

Detecting Period and Burst Durations in Video Streaming by Means of Active Probing

Bjørn J. Villa and Poul E. Heegaard

Abstract—This paper presents a method for detecting periodic behavior and burst durations in IP traffic observed on access links. Periodic behavior in traffic observed on the Internet is growing due to the increasing amount of services with video components. The method is based on active probing and serial correlation of the observed probe packet inter-arrival times on the receiver side. Verification of the method is done based on an experimental implementation, using both controlled lab services and live public services. The services used in the experiments are adaptive video streaming based on Microsoft Smooth Streaming and Netflix. The results show that the method provides accurate results for both the estimator of service period and the burst durations. The findings are similar both for the controlled lab experiments and the live network measurements. The amount of probe traffic required in order for the method to perform satisfactory is very moderate.

Index Terms—Periodic traffic patterns, traffic burst duration, active probing, serial correlations, adaptive video streaming, access links.

I. INTRODUCTION

The amount of services provided to Internet users around the world following an Over-The-Top service delivery model is increasing. This model is based on that the involved network operators are not taking any active measures in order to assure the required Quality of Service (QoS) levels of the specific services. Traffic is carried as part of the best-effort class and will therefore face obvious challenges in terms of being able to meet the end users expectations regarding Quality of Experience (QoE).

The best-effort traffic class on the Internet is a rough place, and the handling of traffic is unpredictable. This can lead to fluctuations in the experienced QoS metrics with the potential of a negative service effect. The engineering approach to this with an Over-The-Top perspective is to make the services adaptive. The type of adaption required is for the service itself to be able to adjust their QoS requirements. The case of fixed QoS requirements, e.g. with regard to bandwidth or even packet loss ratio would not work well in a best-effort network class.

The strong growth of services with video components over the last years, are predicted to become even stronger in the years to come [1]. This success is partly caused by the emerging solutions and standard [2] for adaptive video streaming. The adaptive nature of these services have proven to be very effective, and enabled new business models (e.g. Netflix, YouTube) on the Internet for content which earlier

was considered impossible to deliver in a best-effort QoS class.

An interesting consequence of the growing amount of video services, and adaptive ones in particular – is that the traffic patterns on the Internet are changing. This alone is of course interesting enough, but in the context of adaptive services it may also create some interesting effects. An interesting research question arises from this, as it will become more and more common that adaptive services must adapt according to other adaptive services. Whether this will lead to a continued or reduced success in terms of effectiveness for this type of services remains to see.

A. Problem Statement

The research topic addressed in this paper covers a method to detect and characterize services with a periodic behavior on access links, by means of active probing. As basis for this work we have used adaptive video services, which typically embed this behavior. Thus, the method does not address the specific recognition of a periodic service as being video or something else. The method focuses on detecting the period (if any) and also the burst duration of traffic inside each period. Further on, a basic evaluation of how much probe traffic is needed and also the required sample size is covered.

The investigated method provides estimates for parameters (T_p , T_b) indicated in Fig. 1. The periodic pattern of the video service used in the experiments consists of burst and idle periods. However, it should be noted that even in the idle period there is traffic, but much less than in the burst periods.




Fig. 1. Traffic pattern for a video service.

The value of this work can be viewed as both a general contribution to the discipline of service characterization, but also as a candidate component of future methods for estimating available bandwidth. Obtaining knowledge of periodic components in the cross-traffic and also the duration of burst periods has a potential of increasing effectiveness and accuracy in available bandwidth estimations.

B. Research Approach

The chosen approach for investigating the suggested method of detecting periodic behavior and burst duration, is by means of experiments in a controlled lab environment and a real network with public services. In the lab environment

Manuscript received February 12, 2013; revised April 2, 2013.

The authors are with the Norwegian Institute of Science and Technology, Department of Telematics, Trondheim, Norway (e-mail: {bjorn.villa, poul.heegaard}@item.ntnu.no).

we use an adaptive video streaming service based on the Smooth Streaming platform from Microsoft, while in the live network we are accessing a the public Netflix service. The required background load in the lab environment is provided by means of a CBR type traffic source. The complete measurement setup will be further described in Section V.

C. Paper Outline

The structure of this paper is as follows. Section II provides an overview of related work; Section III provides the relevant characteristics of adaptive video streaming as captured by passive measurements; Section IV describes our method of detecting periodic traffic patterns and burst duration by means of active probing; Section V describes our hybrid passive and active measurement setup; Section VI presents the results based on the active probing and compares them with the passive measurements; Section VII presents our conclusions and an outline of future work is given in Section VIII.

