Offline and Online Network Traffic Characterization

Su Zhang and Mary K. Vernon

University of Wisconsin Madison, Computer Science Department,
1210 W. Dayton St. Madison, U.S.
{zs,vernon}@cs.wisc.edu

Abstract. This paper investigates a new technique called Bayesian-Block-Analysis (BBA) for analyzing the time varying rate of events. The first goal is to evaluate the accuracy of BBA in identifying the rate changes in synthetic traces that have a given interevent times distribution, and known rate change points. We find that BBA is highly accurate on traces with exponential interevent times and known rate changes, and reasonably accurate with more heavier-tailed interevent times. The second goal is to apply BBA to actual network event traces. And for request arrivals or loss rate traces, BBA identifies significant stationary-rate periods which are qualitatively consistent with previous results obtained with less efficient or less accurate techniques. For packet arrivals to gateways, BBA identifies stationary rate periods that are corroborated by binning the data on a new timescale. Finally, we also show BB-online rate estimation is accurate for synthetic as well as actual system traces.

Keywords: Network Traffic Characterization, Bayesian Analysis, EWMA, Rate Estimation, Stationary Rate Period, Loss Rate, Packet Arrival Rate.

1 Introduction

An important problem in analyzing various types of network traffic events – such as request arrivals to a web server, packet losses in a packet flow, or packet arrivals to a gateway – is to determine how the average event rate varies with time. A related question is how the inter-event times are distributed during a period in which the average event rate is stationary. Accurate methods for obtaining such results can yield insight into how traffic varies at different points in the network as well as how to generalize the measured behavior at a given point to create representative workloads for system design. How frequently the event rate varies, for example, may impact the stability of various traffic control algorithms such as the control algorithms in the proposed new RCP [4] and XCP [7] transport protocols. On the other hand, obtaining accurate results for the average event rate vs time is challenging since the inter-event times are highly bursty (even in the case that the arrival process is Poisson), and changes in rate occur at unpredictable times.

For packet arrivals, previous studies [10, 11] showed evidence of burstiness on different time scales. For such arrivals, the self-similar and multifractal models [1, 6, 12] are developed to match the statistical properties of the observed network traffic at both small and large timescales. But those models are not easy to apply because there is very

I.F. Akyildiz et al. (Eds.): NETWORKING 2007, LNCS 4479, pp. 912–923, 2007.
© IFIP International Federation for Information Processing 2007

little intuition about how to modify the parameters to represent a range of workloads that the system should be designed to handle. Recent papers [8] show that packet arrivals on high bandwidth links tend towards Poisson because of the quickly increased multiplexing of the traffic on the network links. But a key question is how to identify the periods in which the event rate is stationary.

A widely used method for estimating arrival rate as a function of time on-line – that is, as each new arrival event occurs – is the exponential weighted mean average (EWMA). This is a weighted sum of the estimated rate at the previous arrival event and the inverse of the most recent inter-arrival time. Recent work by Kim and Noble [9] shows that the EWMA estimates can be agile in detecting rate changes or stable during periods of fixed average rate, but not both simultaneously. They propose several alternative ad hoc filters which improve on EWMA, but still have the same general problem.

Recent studies have used several off-line methods to determine average event rate as a function of time for various traffic traces. Zhang et al. [14] use two change point detection methods, known as "bootstrap" and "rank order" to study, for example, loss event rate vs. time for packet streams. These methods are computationally expensive and require a relatively large sample size, and as noted in their paper, are known to have non-negligible errors in terms of missed change points and false positives, respectively. Other recent work has used ad hoc binning to characterize request arrival rate to Internet media servers [2] and job arrivals to a large Internet compute server [3]. Such binning methods [11, 2, 3] are based on an ad hoc bin size and have two substantial drawbacks. First, periods of fixed rate are only identified by visual inspection and the endpoints are at the predefined bin boundaries. Second, each bin must have a good statistical sample and rate changes that occur within a bin cannot be identified. Finally, [8] apply the Canny Edge Detector algorithm [15] to the curve of cumulative event count versus time for packet arrivals to a high bandwidth (OC-48) Internet link. However, the accuracy of this method is very sensitive to the time unit used for updating the cumulative event count, and the time unit that yields accurate results for a given trace is unknown.

