

Unresponsive Flows and AQM Performance

C.V. Hollot, Yong Liu, Vishal Misra and Don Towsley

Abstract—Routers handle data packets from sources unresponsive to TCP’s congestion avoidance feedback. We are interested in the impact these sources have on AQM’s control of long-lived TCP traffic. In this paper, we combine models of TCP/AQM dynamics with models of unresponsive traffic to analyze the effects on AQM performance.

I. INTRODUCTION

ACTIVE queue management (AQM) schemes are meant to regulate TCP traffic in response to router congestion. This is accomplished by dropping (or marking) packets as a function of queue length. This dropping amounts to feedback, wherein sources are instructed to modify their sending rates. However, this arrival stream may contain packets from sources that are unresponsive to this feedback. For example, while some TCP sources are long-lived and in congestion avoidance phase, other TCP flows, with shorter lifetimes, either never experience loss or in the case of ECN are unresponsive to dropping. Also, by design, UDP flows do not respond to loss unless an application layer congestion control mechanism is built. Consequently, some congesting flows are uncontrolled by AQM, and, this motivates us to analyze the effects and tradeoffs these unresponsive flows have on AQM performance.

A. Impact on Transient Behavior

A good starting point is to ask whether unresponsive flows impact AQM performance to a degree that warrants scrutiny. Recent studies [1] show that unresponsive sources, while contributing to about 70 – 80% of the Internet flows, account for only 10 – 20% of its byte volume. Can this small volume of unresponsive flows significantly impact AQM? We believe so, and our focus is trained on transient response, as opposed to mean behavior, as in [2], where unresponsive flows were studied under high levels of aggregation.

To substantiate this belief concerning network transients, let’s do some back-of-the-envelope computations and consider a single bottle-neck link handling both short-lived and long-lived TCP flows. For simplicity, assume homogeneous flows of which N_s are short-lived, and N_ℓ are long-lived flows. The long-lived flows are in congestion-avoidance phase, while the

short-lived flows never escape slow start.¹ Let W_s and W_ℓ be the congestion window sizes for the short and long flows respectively. Take κ to be the byte-volume ratio between long and short-lived traffic. For example, $\kappa = 4$ means that 80% of total byte-volume is contributed by long-lived flows. In steady state, κ is the ratio of arrival rates, which, in terms of window size is expressed by $N_s W_s \kappa = N_\ell W_\ell$, or

$$N_s = \frac{N_\ell W_\ell}{\kappa W_s}.$$

Now let’s consider the dynamic of “traffic-increase” in an uncongested network. If the loss probability p is small, the back-off probability of a TCP connection, with window size w , is approximately wp . Thus, for each round-trip time, $N_s(1 - W_s p)$ short-lived flows double their window sizes, while $N_\ell(1 - W_\ell p)$ long-lived flows increase their congestion windows by 1. In other words,

$$\frac{\Delta^\uparrow W_s}{\Delta^\uparrow W_\ell} = \frac{N_s(1 - W_s p)W_s}{N_\ell(1 - W_\ell p)} \approx \frac{W_\ell}{\kappa}.$$

Since $W_\ell \propto \sqrt{\frac{1}{p}}$, the above shows that short-lived TCP flows can dominate the dynamic of traffic-increase when congestion is low. For the case of decreasing traffic, the window back-off is either $w - 1$ (for TO loss) or $w/2$ (TD loss). We then have

$$\frac{\Delta_\downarrow W_s}{\Delta_\downarrow W_\ell} \leq \frac{N_s W_s p W_s}{N_\ell W_\ell p W_\ell / 2} = \frac{2W_s}{\kappa W_\ell}$$

showing that long-lived TCP flows dominate the dynamic of traffic-decrease.

B. Impact on AQM Design

Fluid models for describing the dynamics of long-lived TCP flows have proven useful in studying AQM performance [3], [4], [5], [6], differentiated services [7] and adaptive AQM [8]. In this paper we use them to study the interaction between unresponsive flows and an AQM’s ability to control long-lived TCP flows. To accomplish this we modify this fluid model to account for unresponsive flows, and utilize models for various types of unresponsive traffic. We develop a new model for short-lived TCP flows and also include existing ON-OFF UDP and long-range dependent traffic models in our family of unresponsive flows. These models will help describe both the mean and correlation behavior of unresponsive flows, from which we can analyze effects on AQM performance such as responsiveness and (stability) robustness.² Responsiveness

This work is supported in part by DARPA under Contract DOD F30602-00-0554 and by NSF under grants EIA-0080119, ITR-0085848, and ANI-9980552. Any opinions, findings, and conclusions of the authors do not necessarily reflect the views of the National Science Foundation.

C.V. Hollot is with the ECE Department, University of Massachusetts, Amherst, MA 01003; hollot@ecs.umass.edu

Yong Liu and Don Towsley are with the Computer Science Department, University of Massachusetts, Amherst, MA 01003; {yongliu,towsley}@cs.umass.edu

Vishal Misra is with the Computer Science Department, Columbia University, New York, NY 10027; misra@cs.columbia.edu

¹Throughout this paper, short-lived TCP flows are defined as never leaving the slow-start phase of TCP.

²Hereafter, we simply use *robustness*.

refers to the speed to which a system converges to an equilibrium, while robustness refers to a system remaining stable under network parameter variations. For example, in Section IV we will show that the mean sending rate of unresponsive flows reduces the bandwidth available to long-lived TCP traffic, which, in turn makes the AQM more robust, but less responsive. Another performance issue deals with an AQM's response to variations in unresponsive traffic. We will analyze the resulting variations in queue length, loss probability and long-lived TCP sending rate. This will open a door to study a long-standing question of ours; namely,

What role does queue-averaging play in AQM?