II. RELATED WORK

The study of phenomena's observed in Internet service usage and traffic patterns generated are done in different ways. Quite often the phenomenon of interest is required to be studied over time in order to obtain sufficient information. In such cases, the tools and methods available for analyzing a discrete set of time-ordered data (i.e. time series) are quite useful [3]. For the purpose of detecting periodic components in traffic pattern and the duration of such, estimators for serial correlation [4], [5] (also known as autocorrelation) are known to be very efficient.

In the area of available bandwidth estimations along a network path, there are many different approaches to how this can be done using active probing techniques [6], [7]. The active probing results in a time series of observations, such as changes in delay components [8]. These observations are then used as input to the respective algorithms for estimating the available bandwidth. There are many challenges of performing such estimations in an accurate manner [9], among which neglecting the burstiness of cross-traffic is one. Among the more recent methods in this domain, the approach described in [10] is quite interesting. It aims at estimating available bandwidth in real-time and does this by applying a filter-based method. The idea of using filtering as part of the continuous processing of collected observations in our own work is inspired by this.

Especially on access links, the traffic patterns observed are dominated by burst components. The reason for this is composed by different factors. First of all, the strong dominance of TCP based applications [11] has been shown to generate traffic bursts in short time scales [12]. This follows directly from the TCP protocol behavior in terms of timeout events and congestion avoidance. Another important factor which leads to bursty traffic is the strong dominance [1], [13] of services with video components. In our previous work [14] we studied the packet inter-arrival time distributions for this type of services, and suggested a new method to achieve a shaping effect. The typical traffic pattern observed when video services are present has the signature of periodic components, containing a number of bursts.

With regard to the correlation structures in a time series of observed delay components, it has been shown [15] that the TCP protocol tends to generate traffic patterns with this property in sub-second time scales. In addition, the nature of services used also contributes to correlation structures in the time scale of seconds. The authors of [16] showed that they were able to detect Skype traffic by investigating correlation between traffic bursts in network traffic traces. A similar effect was shown in [17] where the focus was on video services.

A more general study of how correlation structures in traffic patterns can be obtained either through sampling or active probing is presented in [18]. In this, estimators for correlations of network traffic are described and experimental results presented. The estimators are applicable for both samples based on passive and active monitoring.

In our earlier work [19], [20], we studied the behavior and characteristics of adaptive video streaming from different perspectives. Further on, in our work in press [14] we presented a new method for achieving a traffic shaping effect for adaptive video streaming, without involvement from network components.

III. CHARACTERISTICS OF VIDEO STREAMING SERVICES

The periodic nature of an adaptive video streaming service is given by the repeated requests for the next segment in a video stream, at a specific quality level.

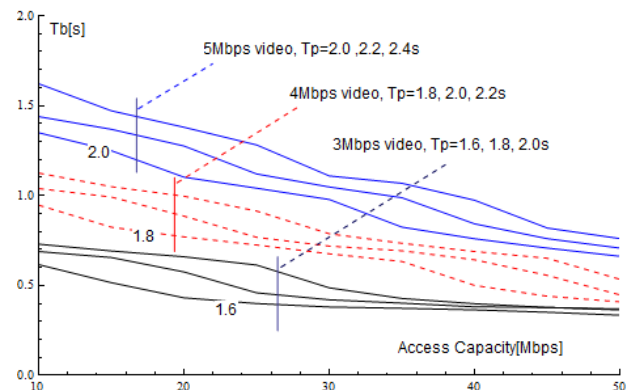


Fig. 2. Passive measured burst durations for MS smooth streaming

The passive measured burst periods for quality levels at 3/4/5Mbps over the access capacity range from 10-50Mbps are shown in Fig. 2. The segment request interval (T_p) used in the experiments for the MS Smooth Streaming [21] based service is between 1.6 and 2.4 sec. The interval granularity available in the specific solution was 0.1s, which gave five different T_p scenarios available for use in the experiments. Only a selection of this is shown in Fig. 2 for the different video stream quality levels and access capacity levels.

For the public Netflix service used in our experiments, we did not have the opportunity to create a similar wide range of scenarios as for the MS Smooth Streaming service. The obvious reason for this would be that it was a live service, and Netflix resources were not involved in our research. The characteristics of the Netflix streaming were studied in depth in [22], and significant differences between Netflix and MS Smooth Streaming solution are described. Relevant to our work are their findings that the Netflix streaming uses two

TCP connections and also that the request intervals relates to data volumes rather than duration.

For a specific movie, accessed from Netflix in both their stated HD quality and also as non-HD the passive measured burst periods by using TCPdump (cf. Fig. 8) are presented in Fig. 3.

