

Impact of Cross Traffic Burstiness on the Packet-scale Paradigm

Rebecca Lovewell and Jasleen Kaur
University of North Carolina at Chapel Hill
{lovewell, jasleen}@cs.unc.edu

Abstract—The packet-scale paradigm is a novel framework for achieving ultra-high speed congestion control. Due to its reliance on finely-controlled inter-packet gaps, the paradigm is expected to be sensitive to transient burstiness in traffic encountered on bottleneck links. This paper uses a first-principles approach to study the impact of cross traffic burstiness on the efficiency of the packet-scale paradigm. It relies on a simple periodic on-off model for cross traffic and studies the interaction of the burstiness timescale, round-trip times, and the smoothing filters adopted by the paradigm. The analysis is validated against ns-2 simulations with a prototype. Our analysis helps gain fundamental insights on the impact of several factors.

I. INTRODUCTION

The packet-scale paradigm is a recently proposed framework for congestion control that promises to scale up the state of the art in ultra-high speed congestion control by several orders of magnitude [1], [2].¹ A key feature of the paradigm is its use of short probing-streams to probe for available network bandwidth at fine timescales, using finely-controlled inter-packet spacings—this gives the paradigm its characteristic properties of scalability, adaptability to dynamic bandwidth, and queue-friendliness. Unfortunately, the excellent adaptability of the paradigm to dynamically varying bandwidth also raises significant concern about its sensitivity to transient short-scale burstiness in the cross traffic encountered on congested links. In particular, before this paradigm can be deployed in practice, it is important to ask: *to what extent does cross traffic burstiness impact the efficiency of the packet-scale paradigm?* In this paper, we use a first-principles analysis approach to partially study this issue.

Most congestion control protocols are analyzed for their steady-state throughput by conducting stochastic analysis of their window growth functions [3], [4], [5]. Most of these analyses incorporate the impact of packet losses and round trip-times (RTTs), considering these to be the major influences on the throughput of a transfer. In contrast, the packet-scale paradigm does not rely on window-based rate-control and relies on fundamentally different mechanisms for controlling the rate of a transfer. Furthermore, due to its agility at fine timescales, the paradigm is mostly limited by dynamics in the bottleneck queues (which, in turn, influence the available bandwidth, avail-bw). In particular, (i) by design, the steady-state throughput of long transfers is independent of path RTT;

and (ii) due to its queue-friendly behavior, packet losses are expected to occur rarely (and are not expected to play a dominant role in limiting throughput). Consequently, existing analysis approaches are not suitable for studying the paradigm.

In this paper, we instead adopt a first-principles approach to analyze the impact of fine-scale traffic burstiness on the paradigm.² Our analysis considers a simple periodic on-off model for the cross traffic burstiness (parametrized by the burstiness timescale) and incrementally incorporates the impact of queue backlogs, RTTs, and the averaging timescales of the paradigm. We validate our analysis against experiments conducted with an ns-2 implementation of a prototype of the paradigm. Our analysis shows that:

- The degree to which traffic burstiness impacts the efficiency of the paradigm depends on both the average cross traffic load, as well as its characteristic ratio of burst-to-idle timescales.
- RTTs have little influence on the worst-case impact of traffic burstiness on the paradigm.
- The impact of cross-traffic burstiness on the paradigm can be alleviated by selecting the two smoothing timescales used by the paradigm such that: (i) these are reasonably larger than the typical burstiness timescale of the cross-traffic encountered, and (ii) the difference between the two smoothing timescales is small.

In the rest of this paper, we formulate the analysis framework in Section II. We develop and validate the analysis in Sections III-VI. We use the analysis to quantitatively study the performance of the paradigm under different conditions and settings in Section VII. We conclude in Section VIII.

II. FORMULATION

The available bandwidth of a link, observed during a time interval (t_1, t_2) , is defined as:

$$B_{avail}(t_1, t_2) = C - \frac{B_{CT}(t_1, t_2)}{t_2 - t_1} \quad (1)$$

where C is the transmission capacity of the link, and $B_{CT}(t_1, t_2)$ is the total amount of traffic that arrives on the link during the interval (t_1, t_2) . The avail-bw of a link reflects its spare capacity. The end-to-end avail-bw of a path is defined as the minimum avail-bw among its constituent links [6].