It has been our contention, see [4], that introduction of queue-averaging within an AQM feedback loop has questionable benefit, and, that a well-tuned AQM would opt for an averaging time constant much less than round-trip times. The opposing viewpoint has been that queue averaging plays a role in the presence of bursty and short-lived traffic and traditional greedy flow analysis does not capture this reality. By virtue of the unresponsive flow models we develop in this paper, we are in a position to conduct a more realistic analysis and provide insights into the problem. In Section IV, we show that queue averaging results in a trade-off between AQM responsiveness, robustness and response to the uncontrolled flows. Our conclusion is threefold:

- (i) For robustness, the queue averaging time constant should be chosen outside the range $(R, \frac{R^2 C}{2N})$ where R stands for round-trip time, C the link capacity and N the number of long-lived TCP flows; see Figure 8.
- (ii) AQM responsiveness is inversely related to the queue averaging time constant; see Figure 8.
- (iii) It is impossible, via selection of the averaging time constant, to sufficiently and simultaneously smooth the variations in queue length and loss probability (or conversely, to both the queue length and long-lived TCP sending rate) due to variations in unresponsive flows; see Figure 9.

The last point is particularly important as it implies that while averaging results in a smooth or stable congestion feedback, it also introduces (considerable) jitter in the queueing delay. That trade-off should be noted in AQM design.

The remainder of the paper is organized as follows. In the next section we incorporate unresponsive traffic into the existing fluid model of the TCP/AQM dynamic. In Section III we develop models for different types of unresponsive flows, and use them in Section IV to analyze the trade-offs involved when one introduces queue-averaging in an AQM loop. We support this analysis in Section V with both fluid and ns simulations.

II. INCORPORATING UNRESPONSIVE FLOWS INTO THE TCP/AQM DYNAMIC

In this section, we modify the fluid model for the TCP/AQM dynamic introduced in [3] to include unresponsive flows. We consider a single bottle-neck queue fed by long-lived,

homogeneous TCP connections ℓ and some unresponsive flows u as shown in Figure 1. This diagram highlights the fact that long-lived flows ℓ are under AQM control while unresponsive flows are not.

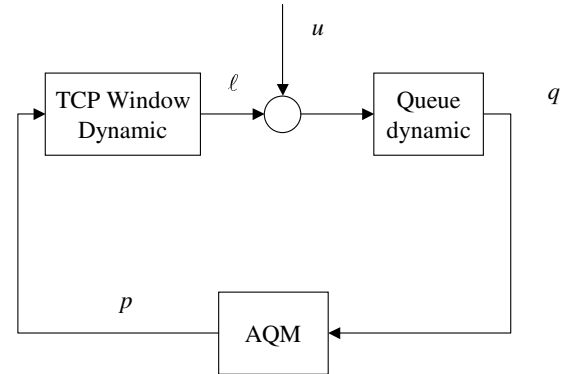


Fig. 1. Block diagram of an AQM system including unresponsive flows u as well as long-lived TCP flows ℓ .

A stochastic differential equation describing the sample path of each long-lived TCP connection is, from [3]:

$$dW_\ell(t) = \frac{dt}{R(q(t))} - \frac{W_\ell(t)}{2} d\eta(t)$$

where W_ℓ is the congestion window size in packets, $R = \frac{q}{C} + T_p$ the round-trip time and η the packet loss modeled by a Poisson process. Here, q , C and T_p denote the queue length (packets), link capacity (packets/sec) and propagation delay (in seconds) respectively. Suppose the number of long-lived TCP connections N is large. The aggregate rate of all long-lived TCP connections is $N\bar{W}_\ell(t)/R(q(t))$, where $\bar{W}_\ell(t)$ is TCP connections' average window size. As in [3], $\bar{W}_\ell(t)$ is described by

$$\frac{d\bar{W}_\ell}{dt} = \frac{1}{R(q(t))} - \frac{\bar{W}_\ell \bar{W}_\ell(t - R_0)}{2R(q(t - R_0))} p(q(t - R_0)) \quad (1)$$

The sample path of the bottle-neck queue length is:

$$\frac{dq(t)}{dt} = -1_{q(t)} C + \frac{N}{R(q(t))} \bar{W}_\ell(t) + u(t),$$

which, in assuming $q(t) > 0$ a.s., becomes

$$\frac{dq(t)}{dt} = -C + \frac{N}{R(q(t))} \bar{W}_\ell(t) + u(t). \quad (2)$$

In the above equation, the dynamics of $\bar{W}_\ell(t)$ is on the time scale of long-lived TCP connections' round trip time. It will be shown in the following section that depending on the characteristics of unresponsive traffic, the variations of $u(t)$ could be on a range of time scales. If we are only interested in the average behavior of the queue on a time scale which is coarser than time scales of $u(t)$'s variations, we can replace $u(t)$ in equation (2) by its mean u_0 . Then,

$$\frac{d\bar{q}(t)}{dt} = -(C - u_0) + \frac{N}{R(\bar{q}(t))} \bar{W}_\ell \quad (3)$$

where it is clear that the link bandwidth is diminished from C to $C - u_0 \triangleq C_{eff}$. This has interesting consequences which we will discuss later, but, we are also interested in how time-variations in the unresponsive traffic affects the management of long-lived flows. To study this we will carry out a linearization on (1) and (2) about the equilibrium $(W_{\ell 0}, p_0, q_0, u_0)$ defined by

$$W_{\ell 0}^2 p_0 = 2; \quad \frac{NW_{\ell 0}}{R_0} = C_{eff}; \quad R_0 = \frac{q_0}{C} + T_p.$$

Linearizing about this equilibrium gives³:

$$\begin{aligned} W_{\ell}(s) &= -P_{win}(s)e^{-sR_0}p(s); \\ \ell(s) &= \frac{N}{R_0}W_{\ell}(s); \\ q(s) &= P_{que}(s)(\ell(s) + u(s)); \end{aligned}$$

where $W_{\ell}(s)$, $p(s)$, $q(s)$, $\ell(s)$ and $u(s)$ refer to the Laplace transforms of the associated time-domain variables⁴, and where the window and queue transfer functions are:

$$P_{win}(s) = \frac{R_0 C_{eff}^2}{s + \frac{2N}{R_0^2 C_{eff}}}; \quad P_{que}(s) = \frac{1}{s + \frac{C_{eff}}{C} \frac{1}{R_0}}. \quad (4)$$

Figure 2 is a block diagram of the linearized AQM feedback system where $C_{aqm}(s)$ denotes the transfer function of an AQM controller. In the next section we turn to the modelling

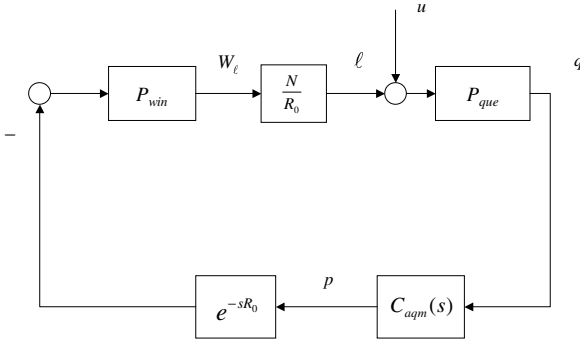


Fig. 2. Linearization of AQM with long-lived TCP traffic ℓ and unresponsive traffic u .

of the unresponsive traffic u .