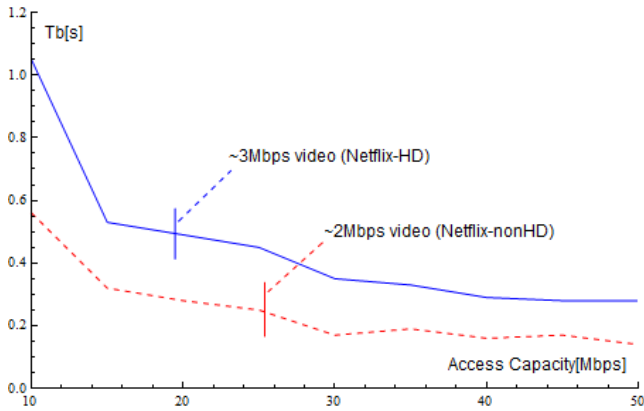


Fig. 3. Passive measured burst period durations for Netflix

In order to get a view on the traffic pattern generated by Smooth Streaming and Netflix in a scenario with a minimum of constraints, we measured the average burst periods using 100Mbps access capacity. In Table I the results obtained for all available video quality levels are presented. It should be noted that the Netflix service may be subject to some limitations outside of our control as it resides outside of the lab environment.

TABLE I: PASSIVE MEASURED BURST PERIODS AT 100MBPS

Burst Period	2Mbps	3Mbps	4Mbps	5Mbps
Smooth Streaming with $T_p=2.0s$	- na -	0.07s	0.09s	0.12s
Netflix	0.13s	0.21s	- na -	- na -

As a consequence of the Netflix request intervals being decided by data volumes, rather than playback time – the periodic nature of this streaming type could be questioned. The effect of using multiple TCP session is also interest in this regard. In Fig. 4, the estimated probability density functions (pdf) for passively measured time gaps between successive HTTP GET messages, per TCP session and combined are shown. The pdf estimations are based on measurement of the request intervals for the Netflix-HD movie at 30Mbps access.

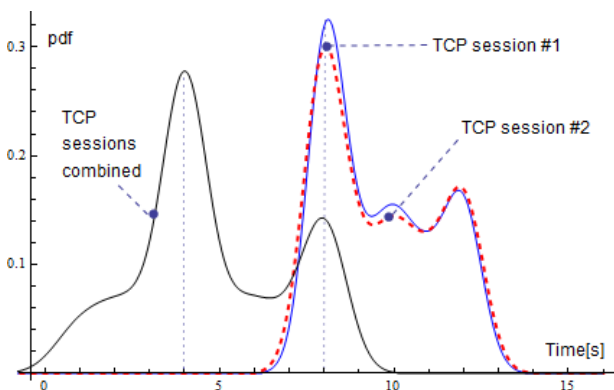


Fig. 4. Estimated PDF for Netflix period durations

It is interesting to see that each of the two TCP sessions has

a location in their distribution around 8 seconds. They also seem to multimodal, with peaks in additions to the dominating one at 8sec. When the traffic from the two sessions is studied together as they appear on the wire, we get a distinct shift in the distribution. The location for the combined pdf is around 4sec. This indicates that the TCP sessions are interleaved in such a way that a periodic pattern could be expected to have a period lower than each of the TCP sessions.

IV. METHOD

The method used in order to detect and estimate periodic behavior and burst duration of traffic on an access link is by mean of active probing, and a serial correlation at different lags for the times series of processed probe packet inter-arrival times.

A. Probe Traffic Generation

The probe traffic is generated according to trace files with configurable packet size and inter-packet times. The specification of the probe traffic is assumed known to the receiver. Further on, the receiver must be able to tell the difference between a probe packet received and any other traffic, and also that it is able to timestamp the received packets. As illustrated in Fig. 5, the time between probe packets sent is noted t_{in} and the measured inter-arrival times between consecutive probe packets at the receiver side is noted $t_{out,i}$. The time between packet pairs is noted t_{gap} .




Fig. 5. Packet-pair based probe traffic.

Depending on the cross-traffic profile, the receiver may see either an increase or decrease in the probe packet spacing (Δt_i). Recognizing that even cross-traffic during the t_{gap} periods may impact the $t_{out,i}$ samples it leads us towards some kind of processing of the samples before using them further.

B. Processing of Received Probe Traffic

In order to simplify the analysis part we found that a basic filtering mechanism for the observed $t_{out,i}$ values was very efficient. The purpose of this is to filter out only those $t_{out,i}$ values with a clear indication of that significant cross-traffic has occurred during the corresponding time interval. The filtering approach is based on using a cumulative moving average CM_i (cf. Eq. 1) of the $t_{out,i}$ samples, rather than the known probe packet spacing t_{in} as basis for deciding whether a sample indicates cross-traffic or not.

$$CM_i = \frac{t_{out,1} + t_{out,2} + \dots + t_{out,i}}{i} \quad (1)$$

The time series T_{out} of observed $t_{out,i}$ values is then passed through the filter (cf. Eq. 2) using the calculated CM_i value as input. The resulting time series is noted T_F and contains elements noted $t_{f,i}$.