¹This material is based upon work supported by the National Science Foundation under Awards CNS-1018596 and OCI-1127413.

²In this paper, we ignore dynamics that occur at timescales even smaller than a probing-stream and focus only on larger timescales.




Fig. 1. Periodic On-off Cross Traffic Model

The packet-scale paradigm attempts to measure the end-to-end avail-bw and adjust the sending rate of the transfer to match it. The Rapid protocol described in [1] is a prototype of the packet-scale paradigm. Below, we briefly describe the key mechanisms that control Rapid behavior in steady-state—for brevity, we provide only details that are relevant to the discussion in this paper. In particular, since our analysis focuses on timescales larger than probing-streams, our description is light on mechanisms that operate at sub-probing-stream timescales.

A. Rapid: Operation in Steady-state

The Rapid sender continuously sends back-to-back probing-streams (p-streams) of length N packets, where each packet is sent at a potentially different probing rate by controlling the inter-packet gaps. The Rapid receiver records the arrival time of packets at the receiver, and returns these to the sender. The sender compares the receiver-gaps to the original gaps and computes an estimate for the avail-bw according to the *highest probing rate for which the inter-packet gaps did not show an increasing trend*. Details of the bandwidth-estimation logic and the structure of the p-streams can be found in [1]. In this paper, we ignore packetization effects and assume fluid behavior at p-stream and lower timescales—this implies that a Rapid p-stream returns a perfect estimate of the avail-bw.

The avail-bw estimate, B_{est} , computed by the sender is then fed through a moving average filter that updates the average sending rate, \hat{R} , used for the next p-stream as follows³:

$$\begin{aligned} \hat{R} &= \hat{R} + \frac{L}{\tau}(B_{est} - \hat{R}), & \text{if } \hat{R} > B_{est}, \\ \hat{R} &= \hat{R} - \frac{L}{\eta}(\hat{R} - B_{est}), & \text{otherwise} \end{aligned}$$

where $L = \frac{N * P}{\hat{R}}$ is the duration in time of the most recent p-stream, P is the packet size, and τ and η are the smoothing parameters employed by Rapid for increasing or decreasing the sending rate, respectively. Effectively, by employing the above moving average filter, when avail-bw increases (or decreases), Rapid linearly increases (or decreases) its average sending rate to the new avail-bw over τ (or η) time units.

B. Cross Traffic Model

According to Equation (1), the avail-bw is determined by the cross traffic arrival process. Note that the avail-bw influences a Rapid p-streams via dynamics in the (shared) packet queue at the bottleneck link. In particular, the queue-size grows when the collective bit-rate of the p-streams and the arriving cross traffic exceeds the link capacity—it is precisely this growth in the bottleneck queue that Rapid relies on for estimating the avail-bw.

³Note that this description is different from that of the Rapid protocol in [1], which does not use η (or equivalently, sets: $\eta = L$).




Fig. 2. Topology for Experimental Validation

Bursty cross traffic creates additional (transient) queue dynamics. Since Rapid is capable of estimating and adapting to avail-bw at fine timescales, it also reacts to the transient queue build-up caused by the cross traffic bursts. In order to study the impact of such bursts on Rapid p-streams and the steady-state throughput it achieves, we consider a simple model of the cross traffic in this paper—a periodic on-off source. Such a traffic source, illustrated in Fig 1, is characterized by the fixed duration of each burst, T_{on} , the inter-burst idle time, T_{off} , and the on-rate, R_{on} . The average cross traffic rate is then give by: $R_{avg} = \frac{R_{on}T_{on}}{T_{on}+T_{off}}$.

Note that for a given R_{avg} , traffic can be made more “bursty” by increasing R_{on} , and the timescale of burstiness can be controlled using T_{on} . Given this model for cross traffic, we analyze Rapid using the following metrics.