III. MODELS OF UNRESPONSIVE FLOWS

Now that we have accounted for unresponsive traffic in the TCP/AQM dynamic, see Figure 2, our present objective is to develop appropriate traffic models that will allow us to evaluate their impact on queue length q , loss probability p and long-lived arrival rate ℓ . We will consider three types of unresponsive flows: short-lived TCP flows, Markov ON-OFF

³As shown in [4], a good approximation to (1) is $\frac{d\bar{W}}{dt} = \frac{1}{R(q)} - \frac{\bar{W}\bar{W}(t)}{2R(t)}p(q(t-R_0))$. We also adopt this approximation when linearizing.

⁴To be precise, these time-domain variables represent perturbations from their equilibrium values. To avoid an expanding notation, we retain the original symbol for the perturbed variable.

UDP sources, and, traffic with long-range dependence. For each class we are interested in computing their mean sending rate u_0 and correlation function, or, its Fourier transform, the power spectral density. We have already seen that u_0 reduces the link capacity seen by long-lived flows, and in Section IV we will study how the TCP/AQM dynamic responds to the variations in unresponsive traffic, as captured by their correlation functions.

A. Short-lived TCP Flows

Assuming short-lived TCP traffic never escape the slow-start phase of TCP, their dynamics can be described by their window sizes increasing exponentially, rather than linearly as experienced in TCP's congestion-control phase. We can model a single short-lived TCP flow with sending rate x by the differential equation:

$$\begin{aligned} \dot{W}_x(t) &= \frac{W_x(t)}{R(q(t))} - \frac{W_x(t)W_x(t-R_0)}{2R(q(t-R_0))}p(q(t-R_0)); \\ x(t) &= \frac{W_x(t)}{R(q(t))}; \quad W_x(0) = 1 \end{aligned}$$

where the first term in the \dot{W}_x equation models an exponential increase with time constant R_0 . The second term models multiplicative decrease, and, since we assume short-lived flows terminate before receiving a congestion signal, we drop this loss-dependent term. The growth of this flow is illustrated in Figure 3 where the window size W_x grows exponentially and terminates after a few round-trip times. An approximation




Fig. 3. Exponential growth of a short-lived flow.

of this behavior is can be approximated by an exponential pulse of the form $x(t) = \frac{1}{R_0}e^{(t-t_i)/R_0}\rho_{\tau}(t, t_i)$ where $\rho_{\tau}(t, t_i)$ denotes the (unit) pulse function with arrival time t_i and duration of τ secs. The duration is typically bounded by several round-trip times and for simplicity we take it to be constant and let M be the flow rate at cessation. The aggregate short-lived traffic $u(t)$ is then the sum of such exponential pulses, each characterized by its arrival time t_i ; see Figure 4. Suppose the inter-arrival times of these short-lived flows is exponentially distributed with mean $\frac{1}{\lambda}$. One way to describe the statistics of this process is to consider its time reversal. This reversal preserves the statistics of u and can be modelled as a *shot noise* process. Shot noise (see [9]) is the response




Fig. 4. Aggregate short-lived TCP flows u as a sum of exponential pulses.

of a linear, time-invariant (LTI) system excited by Poisson impulses $z(t)$ with arrival intensity λ ; $z(t) = \sum_i \delta(t - t_i)$. If the impulse response of the LTI system is $h(t)$, the shot-noise process is simply $u(t) = \sum_i h(t - t_i)$. It can be shown that this process has mean and power spectral density

$$u_0 = \lambda H(0); \quad S_{uu}(\omega) = 2\pi\lambda^2 H^2(0)\delta(\omega) + \lambda |H(j\omega)|^2$$

where $H(j\omega)$ is the Fourier transform of $h(t)$. For our model of short-lived flows, $h(t)$ is the exponentially decaying pulse $h(t) = M e^{-t/R_0} \rho_\tau(t, 0)$ having Fourier transform

$$H(j\omega) = \frac{M}{j\omega + \frac{1}{R_0}} \left(1 - \frac{e^{-j\omega\tau}}{MR_0} \right)$$

Consequently, $u_0 = \lambda(MR_0 - 1)$ and

$$S_{\hat{u}\hat{u}}(\omega) = \frac{\lambda M^2}{\omega^2 + \frac{1}{R_0^2}} \left| 1 - \frac{e^{-j\omega\tau}}{MR_0} \right|^2 \quad (5)$$

where $\hat{u}(t) = u(t) - u_0$. We will use the mean and power spectral density computations in (5) later in Section IV to analyze the effect of unresponsive flows on the TCP/AQM dynamic.