$$\forall t_{out,i} \in T_{out}$$

$$\text{if } t_{out,i} > CM_i \text{ then } t_{f,i} = t_{out,i} \quad (2)$$

$$\text{if } t_{out,i} \leq CM_i \text{ then } t_{f,i} = CM_i$$

The elements of the time series T_F is then used as input to computation of lag- s serial correlation X_s (cf. Eq. 3). For the sake of reducing computation time, and also the time required to detect and estimate period and burst duration, the maximum lag investigated should be set to a reasonable level. The chosen level should match a threshold in time, which reflects the highest potential period of interest to us.

$$x_s = \frac{\sum_{i=1}^{N-s} (t_{f,i} - \bar{t}_f)(t_{f,i+s} - \bar{t}_f)}{\sum_{i=1}^N (t_{f,i} - \bar{t}_f)^2} \quad (3)$$

In our case, since the probe packets are equally spaced in time, the lag value s directly maps over to time by simply multiplying it with the probe packet period.

C. Analyzing Output from Serial Correlation

The output of lag- s serial correlation of the time series T_F gives a new time series consisting of correlation values X_s . Our method of detecting and estimating period and burst duration in the video streaming services used in the experiments is based on inspection of the X_s values. To demonstrate the strengths of this approach, we would first like to present a graphical view of the serial correlation of a theoretical discrete signal F_i (cf. Fig. 6) with a period of 300.




Fig. 6. Theoretical discrete periodical signal

This signal can also be represented as a time series consisting of a repeating pattern of length 300 with values of either 0 or 1. When the lag- s serial correlation is computed for this, we get an output as shown in Fig. 7.

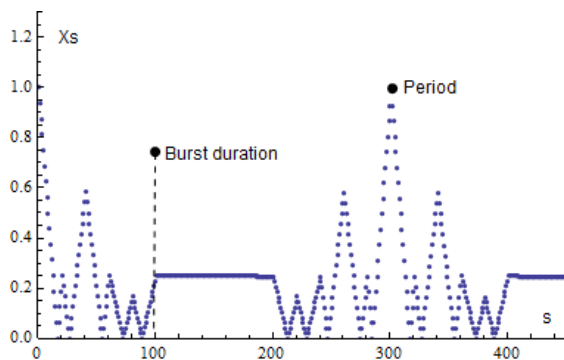


Fig. 7. Serial correlation for theoretical signal

From investigating the graphical view of the serial correlation we can clearly see the period of the original signal indicated by the peak correlation value at lag 300. We can also see the width of 100 for the sub-period in the original signal which contains the three spikes. These capabilities of the serial correlation are well described in [5], and we will use this as basis for analyzing our measurement results for the purpose of validating our method.

V. MEASUREMENT SETUP

The hybrid active and passive measurement setup used for performing the experiments related to this paper contains several components, ranging from the client side over to the server side as shown in Fig. 8. On the client side the video service is accessed by a dedicated Windows based PC, while the client receiving the probe traffic is a separate Linux based PC. On the server side, we have the dedicated Microsoft Smooth Streaming server and also the probe traffic generator. For the purpose of access to the live Netflix service, the lab is connected to the Internet. The access network part consists of commercial off-the shelf products which give access to useful functions for controlling bandwidth similar to those used in commercial networks.




Fig. 8. Hybrid active and passive measurement testbed

The CBR traffic source is based on the Click Modular Router [23], which is a simple solution for generating basic types of traffic. The probe traffic sender and receiver which are used for the active measurements are based on the CRUDE/RUDE tool [24], which has the required capabilities of generating traffic pattern based on trace files. The TCPdump node is used to provide the actual characteristics of the services by means of passive measurements, as presented in Section III. These characteristics serve as basis for verifying the effectiveness and accuracy of the method we are studying.

VI. RESULTS

The results to be presented in this section cover both experiments with Smooth Streaming in a lab environment, and similar experiments using the live Netflix service. The results for the controlled lab service is more elaborate than the Netflix service. The characteristics for both services captured by means of passive measurements (cf. Section III) should be used as reference point when considering the following results.

In the presentation of serial correlation results for the probe traffic inter-arrival times, the lag parameter has been converted to time for the sake of making the results easier to read. The lag to time conversion follows directly from the probe period used.

A. Smooth 5Mbps Streaming on 100Mbps Access

The scenario with potential of having the most burst oriented traffic is the scenario with full access speed (100Mbps) and a single service active using any one of the available T_p values between 1.6s and 2.4s. As we can see from Fig. 9, the computed serial correlations using the method described in Section IV gives very distinct peaks at the expected locations, i.e. at the T_p values being used.

