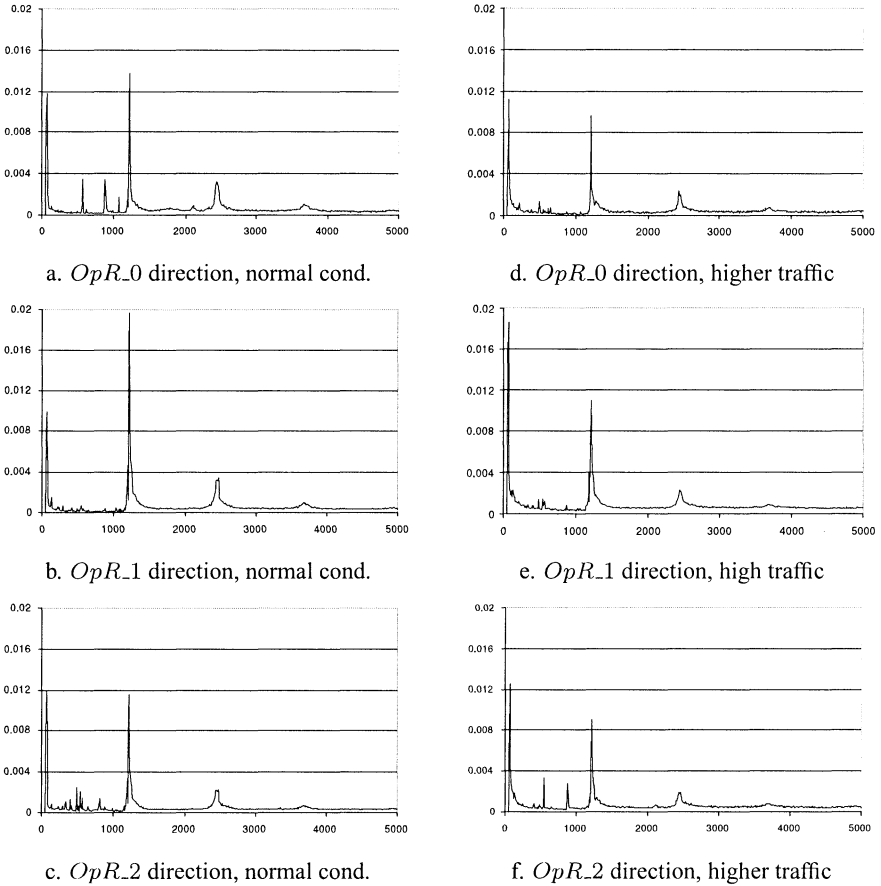


Figure 5. Packet interarrival time PDFs calculated based on real network data

flow, with a normal amount of queuing considering the high aggregation level. Figure 5.b includes a noticeable spike, however it is found at the higher PIT values, which means there were no severe congestion. The visual comparison of Figures 5.e and f does not reveal any major difference, as in both figures

the PDF spikes at the lowest possible PITs, although the spike is significantly higher at the real bottleneck case (Figure 5.e).

In fact, kurtosis – as a metric for bottleneck behavior – has clearly distinguished these situations as well. It provided the positive result of 1.15119 for the bottleneck case visualized by Figure 5.e as opposed to the otherwise normal traffic condition shown in Figure 5.f, where kurtosis was calculated to be -0.13289 (see Table 3).

During the various simulations and measurements it was noticed that PIT kurtosis calculated on the full PIT distribution depends slightly on link capacity. When two links have the same utilization (in percentage), the link having the higher capacity appear to have higher kurtosis. The reason behind this observation is not clearly identified yet. The consequences of this behavior, however, can lead far. In an extreme scenario the kurtosis of a *higher capacity, less utilized* connection can be higher than the kurtosis of a *lower capacity, more utilized* link. This would weaken the accuracy of PIT kurtosis as a bottleneck metric.

To overcome the above obstacle, a correction of PIT kurtosis should be considered for bottleneck detection. This correction should be linked to the analyzed data volume somehow.

Current analysis shows that PIT kurtosis scales well and provides more accurate, less capacity-dependent results when the it is calculated on a subset of the PIT-distribution. Leaving the 10 percent tail out of the analysis and calculating kurtosis up to the 90 percentile point of the PIT-distribution has provided satisfactory results on the available data sets. Future work should verify this observation, and clarify the issue of the slight capacity-dependency of PIT kurtosis.