B. Markov ON-OFF Model of UDP Traffic

We now consider unresponsive UDP traffic modelled by Markov ON-OFF processes. For a Markov ON-OFF process $x(t)$ (with both ON and OFF periods exponentially distributed), a Poisson-driven stochastic differential equation can be used to describe a sample path [10]:

$$dx(t) = (1 - x(t))dN_1(t) - x(t)dN_2(t) \quad x(0) \in \{0, 1\}$$

where N_1 and N_2 are Poisson counters with rates λ and μ respectively. Taking the expectation on both sides gives

$$\frac{d}{dt} E[x(t)] = (1 - E[x(t)])\lambda - E[x(t)]\mu.$$

In steady state, $x_0 = E[x] = \lambda/(\lambda + \mu)$. To compute the correlation, consider

$$dx(t)x(0) = (1 - x(t))x(0)dN_1(t) - x(t)x(0)dN_2(t).$$

Taking the expectation on both sides gives the differential equation

$$\frac{d}{dt} E[x(0)x(t)] = -(\lambda + \mu)E[x(0)x(t)] + \lambda E[x(0)]$$

with initial condition $E[x(0)x(0)] = E[x(0)] = \lambda/(\lambda + \mu)$. Its solution is

$$E[x(0)x(\tau)] = \frac{\lambda}{(\lambda + \mu)^2} (\mu e^{-(\lambda + \mu)\tau} + \lambda), \quad \tau > 0.$$

The autocorrelation function of its rate variation $\hat{x}(t) = x(t) - x_0$ is:

$$R_{\hat{x}\hat{x}}(\tau) = E[\hat{x}(0)\hat{x}(|\tau|)] = \frac{\lambda\mu}{(\lambda + \mu)^2} e^{-(\lambda + \mu)|\tau|}$$

If we have n independent homogeneous Markov ON-OFF flows with peak rate r , the mean of their aggregate rate u is $u_0 = nr\lambda/(\lambda + \mu)$. The autocorrelation function for the rate variation $\hat{u}(t) = u(t) - u_0$ is then

$$R_{\hat{u}\hat{u}}(\tau) = E[\hat{u}(0)\hat{u}(|\tau|)] = \frac{nr^2\lambda\mu}{(\lambda + \mu)^2} e^{-(\lambda + \mu)|\tau|}$$

with power spectral density

$$S_{\hat{u}\hat{u}}(\omega) = \frac{2nr^2\lambda\mu}{(\lambda + \mu)(\omega^2 + (\lambda + \mu)^2)}. \quad (6)$$

C. $M/G/\infty$ Model of UDP Traffic

The autocorrelation of a Markov ON-OFF processes can only represent short-range dependent traffic. However, network traffic has been shown to exhibit long-range dependence [11], [12], and an $M/G/\infty$ model possesses general autocorrelation structure that captures both short and long-range dependencies; see [13]. An $M/G/\infty$ input process $b(k)$ is a busy server process consisting of a discrete-time, infinite server system fed by a Poisson process with rate λ and generic service time σ distributed according to F . For different distributions F , both short-range and long-range dependent traffic can be generated. The auto-covariance of $b(k)$ has been established to be:

$$\text{cov}[b(k+h), b(k)] = \lambda \sum_{i=h+1}^{\infty} P[\sigma \geq i] \quad (7)$$

To model UDP traffic u using this $M/G/\infty$ process, we first assume that UDP sessions arrive according to a Poisson process with rate λ . Let s denote the lifetime of a UDP session. To generate strict-sense long-range dependent traffic, s must have a Pareto distribution. However, as shown in [14], there is little evidence that session lifetime-distributions exhibit a power law in their extreme tails. Therefore, we will consider a truncated Pareto distribution with exponential tail for s as suggested in [14]. The resulting complementary cumulative distribution function (CCDF) for s is then

$$G(t) \triangleq P(s > t) = \begin{cases} 1 & t < t_0 \\ (\frac{t_0}{t})^\alpha & t_0 \leq t < t_1 \\ ce^{-\beta t} & t \geq t_1 \end{cases} \quad (8)$$

where $1 < \alpha < 2$ and $c = t_0^\alpha e^{\beta t_1} / t_1^\alpha$. Let r denote the sending rate of each UDP session and note that the aggregate sending rate of UDP traffic u is proportional to the number of active UDP sessions $N(t)$. If we discretize the system with time step δ , then, the number of active sessions at each time step forms a $M/G/\infty$ process $b(k)$ with arrival rate $\lambda\delta$ and service time σ with CCDF

$$P[\sigma \geq i] = P[s > i\delta] = G(i\delta).$$

Together with (7), we then have:

$$\text{cov}[u(t + \tau), u(t)] \approx r^2 \lambda \int_{\tau}^{\infty} G(t) dt$$

for sufficiently small δ . For UDP flows with service time distribution specified by (8), the auto-covariance of the aggregate rate is then:

$$\text{cov}[u(t + \tau), u(t)] = \begin{cases} k_1 - r^2 \lambda \tau & 0 < \tau < t_0 \\ \frac{r^2 \lambda t_0^\alpha}{(\alpha - 1) \tau^{\alpha - 1}} + k_2 & t_0 \leq \tau < t_1 \\ \frac{r^2 \lambda c}{\beta} e^{-\beta \tau} & \tau \geq t_1 \end{cases}$$

where

$$k_1 = r^2 \lambda t_0 + \frac{r^2 \lambda t_0^\alpha}{(\alpha - 1)} \left(\frac{1}{t_0^{\alpha - 1}} - \frac{1}{t_1^{\alpha - 1}} \right) + \frac{r^2 \lambda c}{\beta} e^{-\beta t_1};$$

$$k_2 = \frac{r^2 \lambda c}{\beta} e^{-\beta t_1} - \frac{r^2 \lambda t_0^\alpha}{(\alpha - 1) t_1^{\alpha - 1}}.$$

Then the autocorrelation of the rate variation $\hat{u}(t)$ is $R_{\hat{u}\hat{u}}(\tau) = \text{cov}[u(t + |\tau|), u(t)]$ and its power spectral density is

$$S_{\hat{u}\hat{u}}(\omega) = \mathcal{F}\{R_{\hat{u}\hat{u}}(\tau)\}. \quad (9)$$

Because of the power decay of its autocorrelation, $M/G/\infty$ traffic has so called self-similarity. Its impact on queue dynamics is on a wide range of time scales. Therefore $u(t)$ should not be replaced by u_0 in (3). Instead, we should have

$$\frac{d\bar{q}(t)}{dt} = -C + u'_0(t) + \frac{N}{R(\bar{q}(t))} \bar{W}_\ell,$$

where $u'_0(t)$ has all of the fine time scale behavior of $u(t)$ smoothed out. In Section V, we will see from simulations that $M/G/\infty$ process has bigger impact on AQM performance than short-range dependent traffic like Markov ON-OFF flows.