C. Metrics

In the ideal absence of burstiness in cross traffic, Rapid should be able to fully utilize the steady-state avail-bw [1]. That is, the *steady-state Rapid throughput*, \mathcal{R} , should be equal to $C - R_{avg}$. In order to see how cross traffic burstiness impacts Rapid throughput, in our analysis we study both \mathcal{R} and the overall link utilization, given by:

$$\mathcal{U} = \frac{R_{avg} + \mathcal{R}}{C} \quad (2)$$

D. Experimental Setup

We use the ns-2 implementation of the Rapid protocol to validate our analysis against. We generate a simple dumbbell topology illustrated in Fig 2. The Rapid transfer and the cross traffic share the 1 Gbps bottleneck link. All other links have a transmission capacity of 10 Gbps. The RTT of the Rapid transfer is controlled by varying the propagation delays on the links attached to the Rapid sender and the Rapid receiver. Sufficient queues are provisioned on all links to avoid packet losses. The cross traffic is driven by a periodic CBR on-off source—the source alternates between an *on* state, in which it send traffic at a fixed bit-rate of R_{on} for an interval of fixed length T_{on} , and an *off* state in which it does not send any traffic for an interval of fixed length T_{off} . The cross traffic parameters are varied across experiments.

III. IMPACT OF QUEUE BACKLOG

We begin our analysis with a simple but important observation—which will be used in subsequent analysis—




Fig. 3. Impact of RTT: Illustration

about the influence of a *residual* queue⁴ at the bottleneck link on the avail-bw estimates yielded by a Rapid p-stream:

Observation 1: The presence of a residual queue at the bottleneck link does not influence the avail-bw estimates yielded by an arriving Rapid p-stream—the avail-bw estimates in this case will simply depend on the “arrival” process of the cross traffic.

This is not an immediately obvious observation, since the packets of a p-stream that arrives behind a residual queue are likely to bunch up together while the residual queue drains. While this is certainly true, the inter-packet gaps at the time of *departure* will still be shaped according to the cross traffic that arrives *along with* the p-stream (and which gets inserted in between the p-stream packets)—the departure gaps are independent of the amount of residual queue encountered by the arriving p-stream. In Appendix IX, we formally establish this for the case when two successive probing packets belong to the same *busy period* of the router link.


IV. ROLE OF RTT

A Rapid sender learns of a change in the avail-bw only after a feedback delay of RTT time units. We first investigate the role that this feedback delay might play in influencing Rapid’s efficiency in utilizing the avail-bw. In order to solely focus on this factor, we set: $\tau = \eta = L$ — this ensures that the Rapid sender sets its sending rate equal to a new avail-bw estimate immediately upon learning about it (no smoothing/averaging delay is involved).

Fig 3 illustrates a typical timeline, which includes the periodic on-off cross traffic arrival rate, the Rapid sending rate, and the queue buildup at the bottleneck link. The cross traffic alternates between bit-rates of R_{on} and 0. Observe that:

- When a cross traffic burst arrives at the bottleneck link at time t_0 , it interacts with the arriving Rapid p-streams—the Rapid sender learns that the avail-bw has decreased to $C - R_{on}$ after a delay of 1 RTT at time: $t_0 + RTT$. The Rapid sender immediately reduces its sending rate from C to $C - R_{on}$ at this time.
- The cross traffic burst lasts till $t_0 + T_{on}$. Note that the queue-backlog at the bottleneck link keeps increasing during the time interval $[t_0, t_0 + T_{on}]$ (since the sum of Rapid and cross traffic arrival rates exceeds C). However, because of Observation 1, the Rapid p-streams arriving in this interval yield avail-bw estimates equal to $C - R_{on}$.

⁴The residual queue refers to buildup which already exists on the queue before the first packet of a Rapid p-stream arrives.