This paper evaluates a possible new approach to determining the periods of stationary network event rate, namely a recently proposed highly efficient Bayesian analysis technique called Bayesian Blocks (BB), which was developed for characterizing the periods of constant brightness in photon counting data [13]. Key advantages of the BB analysis are, (1) it is computationally efficient, (2) the time unit and other input parameter values that yield accurate results are relatively easy to determine, and (3) it can be applied to a small sample size. These properties are particularly useful for on-line use of the technique. The BB analysis assumes a non-stationary Poisson event process, but there are intuitive reasons that it might also be successfully applied to more general arrival processes. Specifically, we make the following contributions:

- We quantify the accuracy of the offline BB analysis by applying it to a large number of *synthetic* event traces with exponential or heavier-tailed distributions of inter-event times and known rate changes. To our knowledge, the accuracy of the BBA technique has not been assessed previously.
- The BB technique was developed for off-line data analysis, but we also apply it on-line, to estimate that rate at each new arrival event (without knowledge of future arrival events).

- The synthetic trace results show that 80-90% of the off-line BB rate estimates are within 30% of the actual rates in synthetic traces, even if the inter-event times have a heavy tail. The on-line BB rate estimation is less accurate than the off-line analysis, but is still significantly more accurate than previous on-line rate estimation methods.
- We apply the BB analysis to a variety of measured network event traces, including packet arrivals to a gateway, request arrivals to Internet servers, and loss events in long Internet packet streams, and illustrate the insights that can be obtained.
- We find that the BB analysis can analyze a trace with over 106 events and over 5000 change points in just a couple of minutes on a modern desktop (Pentium M 1.66HZ).

Section II provides a brief description of the BB analysis technique. Section III presents the accuracy assessment of the BB analysis technique using synthetic traces. Section IV applies both off-line and on-line BB analysis to loss events in long Internet flows as well as Internet server request arrivals. Section V applies the BB analysis to characterize network gateway traffic. Section VI concludes the paper.

2 Background

The BB analysis algorithm is both conceptually simple and computationally efficient, and is described very well in [13]. The interested reader is referred to that paper to understand how the algorithm works. A key feature of the algorithm is that for each interval beginning with the full trace of event times, it considers each possible partitioning of the interval into two intervals with different event rates, using an input parameter called the “odds threshold” (OT) and statistical maximum likelihood measures to decide whether the interval should be partitioned at the point with the highest likelihood of a change point. In addition to the OT parameter, there is a Minimum Interval (MI) parameter that defines the minimum number of events per period of fixed average arrival rate.

3 Accuracy Assessment

3.1 Synthetic Traces

Each synthetic trace contains event times that are generated with one of the following distributions of time between events during each period of fixed average event rate: exponential, lognormal (logn), or 2-stage hyperexponential (h2). The absolute value of average event rate is immaterial because during the BBA we define the time unit for the event times to achieve an average of one event every 25 time units. This time unit is defined such that it is small enough so the probability that multiple arrivals will occur in the same time unit is negligible, since the calculations assume this will never occur, and it is as large as possible so the factorials from the number of time units in the interval do not dominate the calculated likelihood. Hence, the key parameter of the logn and h2 distributions is the coefficient of variation (CV) during each period of fixed average rate. We consider logn and h2 distributions with CV of 1.2, 2, and 3, motivated by results to date for actual event traces which are summarized in Section 4, 5. We expect that BBA accuracy will decrease as CV of the inter-event times increases, and also note that h2 inter-event times with CV=3 have a heavy tail.

Each synthetic trace has the following further parameters: 1) number of changes in average rate (*n*), 2) number of events (*m*) between adjacent changes in rate, and 3) the rate change factor (*rcf*), which is the magnitude of the each change in rate. We consider values of *n* ranging from 1 to 64, and find that the *BBA accuracy is independent of n*, as might be expected. We vary *m* from 200 to 2000 when $CV \leq 1.2$, and otherwise from 1000 to 4000, and find as expected that accuracy increases as *m* increases (benefit from larger sample sizes). Similarly, we expect accuracy to increase as *rcf* increases, and we consider values of $rcf \geq 1.5$. In each trace, each rate change is a rate *increase* with probability 0.5 and is otherwise a rate *decrease*. For each combination of inter-event time distribution and values of *m* and *rcf*, we generate 10,000 traces for evaluating BBA accuracy.