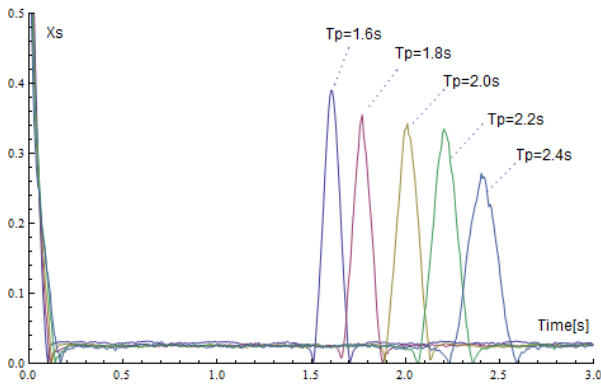


Fig. 9. Period detection by active probing

For this scenario, our method for detecting the periodic nature in the video traffic is quite accurate. The peaks are easy to detect both by means of graphical views and by pure computation.

When we focus in on the lower range of the time axis, we also see the presence of serial correlation between probe packet IAT observations within a burst (cf. Fig. 10). The T_b values based on passive measurements from Table I have been included in the illustration, for each of the T_p cases.

We observe that the serial correlation goes to zero at the point where we have reached the burst duration. The match between the passively measured T_b values, and those indicated by the points where the serial correlation reaches zero are very close.

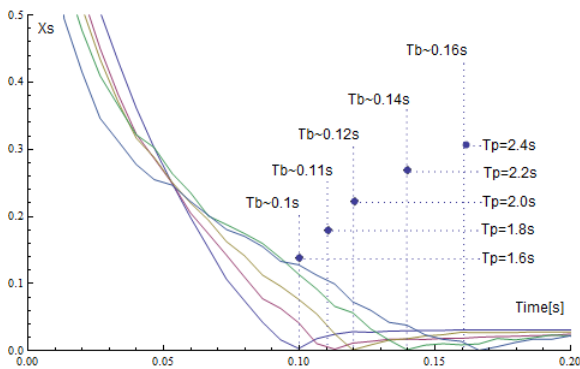


Fig. 10. Estimated burst duration by active probing

In this scenario, both the graphical view and a computational approach would be able to present estimators for T_b based on the output from serial correlation X_s .

B. Smooth 5Mbps Streaming on 25Mbps Access

In the majority of scenarios investigated, the access capacity represents a limiting factor with regard to potential burst rate for the video stream. These are the cases where the access speed is lower than the interface speed on the video server side. This could be considered as the typical scenario in a real use case.

In Fig. 11, the serial correlation results obtained when using a 25Mbps access are shown. For the sake of clarity, only X_s output when using two different T_p values (2.2 and 2.4) are shown.

The peaks in correlation values for the two T_p cases (2.2s and 2.4s) are quite clear. The side lobes surrounding the center peak are in line with the correlation results for the theoretical signal as shown in Fig. 6 and 7. The side lobes

indicate that the main burst period contains sub-burst periods within it. These sub-bursts could be associated with the underlying TCP mechanisms carrying the HTTP encapsulated video stream. When we focus in on the lower range of the time axis (cf. Fig. 12) we see the same type of side lobes with decreasing peak value.




Fig. 11. Period detection by active probing




Fig. 12. Estimated burst duration by active probing

Following the same approach as for the 100Mbps access case, i.e. searching for the point in time when X_s reaches zero is not enough in this case. The analysis must take all the lobes into consideration, and look for X_s reaching zero after the last side lobe. Obviously, this works well based on the graphical view of X_s but could introduce some challenges in a pure computational approach. However, the potential level of precision in burst duration estimation is quite good.

C. Netflix Streaming on 25Mbps and 100Mbps Access

Measurements using the Netflix live stream in both 2Mbps and 3Mbps quality levels were done, over the full range of access capacity levels. The findings were similar for all access capacity levels, except for at the highest level of 100Mbps. In the presentation of the results (cf. Fig. 13) we only show the output from the serial correlation for the 3Mbps stream, when using access capacities of 25Mbps and 100Mbps.

The difference in X_s output noticed at 100Mbps access, is that we instead of a single peak around 4sec get two peaks surrounding this value. Although not studied in depth, the source of this dual peak output is most likely the presence of two TCP sessions. At the highest access level, it seems as if the combined traffic pattern generated by the TCP sessions changes nature. This could be further studied by means of e.g. estimated probability density functions, but this is considered outside of the scope for this paper.