(a) No Queue Buildup

(b) Max Queue Buildup

Fig. 4. Best-case and Worst-case Queue Buildsups

- The Rapid p-streams that arrive at the bottleneck link immediately after $t_0 + T_{on}$ will yield avail-bw estimates equal to C (despite the presence of a residual queue, from Observation 1). However, these p-streams will experience an inflated delay in reaching the receiver (due to the extra queuing delay waiting for the residual bottleneck queue to drain). In fact, the *first* notification that the avail-bw has now increased to C will reach the sender only after time: $t_0 + T_{on} + RTT + D$, where D is the queue drain time of the maximum residual queue, given by: $D = \frac{Q_{max}}{C}$. Therefore, the Rapid sender will increase its sending rate back to C only at time $t_0 + T_{on} + RTT + D$.
- This whole interaction repeats with the arrival of the next cross traffic burst.

The average Rapid throughput can be computed by averaging its sending rate over the interval $[t_0, t_0 + T_{on} + T_{off}]$ as:

$$\begin{aligned} \mathcal{R} &= \frac{C \cdot (T_{off} - D) + (C - R_{on}) \cdot (T_{on} + D)}{T_{on} + T_{off}} \\ &= C - R_{avg} \frac{T_{on} + D}{T_{on}} \end{aligned} \quad (3)$$

Note that while RTT does not appear directly in the above term, it may indirectly influence the maximum queuing delay, D , which adversely impacts the Rapid throughput.

D , in turn, depends on the phase lag between the cross traffic bursts and Rapid’s response. This lag is governed by: $RTT_{rem} = remainder(\frac{RTT}{T_{on} + T_{off}})$. For instance, when $RTT_{rem} = 0$ (as illustrated in Fig 4(a)), there is *no* queue buildup at the bottleneck link since the cross traffic peaks overlap perfectly with the Rapid troughs. Thus, $Q_{max} = 0$ and $D = 0$. On the other hand, Fig 4(b) illustrates that the worst-case queue buildup occurs when the peaks and troughs do not overlap in time at all (for instance, when $RTT_{rem} > T_{on}$ and $T_{off} > T_{on}$). In this case, $Q_{max} = T_{on}R_{on}$, and $D = \frac{T_{on}R_{on}}{C}$.

Note that when the peaks and troughs overlap only partially (for instance, when $RTT_{rem} < T_{on}$), the queue-buildup is smaller than the worst-case and is proportional to RTT_{rem} . Since our objective is to quantify the extent to which burstiness can adversely impact Rapid performance, we focus only on the worst-case queue buildup, for which:

$$\begin{aligned} \mathcal{R} &= C - R_{avg} - \frac{R_{avg}^2 T_{on} + T_{off}}{C} \\ \mathcal{U} &= \frac{R_{avg} + \mathcal{R}}{C} = 1 - \frac{R_{avg}^2 T_{on} + T_{off}}{C} \end{aligned}$$

This implies that the role of RTT is limited to influencing the phase shift between the cross traffic peaks and Rapid troughs. *The worst-case loss in utilization is independent of the path RTT.* Henceforth, we focus only on analyzing the worst-case queue buildups.

V. LARGE-SCALE BURSTS ($T_{on} + D \geq \eta$, $T_{off} - D \geq \tau$)

Note that the parameters τ and η introduce additional delay before Rapid completely adapts to an increase or decrease in avail-bw. For studying the influence of τ and η , we first consider the case in which the cross traffic bursts last long enough to allow Rapid to converge to the avail-bw changes. More precisely, when $T_{on} + D \geq \eta$ and $T_{off} - D \geq \tau$, the Rapid sender will always have enough time to completely adapt to a change in avail-bw, before it changes again (see Fig 5). Observe again:

- At t_0 , a cross traffic burst arrives and the queue starts to build up. Rapid learns of the corresponding decrease in avail-bw after a delay of 1 RTT. It then takes η time units to converge its sending rate to the lower value of $C - R_{on}$.
- The cross traffic burst stops at $t_0 + T_{on}$. Since the queue builds up in the meantime, however, Rapid learns of the corresponding increase in avail-bw only after a delay of $RTT + D$, where D is the time it takes for the residual queue to drain: $D = \frac{Q_{max}}{C}$. It then takes τ time units to increase its sending rate from $C - R_{on}$ to C . All of this occurs before Rapid learns of another decrease in avail-bw.
- A worst-case queue buildup is caused by worst-case phase alignment of cross traffic peaks and Rapid peaks. When such an alignment, the value of Q_{max} further depends on the traffic and smoothing parameters—detailed analysis of the associated conditions is presented in [7]. Here, we focus on only the very worst of these conditions, in which: $D \leq T_{off} - T_{on} - \tau$. Under this condition, Rapid is capable of sending data at a rate of C throughout a cross-traffic burst. When this happens, we have: $Q_{max} = R_{on}T_{on}$ and $D = \frac{R_{avg}}{C}(T_{on} + T_{off})$.

The steady-state Rapid throughput can be time-averaged as:

$$\begin{aligned} \mathcal{R} &= \left(C - \frac{R_{on}}{2}\right) \frac{\tau + \eta}{T_{on} + T_{off}} + C \cdot \frac{T_{off} - \tau - D}{T_{on} + T_{off}} \\ &+ (C - R_{on}) \frac{T_{on} - \eta + D}{T_{on} + T_{off}} \\ &= C - R_{avg} - \frac{R_{avg}^2}{C} \frac{T_{on} + T_{off}}{T_{on}} - R_{avg} \cdot \frac{\tau - \eta}{2T_{on}} \\ \mathcal{U} &= 1 - \left(\frac{R_{avg}}{C}\right)^2 \frac{T_{on} + T_{off}}{T_{on}} - \frac{R_{avg}}{C} \cdot \frac{\tau - \eta}{2T_{on}} \quad (4) \end{aligned}$$

Validation We validate the above analysis against an ns-2 experiment in which: $T_{on} = 100ms$, $T_{off} = 350ms$, $R_{avg} = 111Mbps$, $\tau = 150ms$, $\eta = 75ms$, and $RTT = 125ms$. Fig 6 plots a sample of the theoretical and experimentally observed time series for the Rapid sending rate. We find that the experiment validates our fluid analysis quite well. We have




Fig. 5. Impact of Large-scale Bursts: Illustration




Fig. 6. Large-scale Bursts: Experimental Validation

also conducted several additional experiments (i) by varying the above parameters, and (ii) by using an exponential on-off cross traffic source —our measured values for link utilization match the predicted one very well in *all* cases.

VI. SMALL-SCALE BURSTS ($T_{on} < \eta$, $T_{off} < \tau$)

We next consider the case in which the cross traffic bursts are fairly short and Rapid is unable to adapt to a change in avail-bw, before the latter changes again—this happens when $T_{on} + D < \eta$ and/or $T_{off} - D < \tau$. Fig 7 illustrates a case for which both the conditions hold. In this case,⁵ when avail-bw increases, the Rapid sender is able to reach a max sending rate of only $R_H < C$; and when avail-bw decreases to $C - R_{on}$, Rapid is able to reduce its sending rate to a min value of only $R_L > C - R_{on}$ before the avail-bw increases again. In steady-state, we have:

$$\begin{aligned} R_L &= R_H - (R_H - (C - R_{on})) \left(\frac{T_{on} + D}{\eta}\right) \\ R_H &= R_L - (C - R_L) \left(\frac{T_{off} - D}{\tau}\right) \end{aligned}$$

The worst-case queue buildup occurs when $RTT_{rem} = \frac{T_{on}}{T_{on} + T_{off}}(T_{off} - D)$, such that Rapid peaks are in the worst-case phase alignment with the cross traffic peaks. In this case, the maximum queue buildup can be derived as:

$$Q_{max} = T_{on} \left(\frac{R_H T_{off} - R_L T_{on}}{2(T_{on} + T_{off})} + \frac{R_H}{2} - C + R_{on} \right)$$

Recall that $D = \frac{Q_{max}}{C}$. The steady-state Rapid throughput can then be averaged as:

$$\mathcal{R} = \frac{R_L + R_H}{2}$$

⁵Due to space constraints, we include all other cases of small-scale bursts only in [7].