Results for a representative trace with $n = 2$, m=200, and Poisson inter-event times are shown in Figure 2; results for a trace with $n = 7$, $m = 4000$, and h2 inter-event times with CV=3 are given in Figure 3. Unless stated otherwise, the synthetic traces contain abrupt rate changes. Figure 2(d) provides results for a trace with gradual rate change for comparison.

3.2 Accuracy Measures

The traditional accuracy measures for rate change point detection algorithms (e.g., [9]) are: (1) fraction of known rate changes that are detected within a pre-defined distance of the actual rate change, and (2) rate of false positives. However, these measures do not account for predicted rate changes that occur just before or just after the pre-defined distance, nor do they measure the accuracy of the magnitude of a correctly or falsely predicted rate change. Hence, we choose a more comprehensive measure, f_e where e=15 or e=30, such that f_e is equal to the fraction of actual event times at which the predicted rate is within e% of the actual event rate immediately prior to that event. Note that an undetected rate change will yield incorrectly predicted rate values for each event in the next period of average rate.

3.3 Parameter Value Selection

BBA has two parameters other than the inter-event time unit which we define such that on average an event occurs every 25 time units. The two parameters, Odds Threshold and Minimum Interval (*OT* and *MI*), are such that increasing either value decreases the number of falsely predicted rate changes but also decreases the number of correctly predicted rate changes. We use BBA of the synthetic traces to determine which values of these BBA parameters leads to the highest overall f_e accuracy.

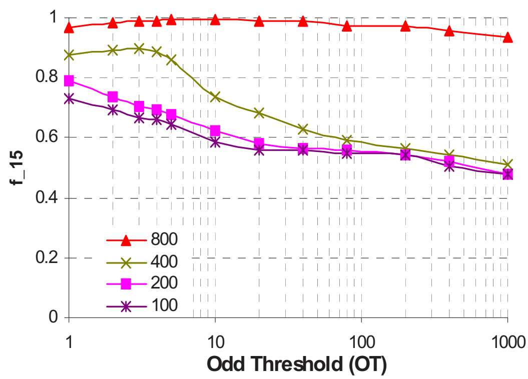

Fig. 1. Accuracy V.S. Odd Threshold (OT)

To conserve space, we omit the experimental results which show that for interevent time distributions with CV ≤ 1.2, the fraction of accurate rate estimates improves slightly as *MI* increases from 2 to 10 (close to the assumption, not much false positives), but does not improve for larger values of *MI*. Similarly, for inter-event time distributions with CV = 3, the overall accuracy on the synthetic traces improves as *MI* increases to 80 (because most false positive periods <80), but does not improve for larger values of *MI*. Hence, for all further results in this paper, we use *MI*=10 when CV ≤ 1.2 and *MI* = 80 when CV=3.

Figure 1 shows the average f_{15} for traces with *rcf* = 1.5 and CV ≤ 1.2. The four curves in the figure correspond to different values of *m*. These results show that a smaller OT produces a higher overall accuracy, as measured by the fraction of the BBA rate estimates that are within 15% of the actual rate. Omitted due to space constraints are results for traces with CV = 3 where for each m ≥ 1000, average accuracy improves gradually as OT increases from 5 to 10,000 and then accuracy decreases for larger values of OT. For the higher CV, a larger OT is needed to avoid false positives. For the further BBA results in this paper, we use *OT* = 4 when CV < 2, and *OT* = 1000 for CV ≥ 2.