D. Power Spectral Density Comparison

We have introduced three different models for unresponsive traffic, characterized by their power spectral densities (5), (6) and (9). All the power spectral densities are low-pass in nature as illustrated in Figure 5, for the parameter settings:

$$\begin{aligned} \text{Short-lived TCP:} & \quad R_0 = 0.1; MR_0 = 8; \lambda = 70 \\ \text{Markov ON-OFF:} & \quad n = 25; r = 40; \lambda = \mu = 5 \\ \text{M/G/\infty:} & \quad t_0 = 0.1; t_1 = 100; \alpha = 1.5; \\ & \quad \beta = 1; \lambda = 62.5; r = 40. \end{aligned}$$

The parameters for short-lived TCP traffic and the Markov ON-OFF flows are intentionally chosen to show near-identical low-pass structure. The corner frequencies are $1/R_0$ and $\lambda + \mu$ respectively which means that short-lived TCP flows

are correlated on a time scale greater than the round-trip time, while the ON-OFF model has correlation on the order of $1/(\lambda + \mu)$. As seen from (5) and (6), these models have similar power spectral densities when the window size of short-lived TCP flows MR_0 is much larger than 1; in this example, $MR_0 = 8$. Now, compare these power spectral densities to the




Fig. 5. Power spectral density comparison: (a) short-lived TCP traffic, (b) Markov ON-OFF UDP flows, (c) $M/G/\infty$ UDP traffic.

$M/G/\infty$ traffic model in Figure 5. Since, the autocorrelation structure of the $M/G/\infty$ traffic model is composed of a power body, then, its power spectral density has an overall lower decay rate. This is bore out in Section V.C, where simulations show that the $M/G/\infty$ traffic model produces larger variations in queue length and dropping probability.

IV. LINEAR ANALYSIS

In the previous section we characterized several models of unresponsive flows by their mean sending rates u_0 and power spectral density functions $S_{\hat{u}\hat{u}}$. Accordingly, we now analyze their impact on the TCP/AQM dynamic. We begin by studying the effect of mean.

A. Effect of mean sending rate

The mean sending rate u_0 of unresponsive traffic has the effect of reducing available bandwidth so that long-lived TCP flows experience a diminished (effective) link capacity of $C_{eff} = C - u_0$. As discussed in [4], the AQM's gain $\frac{C_{eff}^2}{2N}$ would then decrease, making the AQM more robust but less responsive. To study this impact quantitatively, consider the feedback system in Figure 2 with transfer functions P_{win} and P_{que} given in (4). The resulting closed-loop transfer function

$$S \triangleq \frac{1}{1 + \frac{N}{R_0} P_{win} C_{aqm} P_{que} e^{-sR_0}} \quad (10)$$

is commonly referred to as the closed-loop sensitivity function, and, stability and responsiveness of the TCP/AQM dynamic

can be deduced from its frequency response. Roughly speaking, robustness is inversely proportional to the peak frequency response $\|S\|_\infty \triangleq \sup_\omega |S(j\omega)|$, while responsiveness is proportional to S 's unity-gain crossover frequency ω_g defined by $|S(j\omega_g)| = 1$. To show how responsiveness and robustness are affected by the mean sending rate, consider network parameters $C = 3750$ packets/sec, $R_0 = 250$ ms and $N = 60$ TCP sessions, and, a proportional AQM controller $C_{aqm} = 5.8624 \times 10^{-5}$. This controller roughly corresponds to RED with the time constant of queue averaging taken much less than the round-trip time. For a target queue length of 50 packets, we plot in Figure 6 the responsiveness and robustness as a function of the mean sending rate $u_0 \in (0, C)$. As




Fig. 6. This figure shows how AQM robustness increases (top figure) while responsiveness decreases (bottom figure) when the mean send rate u_0 increases.

we can see, AQM robustness increases (top figure) while responsiveness decreases (bottom figure) when the mean send rate u_0 increases (C_{eff} decreases). Therefore, an effect of the unresponsive traffic's mean sending rate is to make AQM less responsive, but more robust.

B. Effect of variation in the unresponsive flows

Now we consider how variation in unresponsive flows affects AQM performance. We will use the autocorrelation functions in (5), (6) and (9) to evaluate performance. We take the AQM dynamic $C_{aqm}(s)$ to be the RED controller as modeled by the low-pass filter:

$$\begin{aligned} \dot{q}_{avg} &= -K_{red}q_{avg} + K_{red}q; \\ p &= L_{red}q_{avg} \end{aligned}$$

with associated transfer function $C_{aqm}(s) = \frac{L_{red}K_{red}}{s+K_{red}}$. A discrete-time implementation of RED is described by the exponential moving-average filter:

$$\begin{aligned} q_{avg}(k+1) &= (1-w_q)q_{avg}(k) + w_q q(k); \\ p(k) &= \frac{max_p}{max_{th} - min_{th}} q_{avg}(k) \end{aligned}$$

where $1-w_q = e^{-K_{red}/C}$; $L_{red} = \frac{max_p}{max_{th} - min_{th}}$. Intuitively, this queue-averaging (or low-pass filtering) should impact the degree to which variations in the unresponsive flows affect longer-lived TCP flows. Indeed, if variations in unresponsive flows exist at time scales much faster than that which AQM loop operates, then, queue-averaging could filter-out those higher-frequency queue-components, which otherwise could communicate, via the AQM, misleading congestion signals. We now analyze this claim.

1) *Closed-loop responses*: To evaluate the variations in queue length q , loss probability p and the arrival rate ℓ , due to variations in the unresponsive flows \hat{u} , we use the defining relationship:

$$S_{yy}(\omega) = |T_{yx}(j\omega)|^2 S_{xx}(\omega) \quad (11)$$

where (x, y) is the input-output pair of a linear system with transfer function $T_{yx}(s)$, and (S_{xx}, S_{yy}) their power spectral densities. Referring to the TCP/AQM block-diagram in Figure 2, we can form the following closed-loop transfer functions:

$$\begin{aligned} T_{q\hat{u}} &= P_{que}S; \\ T_{p\hat{u}} &= C_{aqm}P_{que}S; \\ T_{\ell\hat{u}} &= \frac{N}{R_0}P_{win}C_{aqm}P_{que}S e^{-sR_0} \end{aligned} \quad (12)$$

where the sensitivity S is given in (10). As an example, the power spectral density for the queue length is computed from: $S_{qq}(\omega) = |T_{q\hat{u}}(j\omega)|^2 S_{\hat{u}\hat{u}}(\omega)$ where $S_{\hat{u}\hat{u}}(\omega)$ is taken from (5), (6) or (9). Before we do these computations we consider a trade-off between robustness and responsiveness induced by queue-averaging.