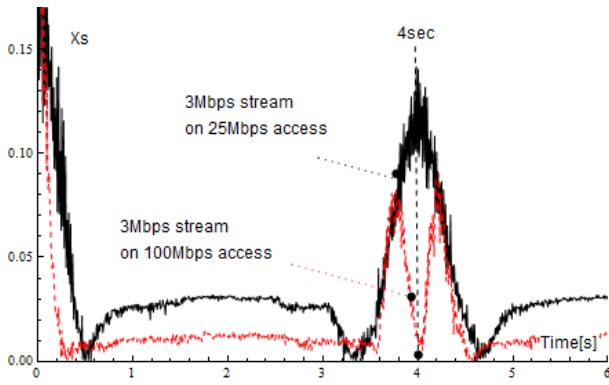


Fig. 13. Period detection by active probing

In order to see if we are able to detect the burst duration for a Netflix stream using two TCP sessions, we focus on the lower range of the time axis for the X_s output (cf. Fig. 14). The average values for the real burst durations at 25Mbps and 100Mbps (Ref. Section III) are indicated at 0.45s and 0.21s.

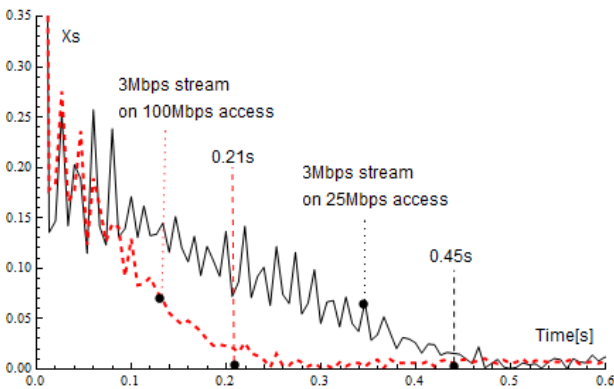


Fig. 14. Estimated burst duration by active probing

The X_s output for both 25Mbps and 100Mbps access indicate burst durations of respectively 0.48s and 0.25s. These estimates are a little bit high, but keeping in mind that we are probing the aggregate of two TCP sessions, the results are quite promising.

D. Active Probing Rate

The amount of probe traffic required in order to detect and estimate the period and burst durations in an accurate manner is of interest and concern. The default probe pattern used in our experiments (cf. Fig. 5) was a sequence of 100 byte packet pairs, with a fixed 0.5ms gap between the packets (t_{in}) and 6.1ms gap between the packet pairs (t_{gap}). This gives a probe packet rate of 300pps, corresponding to about 240Kbps. However, it is important to note that it takes 2 probe packets (cf. Fig. 5) to produce one sample input to the computation of X_s . Thus, the probing rate is half of the probe packet rate.

To study the effect of changing the amount of probe traffic, we performed experiments with probe packet rates down to 160pps and up to 700pps (cf. Fig. 15), by changing the gap between packet pairs (t_{gap}). The findings for both the video service in lab and Netflix were similar, thus we only present the results for Netflix on a 25Mbps access link as illustration.

By investigating the serial correlation output X_s for Netflix (3Mbps level) for the purpose of detecting the period, both the lowest and highest probe packet rates give the same result. The period is detected to be around 4sec, as indicated by the peak in the X_s plots. The main difference between the plots

used to detect the period is that a higher probe packet rate gives a smoother curve but at the same a less distinct peak.

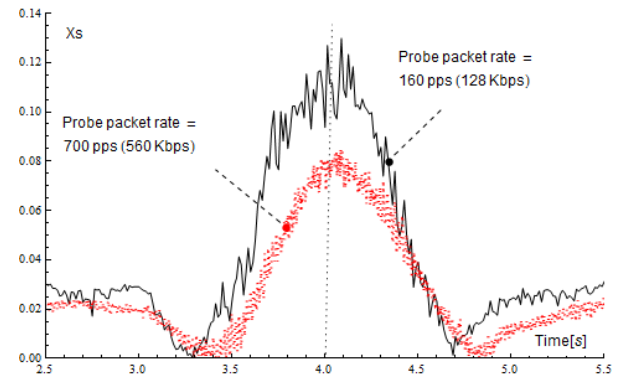


Fig. 15. Period detection, High/Low probing

Focusing in on the lower range of the time axis for the X_s output for respectively high and low probe packet rates (cf. Fig. 16), we see that the same conclusion can be drawn with regards to burst duration based on high and low probe packet rates. However, the higher probe packet rate seems to indicate a higher value for burst duration than the lower. The average value for burst duration, captured using passive measurement (cf. Section III) for Netflix operating at 3Mbps quality level in our measurements was 0.45sec per TCP session. Thus, both the higher and lower probe packet rates give an estimate for the burst duration which is somewhat high, i.e. between 0.55-0.65s.