3.4 BBA Accuracy Results

Figures 2, 3 provide results for particular traces that illustrate the *typical accuracy* of the BB analysis for different types of traces. The results in figures 2 (a, b, c) is for a trace with exponential interevent times (CV = 1), *n* = 2, *m*=200, and *rcf* = 2. The actual rate (equal to 10 for the first 200 events) is shown with a gray diamond symbol at selected points, while the predicted rate at each event is plotted with the solid black curve. Figure 2(a) shows that the BBA predictions closely match with the actual event rates throughout the trace. Figure 2(b) shows that a simulated BBA “on-line” that is at each event time BBA predicts the rate using all of the prior event times, most predicted rates still match fairly closely with the actual rate. In particular, the BBA-online is significantly better at predicting the location of rate changes and has overall more stable estimates than EWMA estimates which are illustrated in Figure 2(c). A

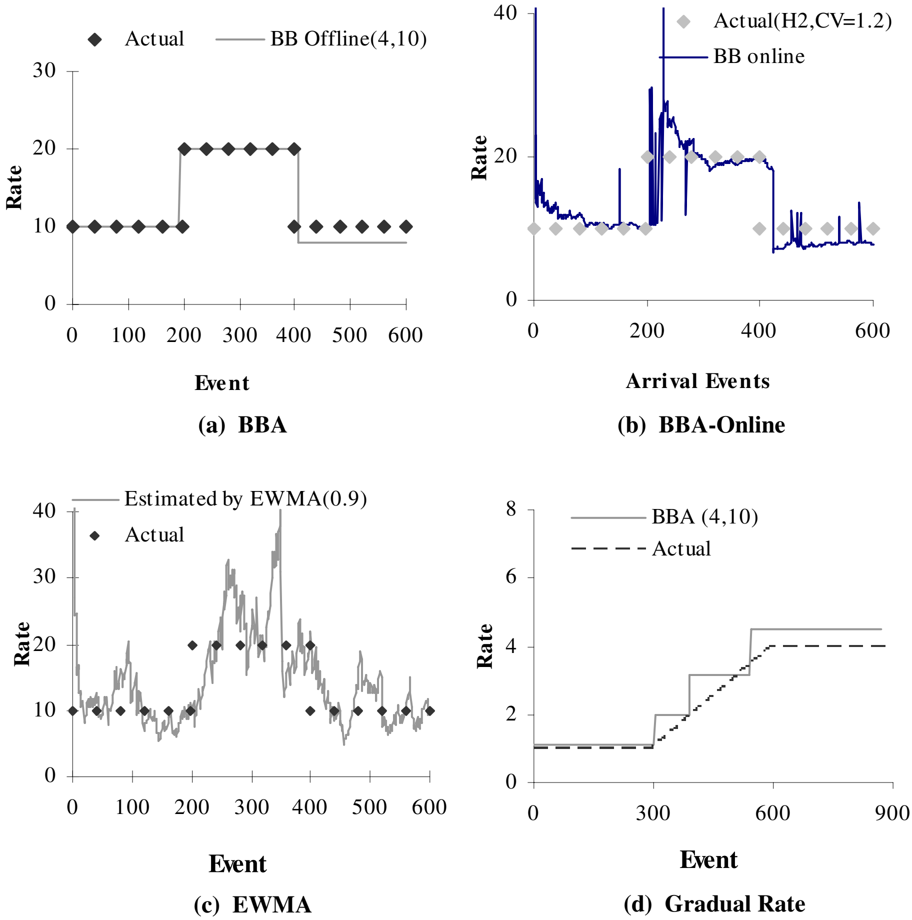

Fig. 2. Representative Accuracy (CV=1, m=200, RCF=2 for a, b, c)

key result of our experiments is that *BBA-online is significantly more accurate than EWMA*, over all possible values of the EWMA weight parameter.

Figure 2(d) provides results for a trace similar to previous ones in Figure 2, but with one gradual rate change instead of two abrupt rate changes. Again BBA is quite accurate, and BBA online estimates (omitted to conserve space) are significantly more stable than EWMA.

Figure 3 shows that BBA is also reasonably accurate for most of the time when the inter-event time distribution has CV=3 and m = 4000. However, the BBA estimates for such traces have a reasonably high number of falsely predicted rate "spikes", having width less than 100 events and rcf >2. Since those rate spikes can easily be "erased", we consider measures of f_e both with the spikes included and with them erased and replaced by the higher of the two rates before and after the spike.