Fig. 7. Impact of Small-scale Bursts: Illustration




Fig. 8. Small-scale Bursts: Experimental Validation

We solve for the above quantities using mathematical software. The closed-form solution is too unwieldy to present here.

Validation We validate the above analysis against an ns-2 experiment in which: $T_{on} = 100ms$, $T_{off} = 200ms$, $R_{avg} = 111Mbps$, $\tau = 200ms$, $\eta = 175ms$, and $RTT = 100ms$. Fig 8 plots a sample of the theoretical and experimentally observed time series for the Rapid sending rate. We find that the experiment validates the analysis quite well. Our experiments with other parameter settings and with exponential on-off cross traffic also match the analysis equally well.

VII. DISCUSSION

We next use the models we have developed to quantitatively study the impact that cross traffic burstiness can have on Rapid and the extent to which it can be alleviated. Sections V-VI establish that the magnitude of τ and η relative to T_{on} and T_{off} determines whether the cross-traffic interacts with Rapid as “large-scale” or “small-scale” bursts. Further, Equation (4) shows that the inability of Rapid to fully utilize the bottleneck link due to large-scale cross traffic bursts: (i) does not depend individually on τ or η , but rather on $\tau - \eta$, the difference between the two smoothing parameters; (ii) is proportional to the relative load of the cross traffic, $\frac{R_{avg}}{C}$; (iii) is inversely proportional to T_{on} ; and (iv) is proportional to the ratio: $\frac{T_{off}}{T_{on}}$. While the small-scale bursts analysis yields closed-form expressions that are not as simple, we use mathematical software to quantify the influence of the above factors and find that these are still the dominant ones.

Role of $\tau - \eta$: Fig 9(a) plots \mathcal{U} as a function of $\tau - \eta$. C is set to 1 Gbps and the different curves correspond to different values of T_{on} , T_{off} , and R_{avg} .⁶ $\tau \in \{50ms, 500ms, 1s\}$, while η is varied over $[1ms, \tau]$. We find that in all cases:

- \mathcal{U} decreases with increase in $\tau - \eta$.

⁶Note that R_{on} can not be larger than C —this limits the maximum feasible choice of R_{avg} for a given $\frac{T_{off}}{T_{on}}$. Similarly, we have: $\tau - \eta \leq \tau$ —this limits the x-range in Fig 9(a).

- The loss in utilization due to a large $\tau - \eta$ is influenced somewhat, but not significantly, by the values of R_{avg} , T_{on} , and T_{off} .
- For a given value of $\tau - \eta$, a larger value of τ seems to alleviate the loss in utilization due to cross-traffic burstiness. However, the $\tau = 500ms$ and $\tau = 1s$ curves are fairly close to each other—this suggests that the influence of τ seems to be limited to whether the bursts are “large-scale” or “small-scale” in comparison to τ . We further explore this issue below.

Observe that τ and η are the only model parameters that are under the control of Rapid. The above figures indicate that independent of the cross traffic load and burstiness, these two parameters should be set such that their difference is small. In the rest of this discussion, we set: $\tau - \eta = 1ms$.⁷

Role of $\frac{R_{avg}}{C}, \tau$: Fig 9(b) plots \mathcal{U} as a function of $\frac{R_{avg}}{C}$.⁸ The different curves correspond to 4 different values of (T_{on}, T_{off}) and 2 values for τ ($\tau - \eta = 1ms$).⁹ We find that:

- \mathcal{U} decreases with increase in $\frac{R_{avg}}{C}$. This trend is more pronounced for smaller values of $\frac{T_{on}}{T_{off}}$.
- For a given combination of T_{on} and T_{off} , Rapid performs better when bursts are “small-scale” (τ, η are sufficiently larger than T_{on}, T_{off}). Further, as long as the bursts are small-scale, the value of τ does not influence Rapid performance much — in fact, the curves for $\tau = 250ms, 500ms$ (not plotted to avoid clutter) are fairly close to the $\tau = 1s$ for all of the cases plotted in this figure.