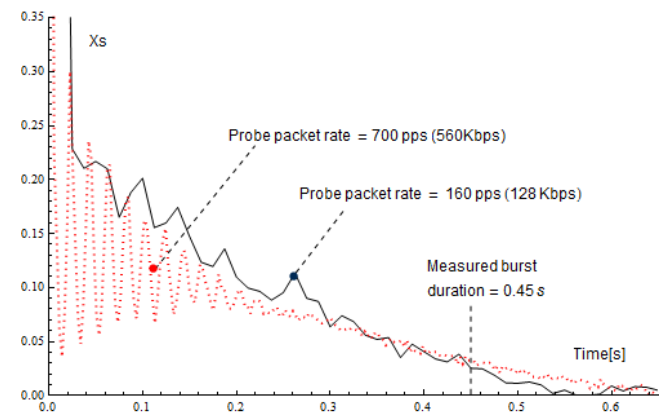


Fig. 16. Estimated burst duration, High/Low probing

Keeping in mind that the Netflix stream contains two TCP sessions, and that the active probing is influenced by both of these the results for burst duration detection are considered quite good.

E. Sample Size

When performing probing of real traffic it is important to reach a state where estimators for the parameters of interest can be presented as fast as possible. In our case, where serial correlation on a time series of observed packet pair gaps is performed, we are concerned about how long the time series must be in order for our estimators of both period and burst duration to appear. Obviously, detecting a periodic pattern in a time series requires us to at least study a time series of length greater than the repeating pattern.

A range of experiments were performed for the video streams available (Smooth Streaming and Netflix) over the full range of access capacities. As illustration of the findings,

the results for 5Mbps Smooth Streaming on a 25Mbps access is shown in Fig. 17.

The findings for this specific case was that a sample size of about 10sec was required in order to get a clear peak in the X_s output at the correct time value ($T_p=2.0$). If the sample size was reduced to only 5 seconds, the peak in X_s output is shifted down to 1.65 seconds, which is not correct.




Fig. 17. Period detection for different sample sizes

The minimum sample size varies depending on the type of video stream and access capacity, but in all scenarios investigated in our work a sample size of 10 seconds of probe traffic was sufficient for detecting both period and burst duration.

VII. CONCLUSIONS

The results from our experimental evaluation of the suggested method for detection of periodic behavior and burst duration in cross traffic on access links are considered quite promising. The use of video as the cross-traffic service component in our experiments strengthens the significance of our findings due to the popularity of such services on the Internet.

The use of serial correlation as a tool for analyzing a time series of observations is clearly a strong approach for the purpose of detecting periodic behavior. In all our experiments, the match between results from passive measurements (TCPdump) and active measurements (probing) have been very good. This applies to both the Smooth Streaming lab service, and also the live Netflix service. In the lab scenario, the addition of CBR background traffic does not change these findings.

Analyzing the output of the serial correlation in order to estimate the burst duration is more challenging when the access capacity is reduced, and the side lobes of the serial correlation appears. In these cases, a certain error margin should be expected. However, it is clear that the information is available in the serial correlation output.

The amount of probe traffic required for the method to perform well is considered to be acceptable. The experiments showed that using probe rates down to 128Kbps gave both period and burst information with reasonable accuracy. The results also showed that probing too much did not improve the accuracy to such an extent that it should be considered worth the additional bandwidth.

The required sample size for the method to give estimators for period and burst duration was only subject for a basic

evaluation. However, even in this domain the initial findings were positive. A required sample size of around 10 seconds does not directly qualify the method for all purposes. However, one should keep in mind that the periodic components of the services used were in the order of seconds. Therefore, it remains an open issue how fast this type of method can be expected to operate.

VIII. FUTURE WORK

The use of a more complex cross-traffic mix than just video services and CBR traffic is interesting to study. How well our method performs in such cases is an open issue and should therefore be investigated further.

Finding the optimal probe traffic pattern (packet pairs, packet trains etc.) was not subject for study in our work. There is most likely room for further improvements if this topic is investigated in more depth.

In the analytical part of our work we have relied a lot on the manual interpretation of graphical views of the output from serial correlation. If the method is considered for real implementation, some more work must be done in order to develop a more computational oriented approach.

Using this type of cross-traffic characterization as our method provides is also considered interesting to use as input to algorithms for available bandwidth estimations. We plan to investigate this topic for the purpose of developing a new method of performing available bandwidth estimation by means of stratified probing.