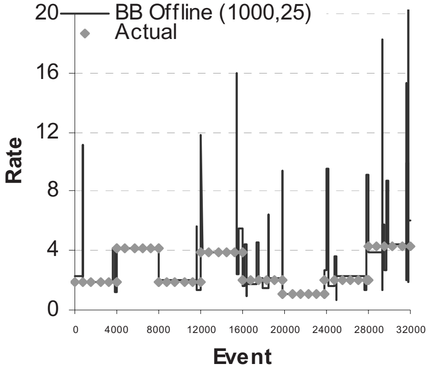

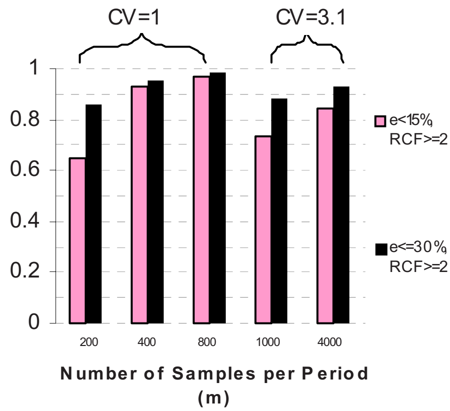

Fig. 3. Representative Accuracy for Abrupt Rate Change, H2, CV=3.1

Fig. 4. Quantitative Accuracy

Figure 4 summarizes the overall average values of f_{15} and f_{30} as a function of CV, m, rcf, and e. The higher value in each pair of bars is for the case that e=30 and the "spikes" are erased. Note that for traces with CV $\leq$ 1.2 and rcf $\geq$ 2, $f_{30} > 90\%$ when m $\geq$ 400 and $f_{30} > 80\%$ when m $\geq$ 200. For traces with CV = 3 and rcf $\geq$ 2, $f_{30} > 90\%$ when m $\geq$ 4000 and $f_{30} > 80\%$ when m $\geq$ 1000. Under these conditions, we can use BBA to achieve fairly accurate analysis of the event rate in the actual event traces, with the understanding that 10-20% of the BBA rate estimates have error greater than 30%. In the next section we analyze several actual event traces and interpret the results in light of these accuracy results for the synthetic traces.

4 Media Server Arrivals

4.1 Off-Line Characterization of Media Server Load

In this section we apply the BB *offline* analysis to characterize client session arrivals to a media server (called BIBS), which were previously found to be Poisson [2], and to characterize the client interactive requests to a different media server (called eTeach), which were found to have a Pareto inter-arrival time distribution during periods of approximately stable arrival rate. In the previous work, stationary periods were determined (approximately) by visual inspection of binned request arrival counts. Figure 5 (a) shows the result of BBA applied to a BIBS one day trace with highest load.

We observe that the stationary periods identified by the BB analysis agree with the binned measures of number of arrivals in each hour, yet are significantly more precise. In particular, the BB analysis (a) clearly identifies the (most likely) endpoints of the stationary periods, (b) reveals longer intervals of constant rate than is apparent in the binned data, and (c) more precisely identifies the peak rates (e.g. at 9pm in