This suggests that even though Rapid performance is relatively independent of the individual values of τ and η (and depends mostly on $\tau - \eta$), it does get strongly influenced by whether or not these parameters are larger than the traffic bursts and idle durations. Using large values for τ and η can help ensure that bursts are relatively “small” in scale and Rapid performs well.

Role of $T_{on}, \frac{T_{on}}{T_{off}}$: Fig 9(c) plots \mathcal{U} as a function of $\frac{T_{on}}{T_{off}}$ — here, $T_{on} = 25ms, 100ms$ and T_{off} is varied to get different values for this ratio. $C = 1Gbps$ and the different curves correspond to different values of $\frac{R_{avg}}{C}$. We find that:

- The link utilization achieved increases with the ratio $\frac{T_{on}}{T_{off}}$. In most cases, the influence flattens out beyond a ratio of $\frac{T_{on}}{T_{off}} = 4$.
- For a given $\frac{T_{on}}{T_{off}}$, the impact of the individual choice of T_{on} is minimal. This is a very interesting observation—the timescale of cross traffic burstiness does *not* significantly influence Rapid’s performance (as long as the

⁷It can be shown that to ensure the stability of the paradigm, we need to select: $\eta < \tau$.

⁸ \mathcal{U} is independent of the individual values of R_{avg} and C , but is influenced by the ratio $\frac{R_{avg}}{C}$, for both large-scale and small-scale bursts. This is a significant property — it implies that the impact of burstiness does not change as we consider upcoming networks with higher transmission capacities of 40-100 Gbps.

⁹Note again that $R_{on} < C$ limits the maximum feasible choice of R_{avg} for a given $\frac{T_{off}}{T_{on}}$.

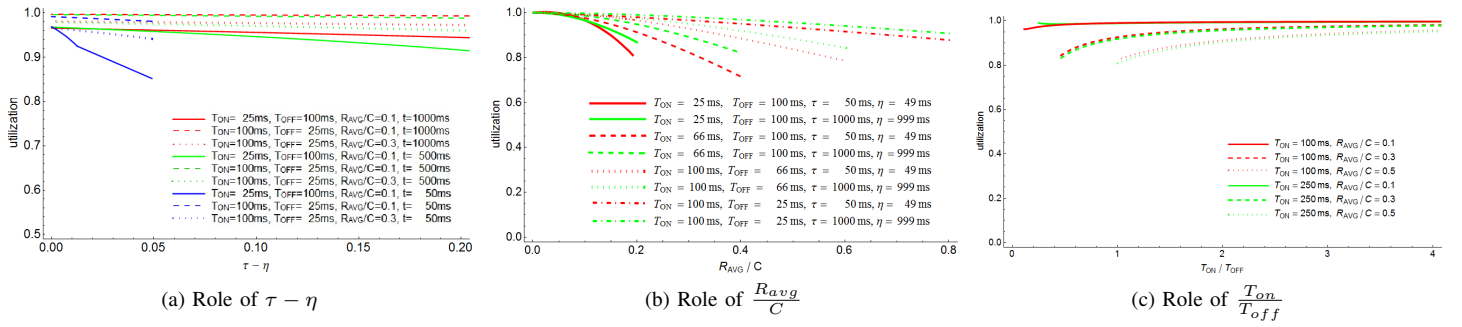


Fig. 9. Impact of Cross-traffic Bursts: Influence of Factors

smoothing timescales τ and η are larger), but the *relative* durations of burst and idle periods do!

Our use of the model in this section reveals that once the Rapid protocol parameters are configured ideally (smaller $\tau - \eta$, large τ), the impact of cross traffic is dependent mostly on the relative load of cross traffic, $\frac{R_{avg}}{C}$, and the relative burstiness-to-idleness ratio: $\frac{T_{on}}{T_{off}}$. With ideal parameter configuration, the ability of Rapid to full utilize the bottleneck link can vary from 80% to 95%, depending on the combination of these two factors.