2) *Trade-off between robustness and responsiveness*: Queue-averaging in an AQM creates a robustness/responsiveness trade-off. For example, let's consider TCP/AQM dynamics for two values of the averaging parameter: $K_{red} = 0.09, 100$ (or $w_g = 2.5 \times 10^{-5}, 0.0267$) as shown in Figure 7. The network parameters are $C = 3750$ packets/sec, $R_0 = 250$ ms, $N = 60$ TCP sessions and $L_{red} = 5.8624 \times 10^{-5}$. The dashed curve (corresponding to more queue-averaging $K_{red} = 0.09$) exhibits less responsiveness and more robustness.⁵ In Figure 8, we plot robustness and responsiveness against queue-averaging K_{red} . The bottom plot confirms intuition that AQM responsiveness is inversely related to averaging. On the other hand, the top plot indicates that robustness is not monotonic. This can be explained: For small amounts of queue-averaging (large values of K_{red}), the low-pass filter contributes little additional phase delay to the feedback loop - hence, the loop is relatively stable. However, when K_{red} is comparable to the poles of P_{win} and P_{que} in (4), or, equivalently, when the time constant of averaging ($\frac{1}{K_{red}}$) is near the interval $(\frac{R_0 C}{C_{eff}}, \frac{R_0 C_{eff}^2}{2N}) = (0.35, 2.38)$, stability margins decrease;

⁵Recall that robustness is inversely proportional to the peak in $|S(j\omega)|$, while responsiveness is proportional to the frequency ω_g for which $|S(j\omega_g)| = 1$




Fig. 7. Comparison of sensitivity functions S for two values of queue-averaging parameter, illustrating the trade-off between responsiveness and robustness. The dashed curve is for $K_{red} = 0.09$ ($w_g = 2.5 \times 10^{-5}$), while the solid one is for $K_{red} = 100$ ($w_g = 0.0267$)

this interval is marked in Figure 8. Further queue-averaging makes the TCP/AQM dynamic more stable, but sluggish. With these trade-offs in hand, we now return to analyzing the effects of the variations in unresponsive flows.




Fig. 8. Trade-off between robustness (top figure) and AQM responsiveness (bottom figure) as a function of queue averaging.

3) *Variance in unresponsive flows:* We now analyze how variations in unresponsive flows distributes across queue length q , loss probability p and TCP arrival rate ℓ . Specifically, we use (11) and (12) to compute the variance of these responses. For example, the variance of the queue response

is

$$var_q \triangleq \mathcal{R}_{qq}(0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} |T_{q\hat{u}}(j\omega)|^2 S_{\hat{u}\hat{u}}(\omega) d\omega$$

where $\mathcal{R}_{qq}(\tau)$ is q 's autocorrelation function. Using the above we can compute variances var_q , var_p and var_ℓ in response to 10 Markov ON-OFF (UDP) sources having peak rate $r = 450$ kbps, as a function of the queue-averaging parameter K_{red} ; see Figure 9. We use the same parameters $C = 3750$ packets/sec, $R_0 = 250$ ms, $N = 60$ TCP sessions, $L_{red} = 5.8624 \times 10^{-5}$ and $\lambda + \mu = 10$; see Figure 9. In the case of the




Fig. 9. Variance in queue length q (top), loss probability p (middle) and long-lived TCP arrival rate ℓ (bottom) as a function of queue-averaging

queue and arrival rate, the variance is not a monotonic function of the queue-averaging parameter. The variance peaks for values of time constant $\frac{1}{K_{red}}$ near the range $(\frac{R_0 C}{C_{eff}}, \frac{R_0 C_{eff}^2}{2N}) = (0.35, 2.38)$, where robustness dips in Figure (8). We notice that the variance in the queue *does not go to zero* as $K_{red} \rightarrow 0$ ($w_q \rightarrow 0$). This occurs since $\lim_{K_{red} \rightarrow 0} |S(j\omega)| = 1$ so that

$$var_q \rightarrow \frac{1}{2\pi} \int_{-\infty}^{\infty} |P_{que}(j\omega)|^2 S_{ss}(\omega) d\omega \neq 0$$

as $K_{red} \rightarrow 0$. On the other hand, var_p and var_ℓ tend to zero as $K_{red} \rightarrow 0$ since T_{ps} and T_{ls} depend explicitly on C_{aqm} and

$$\int_{-\infty}^{\infty} |C_{aqm}(j\omega)|^2 d\omega \rightarrow 0.$$

This illustrates the property that it is impossible, via selection of the averaging time constant K_{red} (or w_q), to sufficiently and simultaneously smooth the variations in queue length and loss probability (or conversely, to both the queue length and long-lived TCP sending rate) due to variations in unresponsive flows;

V. SIMULATIONS

To verify results established in previous sections, we conducted both fluid and packet level simulations with different unresponsive traffic models: short-lived TCP flows, Markov ON-OFF flows, $M/G/\infty$ flows and CBR flows. For fluid simulations, we use Matlab's Simulink tool to simulate the nonlinear fluid-flow model of TCP/AQM dynamics in (1) and (2). UDP traffic rate $u(t)$ is imported from trace files generated according to different unresponsive traffic models. Packet level simulations were done in ns to compare with Simulink results. Due to space limit, we only present simulation results for Markov ON-OFF flows and $M/G/\infty$ flows here. Interested readers can contact us for more results.

A. Simulation Configuration

We consider a simple network topology as in Figure 10. There are two groups of flows sharing the bottle-neck link




Fig. 10. Simulation Topology

between B_1 and B_2 . One group is TCP flows sending data from S_1 to R_1 . The other is composed of unresponsive flows originating at S_2 and destined for R_2 . The one-way propagation delay is $123ms$ for both groups and the packet size is set to 500 bytes. The bottle-neck link's capacity is 3750 packets/sec and the router is managed by RED with parameters: $maxthresh_ = 1800$, $thresh_ = 100$, $linterm_ = 10$ and $gentle_ = 1$. We vary the queue averaging parameter (w_q) from 1 to 0.00001 (equivalently, K_{red} from 3750 to 0.0375) to study the performance of the RED.