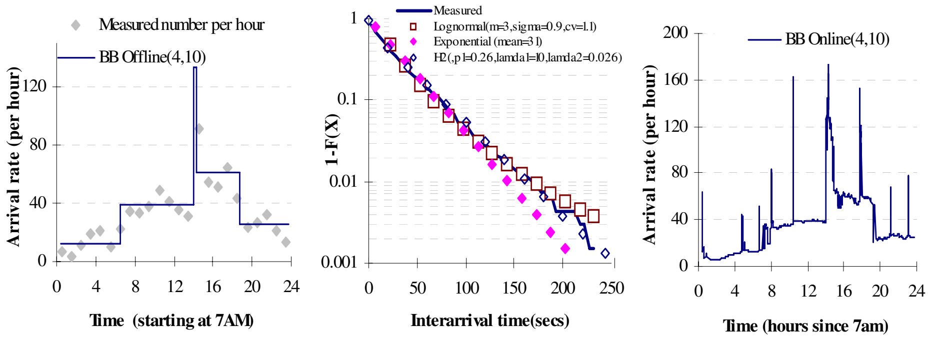

(a) Predicted & Measured Rate VS Time

(b) Interarrival time distribution for Hour 14.3-18.7

(c) Predicted Rate VS Time

Fig. 5. Actual & BBA Predicted Media Server Request Arrival Rate

Figure 5(a)) which can be obscured by other analysis techniques and may be important for system design or configuration. The precise stationary periods also delineate the samples that can be used to obtain more detailed measures of the client arrival process. Further more, the BB results provide a simple global characterization of the observed client session requests as a time varying Poisson process with relatively infrequent, abrupt rate changes. The stationary intervals that have the highest rate, lowest rate, and largest duration are potentially of greatest significance for system design, configuration and optimization.

Figure 5(b) shows the typical results for the request interarrival time distribution in each of the twenty highest rate periods found in the high rate days for the BIBS server. Similar results were obtained for the ten highest rate periods on the high rate days for the eTeach server. In both cases, 97-98% of the measured interarrival times fit the exponential distribution. On the other hand, the full distribution is more precisely modeled by a two-stage hyperexponential with a slightly heavier tail than the exponential. For both servers, this characterization is more precise than in the previous ad hoc analysis [2] because the periods of stationary rate could not be precisely delineated in the ad hoc analysis. The results for the eTeach server are also somewhat significant because several previous *ad hoc* characterizations of interactive client requests to Web servers (e.g. [2] and citations therein) have found the interarrival distributions to deviate significantly from the Poisson. Thus, the BB analysis provides a substantially new characterization of the interactive client requests to a Web media server, namely that these requests are nearly Poisson during the stationary periods, with a relatively small number of rate changes per day.

4.2 On-Line Estimation of Media Request Rate

In this section we evaluate the BB analysis as a technique for estimating the media server request arrival rate during system operation. Such arrival rate estimates might be needed by media caching algorithms or by the Patching streaming protocol to compute the maximum duration of the patch streams. For the BIBS request trace that was characterized off-line in Figure 5(a), Figures 5(c) provides the on-line rate estimation at each client arrival using the BB analysis..

The key result is that, generally, the BB online analysis provides significantly more stable and accurate online estimates than EWMA (EWMA results are similar as in Figure 2(c)). Very short transient spikes in the online rate estimate are due to local fluctuations in the interevent time. In some cases, it may be useful to take action in response to the high rate estimate. For example, the spikes could indicate that the file should be temporarily cached in memory, since disk bandwidth is limited and the file need not be cached for long. In other cases, it may be appropriate to ignore a temporary spike in estimated arrival rate, pending further evidence that the new rate will be sustained. We note that distinguishing the temporary spikes from the stationary rate estimates may be easier in the BB estimates than in the EWMA estimates, since the BB rate estimates are significantly more stable before and after the spikes. Thus, this BB-online analysis shows promise for online rate estimation.

5 Internet Packet Traffic and Loss Rate

5.1 Loss Events in UDP Flows

We use the off-line BB analysis to characterize the loss events observed during two different 24-hour UDP packet flows. Figure 6 provides results for a flow which was transmitted between two sites in a metropolitan area network. We have similar observations on loss event traces as in media server analysis. Especially we identified up to 1.5 hours stationary loss rate period which could be well modeled by a Poisson process. And BB-online also is generally significantly more stable and accurate than previous rate estimation methods (e.g. ALI(8)).