REFERENCES

- [1] C. Systems, *The Zettabyte Era*, Whitepaper, pp. 1-17, May 2012.
- [2] I. Sodagar, "The MPEG-DASH standard for multimedia streaming over the internet," *Multimedia, IEEE*, vol. 18, no. 4, pp. 62-67, April 2011.
- [3] P. Brockwell and R. Davis, *Time Series: Theory and Methods*, Springer, vol. 1, pp. 1-567, 2009.
- [4] P. Lewis and E. Orav, *Simulation Methodology for Statisticians, Operations Analysts, and Engineers*, vol. 1, pp. 1-398, Chapman and Hall/CRC, 1988.
- [5] R. Yarlagadda, *Analog and Digital Signals and Systems*, Springer, vol. 1, pp. 1-569, 2010.
- [6] R. Prasad, C. Dovrolis, M. Murray, and K. Claffy, "Bandwidth estimation: metrics, measurement techniques, and tools," *Network, IEEE*, vol. 17, no. 6, pp. 27-35, Nov.-Dec. 2003.
- [7] E. Goldoni and M. Schivi, "End-to-end available bandwidth estimation tools, an experimental comparison," in *Traffic Monitoring and Analysis*, Springer Berlin Heidelberg, 2010, vol. 6003, pp. 171-182.
- [8] Y. Cheng, V. Ravindran, and A. L. Garcia, "Internet traffic characterization using packet-pair probing," in *Proc. INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, May 2007, pp. 1766-1774.
- [9] M. Jain and C. Dovrolis, "Ten fallacies and pitfalls on end-to-end available bandwidth estimation," in *Proc. the 4th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '04. New York, NY, USA: ACM, 2004, pp. 272-277.
- [10] E. Bergfeldt, S. Ekelin, and J. M. Karlsson, "Real-time available-bandwidth estimation using filtering and change detection," *Computer Networks*, vol. 53, no. 15, pp. 2617-2645, 2009.
- [11] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On dominant characteristics of residential broadband internet traffic," in *Proc. the 9th ACM SIGCOMM Conference on Internet Measurement Conference*, ser. IMC '09. New York, NY, USA: ACM, 2009, pp. 90-102.
- [12] H. Jiang and C. Dovrolis, "Why is the internet traffic bursty in short time scales," in *Sigmetrics*, ACM Press, 2005, vol. 1, pp. 241-252.
- [13] A. Rao, A. Legout, Y.-S. Lim, D. Towsley, C. Barakat, and W. Dabbous, "Network characteristics of video streaming traffic," in *Proc. the Seventh Conference on Emerging Networking EXperiments*

- and Technologies, ser. CoNEXT '11. New York, NY, USA: ACM, 2011, pp. 1-12.
- [14] B. J. Villa and P. E. Heegaard, "Group based traffic shaping for adaptive http video streaming," presented at the 27th IEEE International Conference on Advanced Information Networking and Applications (AINA-2013), Barcelona, Spain, March 2013.
- [15] D. R. Figueiredo, B. Liu, V. Misra, and D. Towsley, "On the autocorrelation structure of tcp traffic," *Computer Networks*, vol. 40, no. 3, pp. 339-361, 2002.
- [16] K. Suh, D. R. Figueiredo, J. Kurose, and D. Towsley, "Characterizing and detecting skype-relayed traffic," in *Proc. IEEE INFOCOM '06*, pp. 1-12, April 2006.
- [17] B. Kulapala, P. Seeling, and M. Reisslein, "Comparison of traffic and quality characteristics of rate-controlled wavelet and det video," in *Proc. ICCCN*, 2004, pp. 247-252.
- [18] A. Rizk, Z. Bozakov, and M. Fidler. (2012). H-probe: estimating traffic correlations from sampling and active network probing. [Online]. pp. 1-18. Available: <http://arxiv.org/abs/1208.2870>.
- [19] B. Villa and P. Heegaard, "Improving perceived fairness and QoE for adaptive video streams," in *Proc. ICNS 2012*, pp. 149-158, March 2012.
- [20] B. J. Villa, P. E. Heegaard, and A. Instefjord, "Improving fairness for adaptive http video streaming," in *Proc. the 18th EUNICE/IFIP*, vol. 1, pp. 183-193, June 2012.
- [21] A. Zambelli, "IIS Smooth streaming technical overview," Tech. Rep., March 2009
- [22] S. Akhshabi, A. C. Begen, and C. Dovrolis, "An experimental evaluation of rate-adaptation algorithms in adaptive streaming over http," in *ACM Multimedia Systems (MMSys)*, vol. 1, pp. 157-168, Feb. 2011.
- [23] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The click modular router," *ACM Trans. Comput. Syst.*, vol. 18, no. 3, pp. 263-297, Aug. 2000.
- [24] J. Laine, S. Saaristo, and R. Prior. (Jan 2000). Real-time udp data emitter (rude) and collector for rude (crude). [Online]. Available: <http://rude.sourceforge.net/>.



Bjørn J. Villa received his MSc degree in telematics at NTNU in 1995. After 15 years in the industry, he is now working on his PhD at NTNU. His research interest cover methods for optimizing QoS and QoE for Internet services, and he has published a number of research papers in this domain. A key service used in the research has been adaptive video streaming.



Poul E. Heegaard received his MSc degree in telematics at NTNU in 1989, and his PhD degree in 1998. Heegaard is a Professor at ITEM/NTNU, and since 2009 he has been the head of the department. His research interests cover performance, dependability and survivability evaluation and management of communication systems.