VIII. CONCLUDING REMARKS

The packet-scale paradigm is a promising framework for ultra-high speed networks. However, it can deliver on its promise only if network “noise” does not adversely impact its bandwidth estimation process. This paper conducts the first analysis of the interaction of the packet-scale paradigm with one source of noise: traffic burstiness. The analysis relies on a first-principles approach to gain several simple but fundamental insights on the influence of network, traffic, and protocol parameters on the efficiency of the paradigm.

While this analysis is an important first step, it is only a small one in understanding the interaction of the packet-scale paradigm with sources that can introduce “noise” in the inter-packet gaps. We are currently extending this analysis in the following directions: (i) An important source of “noise” is non-bottleneck buffering, which occurs in almost all store-and-forward systems that packets pass through. This buffering occurs at very fine timescales—the packet-scale paradigm is uniquely likely to get impacted by this due to the fine timescales at which it operates. (ii) On a similar note, bottleneck burstiness that occurs at timescale smaller than probing-streams are likely to impact the paradigm in ways that this paper does not study. We are currently relying on packet-level queuing theory to study these issues. (iii) Obviously, the periodic on-off cross traffic model is unrealistic. However, we believe it has helped us understand the relative role of several factors and it is likely this will hold as we study more complex and realistic models for traffic burstiness.

REFERENCES

[1] V. Konda and J. Kaur, “RAPID: Shrinking the Congestion-control Timescale,” in *Proceedings of the IEEE INFOCOM*, April 2009.
 [2] “Packet-scale Congestion Control,” <http://rapid.web.unc.edu>.

[3] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, “Modeling TCP Throughput: A Simple Model and its Empirical Validation,” in *Proceedings of ACM SIGCOMM*, September 1998.
 [4] B. Sikdar, S. Kalyanaraman, and K. Vastola, “Analytic Models for the Latency and Steady-State Throughput of TCP Tahoe, Reno and SACK,” in *Proceedings of IEEE Globecom*, November 2001.
 [5] M. Weigle, L. Chen, J. Kaur, and V. Kulkarni, “Generalized Stochastic Performance Models for Loss-based Congestion Control,” *Computer Communications Journal*, vol. 33, pp. 513–525, March 2010.
 [6] C. Dovrolis and M. Jain, “End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput,” *IEEE/ACM Transactions in Networking*, August 2003.
 [7] R. Lovewell and J. Kaur, “Impact of Cross Traffic Burstiness on the Packet-scale Paradigm - Extended Analysis,” Department of Computer Science, University of North Carolina at Chapel Hill, Tech. Rep. TR11-007, September 2011.

IX. INFLUENCE OF A RESIDUAL QUEUE

We formally derive Observation 1 for the case when all packets of a p-stream belong to the same *busy period* of the router link.

Let a_1 and a_2 be the arrival times at the bottleneck link of two successive packets, p_1 and p_2 , of a p-stream. Let d_1 and d_2 be their departure times. Since p_1 and p_2 belong to the same busy period, we have:

$$d_2 - d_1 = \frac{P}{C} + \frac{B_{CT}(a_1, a_2)}{C} \quad (5)$$

where P is the packet size of p_2 , and $B_{CT}(a_1, a_2)$ is the amount of cross-traffic that arrives in the interval (a_1, a_2) (and gets inserted in between the two p-stream packets).

If p_2 is sent at a probing rate of r_2 , then we have: $a_2 - a_1 = \frac{P}{r_2}$. Let R_{CT} denote the average cross traffic arrival rate in the interval (a_1, a_2) . Then we have: $B_{CT}(a_1, a_2) = R_{CT}(a_2 - a_1)$. Thus, we have:

$$d_2 - d_1 = \frac{P}{C} \left(1 + \frac{R_{CT}}{r_2} \right) \quad (6)$$

Rearranging and solving, we get: $d_2 - d_1 > a_2 - a_1$ iff $r_2 > C - R_{CT}$. Hence, the relation between the departure gaps and the arrival gaps (and consequently, the avail-bw estimates yielded by the p-stream) depend only on the cross traffic that arrives along with the p-stream packets (and is independent of the amount of residual queue).