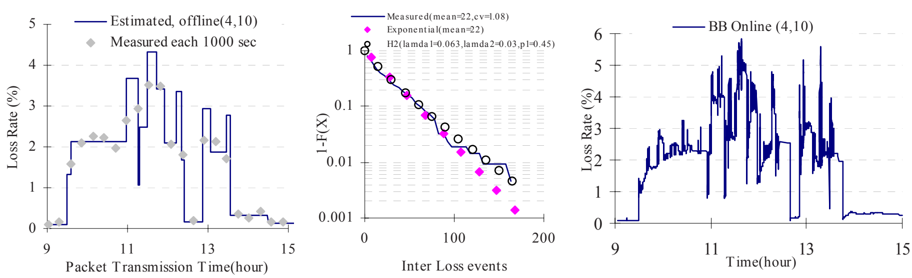

(a) Predicted & Measured Rate VS Time

(b) Inter Loss event distribution for Hour 9.6-11

(c) Predicted Rate VS Time

Fig. 6. Actual & BBA Predicted Loss Rate

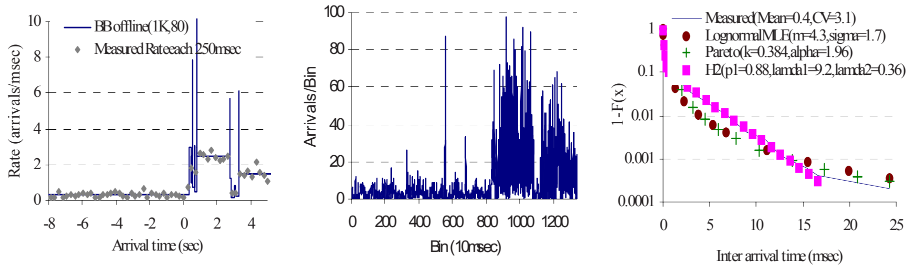

(a) Stationary Rate Period of 5 sec interval with 1st highest rate

(b) Binned Packet Arrivals Count

(c) Representative Interarrival Distribution

Fig. 7. BBA on IPEX Gateway Trace

5.2 Packet Arrivals to the IPEX Gateway

In this section, we apply the off-line BB analysis to characterize packet arrivals in an IPEX gateway [5] trace (32 hours) and a heavily multiplexed OC48 trace studied by [8] (1 hour).

In Figure 7(a) we provide the BBA results including a 5-second interval with the largest number of packet arrivals (8699). The BBA result is commensurate with measured rate from binned packet arrival counts during each 250 milliseconds interval. Notably, the BB analysis finds intervals of several seconds in duration during which the packet arrival rate is estimated to be stationary, punctuated by abrupt (and large) rate changes.

Figure 7(b) shows the packet arrival counts in 10 milliseconds for the same period in 7(a). In contrary to (a), the stationary intervals are not evident. And the highly variable measures are similar to what previous work showed using same binning method. This result suggested that previous binning results could not reveal the stationary rate periods because (1) It is highly dependent on the choice of the bin size. Too small bins cause high variability due to statistical fluctuations in the small samples and too coarse bins make the bin boundaries not likely aligning with endpoints of the stationary periods. (2) Also many bins plotted tightly causes visual effect of high variability. In contrast we selected the 250 msec as the bin size in 8(a) based on the estimated rate from BBA which reveals multi-second stationary rate periods.

Figure 7(c) shows that the packet interarrival times during a stationary period are well modeled by a two-stage hyperexponential distribution, rather than by a Pareto or lognormal distribution.

5.3 Packet Arrivals to OC48 Trace

We further analyzed the OC48 trace which does not have as highly variable packet arrivals as IPEX trace due to the heavy multiplexing. Figure 8(a) shows the BBA result from 1 minute interval with most number of arrivals in one hour. BBA reveals

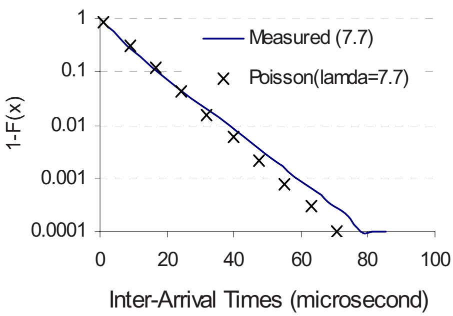

(a) OC48, BBA, 30see with most arrivals starting at 9:20AM, 8/14/2002

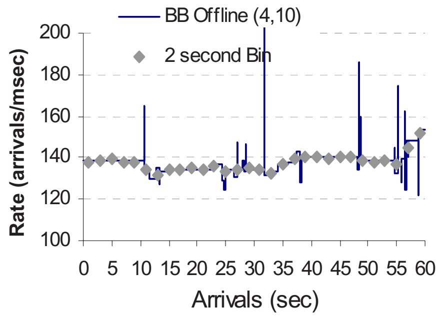

(b) Interarrival Distribution for Period 38.2-49 sec

Fig. 8. OC48, BBA

stationary rate period as long as 8 seconds with binned data commensurate with its result. This result is contrast to the findings of sub-second intervals from [8]. This is probably due to the limitations of Canny Edge Detection used in [8] which is sensitive to the choice to the parameter.