All simulations begin with 60 long-lived TCP flows. After the network stabilizes, we introduce unresponsive flows at time 50 seconds. Unresponsive flows are shut up at time 150 seconds. In order to test AQM responsiveness, we increased the network load level at time 150 seconds by introducing 15 new long-lived TCP flows.

B. Simulations for Markov ON-OFF Flows

We use 10 Markov ON-OFF flows to generate unresponsive traffic. The average ON and OFF periods for UDP flows is $200ms$. When a UDP flow is ON, it generates traffic at rate of $450kbps$. On average, the UDP traffic takes away 15% of the available bandwidth at the bottle-neck link.

1) *Simulink Fluid Simulation*: Figures 11 and 12 show the UDP traffic's impact on queue length and loss probability under different queue averaging parameters.


Oscillations in both queue length and loss probability are composed of two parts: high frequency part and low frequency part. The high frequency part is contribution from rate variation in Markov ON-OFF flows. The low frequency part is determined by the responsiveness of the control loop. It happens whenever the system is driven away from its operation point, both at the beginning of the simulation, and, when the load level changed at time 50, 100 and 150 seconds. When the queue averaging parameter w_q gets smaller (more averaging), the magnitude of high frequency oscillation in both queue and loss probability gets smaller. This is because system robustness increases. At the same time, we see more and more low frequency oscillations. It takes the system longer to reach its operating point after being perturbed. This is because the AQM dynamic becomes less responsive. Both behaviors are consistent with the analysis in Figure 8.

Finally we measure the variance in queue length q , loss probability p and long-lived TCP arrival rate ℓ when the UDP traffic is ON. Figure 13 shows these variances as a function of K_{red} and the averaging parameter w_q . In comparing with the linear analysis of Figure 9 we see that the trends are similar, but that there are differences. First, for the nonlinear fluid simulation, the effects of averaging occur for larger values of K_{red} (or w_q). It's as if the nonlinear fluid model was operating at smaller round-trip times. Second, the absolute values of the variances differ. Third, the variance in TCP arrival rate, see var_ℓ in Figure 13, does not go to zero as $K_{red} \rightarrow 0$. The linear analysis in Figure 9 predicts $var_\ell \rightarrow 0$. This occurs since the linear analysis does not account for the hard constraint on the loss probability: $p \in [0, 1]$.


2) *ns Simulation*: We repeated previous simulations in ns. We constructed a topology as in Figure 10 and used ns's Exponential ON/OFF traffic generator to generate Markov ON-OFF flows. Figure 14 shows the evolution of the bottle-neck queue. Figure 15 depicts the RED's dropping probability. Both of them agree with results from fluid simulations in Simulink. Again, we can see the system's trade-off between stability and responsiveness.

C. Simulations for $M/G/\infty$ Traffic


For $M/G/\infty$ model, we set the arrival rate of UDP sessions to be $\lambda = 25$ sessions/sec. The duration of each session assumes the distribution in (8), with $t_0 = 73ms$, $\alpha = 1.5$, $\beta = 1$, $t_1 = 100s$. The average length of generated UDP sessions is again $200ms$. On average we have 5 sessions active simultaneously. The sending rate of an active UDP session is $r = 430kbps$. The average rate of $M/G/\infty$ traffic is 15% of the bottle-neck bandwidth. All these first order traffic statistics are set close to previous Markov ON-OFF experiments. Figures 16 and 17 show the results from ns. Although the trends remain the same as in the previous experiments, we can see much bigger oscillations in both queue length and dropping probability. This is because the



(a) $w_q = 1$ ($K_{red} = 3750$)




(b) $w_q = 0.001$ ($K_{red} = 3.75$)




(c) $w_q = 0.00001$ ($K_{red} = 0.0375$)


Fig. 11. Simulink: Queue evolution with 15% UDP Traffic. Queue-averaging increases with smaller w_q or K_{red} .



(a) $w_q = 1$ ($K_{red} = 3750$)



(b) $w_q = 0.001$ ($K_{red} = 3.75$)



(c) $w_q = 0.00001$ ($K_{red} = 0.0375$)

Fig. 12. Simulink: RED loss probability with 15% UDP Traffic. Queue-averaging increases with smaller w_q or K_{red} .

power in $M/G/\infty$ traffic is bigger than Markov ON-OFF traffic as shown in Figure 5. It tells us that we really need to model long-range dependent unresponsive traffic carefully when studying AQM performance.

VI. CONCLUSIONS

This paper has made four contributions. First, it has incorporated unresponsive traffic into fluid models of the TCP/AQM dynamic. Second, the paper has introduced a shot noise model for short-lived TCP flows and power spectral analysis has been carried out for different unresponsive traffic models, namely, short-lived TCP model, Markov ON-OFF model, and $M/G/\infty$ model. Third, through linear analysis, fluid-model simulation and ns simulations, this paper evaluated the impact unresponsive traffic had on AQM performance. Finally, the paper offered insight into the tradeoffs between queue averaging and various AQM performance criteria.

REFERENCES

- [1] <http://www.caida.org>, CAIDA.




Fig. 13. Variance in queue length q (top), loss probability p (middle) and long-lived TCP arrival rate ℓ (bottom) as a function of queue-averaging from Simulink simulations.