4. Conclusions

Bottleneck detection based on passive measurements can be supported by analyzing Packet interarrival time (PIT) distribution. The more packets experience queuing at a network node, the more of them leave the node back-to-back. Under normal conditions the PIT probability distribution function (PDF) is relatively "flat": spikes due to typical packet lengths and minimal following times are visible, but appear to be small. As queuing turns into severe congestion, the spike around the lowest possible PIT value gets dominant. This fact is indicated by the fourth central moment – kurtosis – also: it gets more positive. Beside this, the third central moment, skewness gets more and more positive, also. Kurtosis of distributions more flat than normal distribution have negative kurtosis, whereas peaked distributions characterize themselves with positive kurtosis. Applying this to packet interarrival times distributions, positive kurtosis suggests serious bottleneck behavior, while negative kurtosis is

a property of underutilized links. Values being very close to zero is hard to evaluate, but probably hide serious queuing in the path.

Both skewness and kurtosis appeared to acquire higher values as the available bandwidth of a link decreased in the OPNET-based simulation environment. While both measures performed well as relative metrics, only *kurtosis* is powerful enough to distinguish bottleneck links from underutilized connections. Current studies show that kurtosis is slightly dependent on capacity, hence the metric should be refined for more accurate bottleneck detection: kurtosis should be calculated up to the 90 percentile of the PIT distribution.

Analysis of real measurement data has also been carried out, supporting that PIT kurtosis can be a powerful metric of detecting bottlenecks. The usability of PIT skewness and kurtosis in network performance analysis is for further study. The idea of PIT skewness and kurtosis to be used for bottleneck detection has been first submitted as an article for EUNICE 2005. The generalized idea and the more detailed methodology of bottleneck detection based on passive measurement at an aggregated link has then appeared later, at [Varga and Kún, 2005].

References

- Carter, Robert L. and Crovella, M. E. (1996). Measuring bottleneck link speed in packet-switched networks. *Performance Evaluation*, 27-28:297–318.
- Darlington, R.B. (1970). Is kurtosis really peakedness? *American Statistician*, 24(19-22).
- Kang, S., Liu, X., Dai, M., and Loguinov, D. (2004). Packet-pair bandwidth estimation: Stochastic analysis of a single congested node. In *Proceedings of IEEE ICNP 2004*.
- Katabi, D. and Blake, C. (2002). Inferring congestion sharing and path characteristics from packet interarrival times. Technical report, MIT-LCS-TR-828, MIT.
- Kenney, J. F. and Keeping, E. S. (1951). *Mathematics of Statistics*, volume 2. Princeton, NJ: Van Nostrand, 2nd edition.
- Keshav, S. (1991). A control-theoretic approach to flow control. In *Proceedings of SIGCOMM*.
- Moldován, I., Dang, T. Dinh, Bíró, J., Satoh, D., and Ishibashi, K. (2004). Bottleneck links detection method based on passive monitoring. In *Iasted CIIT 2004*.
- Molnár, S. and Miklós, Gy. (1998). Peakedness characterization in teletraffic. In *IFIP TC6, WG6.3 conference PICS'98*.
- Pearson, K. (1905). Das fehlergesetz und seine verallgemeinerungen durch fechner und pearson. *Biometrika*, 169–212.
- Riedl, A., Perske, M., Bauschert, T., and Probst, A. (2000). Dimensioning of ip access networks with elastic traffic. In *First Polish-German Teletraffic Symposium (PGTS 2000)*.
- Varga, P. and Kún, G. (2005). Utilizing higher order statistics of packet interarrival times for bottleneck detection. In *Proceedings of IFIP/IEEE E2EMon*.
- Varga, P., Kún, G., Fodor, P., Bíró, J., Satoh, D., and Ishibashi, K. (2003). An advanced technique on bottleneck detection. In *IFIP WG6.3 workshop, EUNICE 2003*.
- Varga, P., Moldován, I., Dang, T. Dinh, Simon, Cs., Kún, G., and Tatai, P. (2004). Developing a passive measurement-based methodology for detecting network bottlenecks in ip networks. Technical report, Study for Hungarian Telecom, in Hungarian.