Also the packet inter-arrival times during the stationary rate period we found in OC48 traces could be modeled approximately by Poisson as shown in Figure 8(b).

Further research is needed, including analysis of packet traces for other sites before definitive conclusions can be drawn, but it appears that overall, the BB analysis results provide significant new insights into the local and medium timescale behavior of the packet arrival process that are relevant for creating network traffic workloads for system design.

6 Conclusion

This paper has investigated a recently proposed efficient technique, called Bayesian Blocks (BB), for characterizing the time-varying rate in a bursty event stream. Key properties of the BB analysis are that it is simple to apply, computationally efficient, and requires a relatively small sample size. The accuracy of off-line BB analysis was assessed by applying the technique to a variety of synthetic traces with known rate changes. Off-line BB analysis was found to be able to accurately identify each period of constant rate as well as each period of stationary average rate during which the interevent times have a heavy-tailed distribution. BB analysis was applied to a variety of measured event traces of interest and all of the event traces we analyzed were found to have significant periods of stationary rate. Finally, we found BB-online to be significantly more accurate than previous on-line rate estimation methods.

Future work includes applying the BB method to further traces such as TCP connection arrivals and FTP data connection arrivals. We are also interested in evaluating BB and other methods for detecting changes in the average round trip times and changes in the available bandwidth for network flows, which are needed for high performance rate control. Finally, further development of the BB-online algorithm is a topic of our current research.

References

1. P. Abry and D. Veitch, Wavelet Analysis of Long-range Dependent Traffic, IEEE Trans. on Info. Theory, Vol. 44, pp. 2-15 (1998)
2. J. M. Almeida, J. Krueger, D. L. Eager, and M. K. Vernon, Analysis of Educational Media Server Workloads, Proc. Nossdav'01, (2001).
3. S. Chiang and M. K. Vernon, Characteristics of a Large Shared Memory Production Workload", Proc. 7th Workshop on Job Scheduling Strategies for Parallel Processing (2001)
4. Nandita Dukkipati, Nick McKeown, Processor Sharing Flows in the Internet, Stanford HPNG Tech report (2004)
5. IPEX: Collaboration for Internet Traffic Measurement, http://www.xiwt.org/ipex/ipex.html
6. A. Feldmann, A. C. Gilbert, And W. Willinger, Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic, SIGCOMM (1998)
7. Dina Katabi, Mark Handley Charlie Rohrs, Congestion Control for High Bandwidth-Delay Product Networks, Sigcomm (2005)
8. T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido, A nonstationary Poisson View of Internet Traffic, Proc. Infocom 2004, Hong Kong, March (2004)
9. M. Kim and B. Noble, Mobile Network Estimation, MOBICOM (2001)
10. W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, On the Self-similar Nature of Ethernet Traffic, IEEE/ACM Trans.on Networking, Vol. 2, pp. 1-15 (1994)
11. V. Paxson and S. Floyd, Wide-Area Traffic: The Failure of Poisson Modeling, IEEE/ACM Trans. on Networking, Vol. 3, No. 3, pp. 226-244, June (1995)
12. R. Riedi, M. Crouse, V. Ribeiro, and R. Baraniuk, A Multifractal Wavelet Model with Application to Network Traffic, IEEE Trans. on Information Theory, Vol. 45, No. 3, pp. 992-1018(1999)
13. J. D. Scargle, Studies in Astronomical Time Series Analysis v. Bayesian Blocks, A New Method to Analyze Structure in Photon Counting Data, The Astrophysical Journal, Vol. 504, Sept. 1, pp. 405-418.(1998)
14. Y. Zhang, N. Duffield, V. Paxson, and S. Shenkar, On the Constancy of Internet Path Properties, Proc. Internet Measurement Workshop, San Francisco, (2001)
15. J. Canny. A Computational Approach To Edge Detection. In IEEE Transactions Pattern Analysis and Machine Intelligence (1986)