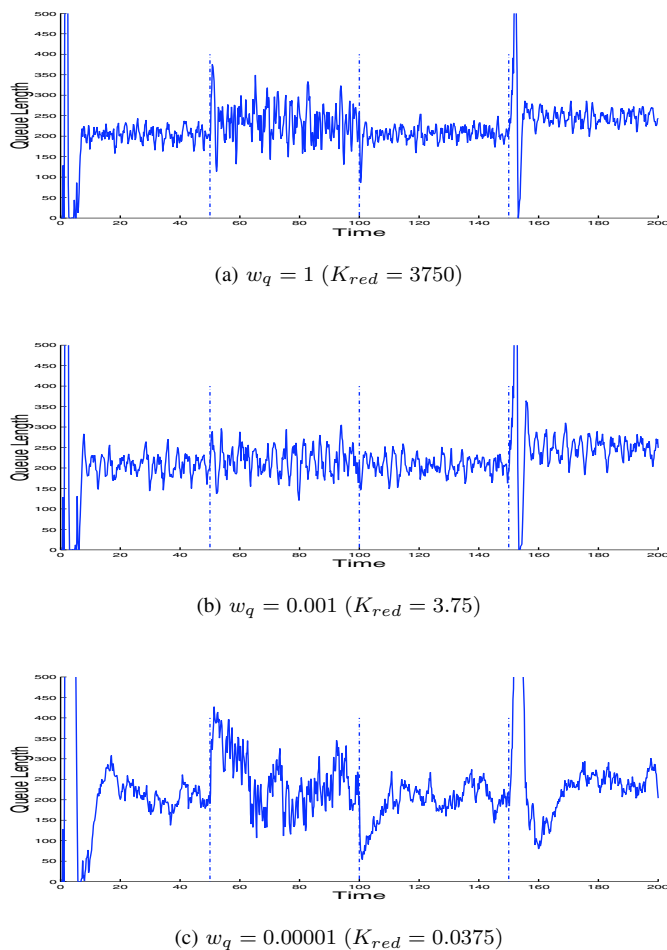


Fig. 14. ns: Queue evolution with 15% UDP traffic

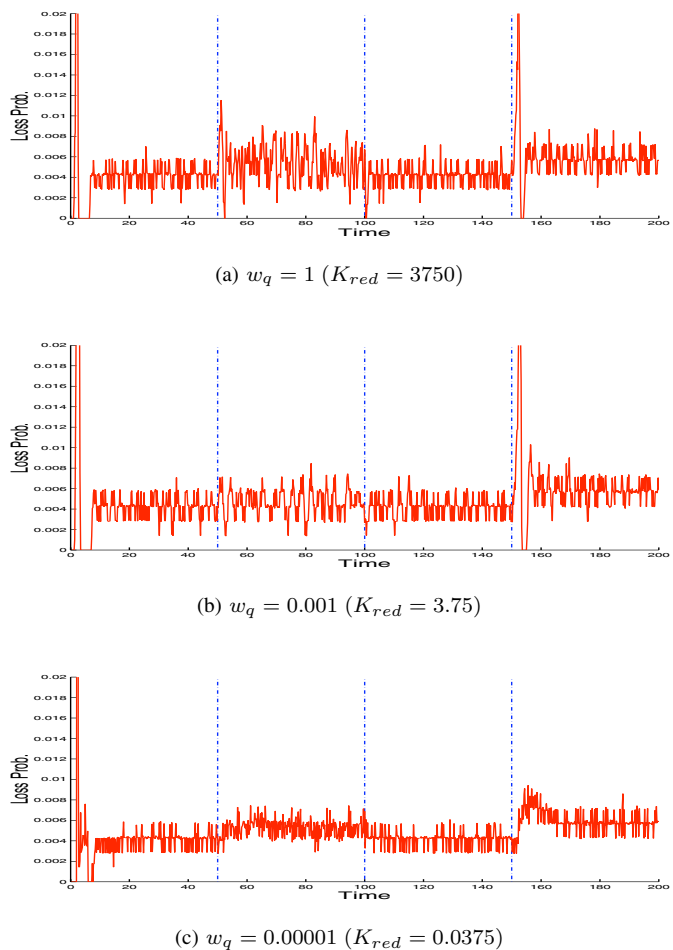




Fig. 15. ns: RED loss probability with 15% UDP traffic


- [2] S. Shakkottai and R. Srikant, "Mean FDE models for Internet congestion control under a many-flows regime," Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Tech. Rep., 2001.
- [3] V. Misra, W. Gong, and D. Towsley, "Fluid based analysis of a network of AQM routers supporting TCP flows with an application to RED," in *Proceedings of ACM/SIGCOMM*, 2000.
- [4] C. Hollot, V. Misra, D. Towsley, and W. Gong, "Analysis and design of controllers for AQM routers supporting TCP flows," *IEEE Transactions on Automatic Control*, vol. 47, no. 6, 2002.
- [5] S. Kunniyur and R. Srikant, "A time-scale decomposition approach to adaptive ecn marking," in *Proceedings of IEEE/INFOCOM*, 2001.
- [6] S. Low, F. Paganini, and J. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, 2002.
- [7] Y. Chait, C. Hollot, V. Misra, D. Towsley, H. Zhang, and J.-C. Lui, "Providing throughput differentiation for tcp flows using adaptive two-color marking and two-level aqm," in *Proceedings of IEEE/INFOCOM*, 2002.
- [8] H. Zhang, C. Hollot, D. Towsley, and V. Misra, "A self-tuning structure for adaptation in tcp/aqm networks," Department of Computer Sciences, UMass, Amherst, Tech. Rep., July 2002.
- [9] A. Papoulis, Ed., *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, 1965.
- [10] R. Brockett, "Lecture notes: Stochastic control," Harvard University, Tech. Rep., 1983.
- [11] M. E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web traffic: Evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, December 1997.
- [12] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking* 2, vol. 6, pp. 1–15, 1994.
- [13] M. Parulekar and A. Makowski, "M/G/∞ input processes: A versatile class of models for network traffic," in *Proceedings of INFOCOM*, 1997.
- [14] W. Gong, Y. Liu, V. Misra, and D. Towsley, "On the tails of web file size distributions," in *Proceedings of 39th Allerton Conference on Communication Control and Computing*, Oct 2001.



(a) $w_q = 1$ ($K_{red} = 3750$)




(b) $w_q = 0.001$ ($K_{red} = 3.75$)




(c) $w_q = 0.00001$ ($K_{red} = 0.0375$)


Fig. 16. ns: Queue Evolution with 15% $M/G/\infty$ Traffic.



(a) $w_q = 1$ ($K_{red} = 3750$)



(b) $w_q = 0.001$ ($K_{red} = 3.75$)



(c) $w_q = 0.00001$ ($K_{red} = 0.0375$)

Fig. 17. ns: RED Loss Probability with 15% $M/G/\infty$ Traffic.