

Dual-Resource TCP/AQM for Processing-Constrained Networks

Minsu Shin, *Student Member, IEEE*, Song Chong, *Member, IEEE*, and Injong Rhee, *Senior Member, IEEE*

Abstract—This paper examines congestion control issues for TCP flows that require in-network processing on the fly in network elements such as gateways, proxies, firewalls and even routers. Applications of these flows are increasingly abundant in the future as the Internet evolves. Since these flows require use of CPUs in network elements, both bandwidth and CPU resources can be a bottleneck and thus congestion control must deal with “congestion” on both of these resources. In this paper, we show that conventional TCP/AQM schemes can significantly lose throughput and suffer harmful unfairness in this environment, particularly when CPU cycles become more scarce (which is likely the trend given the recent explosive growth rate of bandwidth). As a solution to this problem, we establish a notion of dual-resource proportional fairness and propose an AQM scheme, called Dual-Resource Queue (DRQ), that can closely approximate proportional fairness for TCP Reno sources with in-network processing requirements. DRQ is scalable because it does not maintain per-flow states while minimizing communication among different resource queues, and is also incrementally deployable because of no required change in TCP stacks. The simulation study shows that DRQ approximates proportional fairness without much implementation cost and even an incremental deployment of DRQ at the edge of the Internet improves the fairness and throughput of these TCP flows. Our work is at its early stage and might lead to an interesting development in congestion control research.

Index Terms—CPU capacity, efficiency, fairness, proportional fairness, TCP-AQM, transmission link capacity.

I. INTRODUCTION

ADVANCES in optical network technology enable fast pace increase in physical bandwidth whose growth rate has far surpassed that of other resources such as CPU and memory bus. This phenomenon causes network bottlenecks to shift from bandwidth to other resources. The rise of new applications that require in-network processing hastens this shift, too. For instance, a voice-over-IP call made from a cell phone to a PSTN phone must go through a media gateway that performs audio transcoding “on the fly” as the two end points often use different audio compression standards. Examples of in-network

processing services are increasingly abundant from security, performance-enhancing proxies (PEP), to media translation [1], [2]. These services add additional loads to processing capacity in the network components. New router technologies such as extensible routers [3] or programmable routers [4] also need to deal with scheduling of CPU usage per packet as well as bandwidth usage per packet. Moreover, the standardization activities to embrace various network applications especially at network edges are found in [5] and [6] as the name of Open Pluggable Edge Services.

In this paper, we examine congestion control issues for an environment where both bandwidth and CPU resources can be a bottleneck. We call this environment *dual-resource* environment. In the dual-resource environment, different flows could have different processing demands per byte.

Traditionally, congestion control research has focused on managing only bandwidth. However, we envision (also it is indeed happening now to some degree) that diverse network services reside somewhere inside the network, most likely at the edge of the Internet, processing, storing or forwarding data packets on the fly. As the in-network processing is likely to be popular in the future, our work that examines whether the current congestion control theory can be applied without modification, or if not, then what scalable solutions can be applied to fix the problem, is highly timely.

In our earlier work [7], we extended proportional fairness to the dual-resource environment and proposed a distributed congestion control protocol for the same environment where end-hosts are cooperative and explicit signaling is available for congestion control. In this paper, we propose a scalable active queue management (AQM) strategy, called *Dual-Resource Queue* (DRQ), that can be used by network routers to approximate proportional fairness, without requiring any change in end-host TCP stacks. Since it does not require any change in TCP stacks, our solution is incrementally deployable in the current Internet. Furthermore, DRQ is highly scalable in the number of flows it can handle because it does not maintain per-flow states or queues. DRQ maintains only one queue per resource and works with classes of application flows whose processing requirements are *a priori* known or measurable.

Resource scheduling and management of one resource type in network environments where different flows could have different demands are a well-studied area of research. Weighted-fair queuing (WFQ) [8] and its variants such as deficit round robin (DRR) [9] are well known techniques to achieve fair and efficient resource allocation. However, the solutions are not scalable and implementing them in a high-speed router with many flows is difficult since they need to maintain per-flow queues and states. Another extreme is to have routers maintain simpler

Manuscript received January 15, 2006; revised August 21, 2006, and January 12, 2007; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor F. Paganini. This work was supported by the center for Broadband OFDM Mobile Access (BrOMA) at POSTECH through the ITRC program of the Korean MIC, supervised by IITA (IITA-2006-C1090-0603-0037). An early version of this paper was presented at the IEEE INFOCOM 2006, Barcelona, Spain, 2006.

M. Shin and S. Chong are with the School of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 305-701, Korea (e-mail: msshin@netsys.kaist.ac.kr; song@ee.kaist.ac.kr).

I. Rhee is with the Department of Computer Science, North Carolina State University, Raleigh, NC 27695 USA (e-mail: rhee@csc.ncsu.edu).

Digital Object Identifier 10.1109/TNET.2007.900415

queue management schemes such as RED [10], REM [11] or PI [12]. Our study finds that these solutions may yield extremely unfair allocation of CPU and bandwidth and sometimes lead to very inefficient resource usages.

Some fair queueing algorithms such as Core-Stateless Fair Queueing (CSFQ) [13] and Rainbow Fair Queueing (RFQ) [14] have been proposed to eliminate the problem of maintaining per-flow queues and states in routers. However, those schemes are concerned about bandwidth sharing only and do not consider joint allocation of bandwidth and CPU cycles. Estimation-based Fair Queueing (EFQ) [15] and Prediction Based Fair Queueing (PBFQ) [16] have been also proposed for fair CPU sharing but they require per-flow queues and do not consider joint allocation of bandwidth and CPU cycles either.

Our study, to the best of our knowledge, is the first in examining the issues of TCP and AQM under the dual-resource environment and we show that by simulation DRQ achieves fair and efficient resource allocation without imposing much implementation cost. The remainder of this paper is organized as follows. In Section II, we define the problem and fairness in the dual-resource environment, in Sections III and IV, we present DRQ and its simulation study, and in Section V, we conclude our paper.

II. PRELIMINARIES: NETWORK MODEL AND DUAL-RESOURCE PROPORTIONAL FAIRNESS

A. Network Model

We consider a network that consists of a set of unidirectional links, $L = \{1, \dots, L\}$, and a set of CPUs, $K = \{1, \dots, K\}$. The transmission capacity (or bandwidth) of link l is B_l (bits/sec) and the processing capacity of CPU k is C_k (cycles/sec). These network resources are shared by a set of flows (or data sources), $S = \{1, \dots, S\}$. Each flow s is associated with its data rate r_s (bits/sec) and its end-to-end route (or path) which is defined by a set of links, $L(s) \subset L$, and a set of CPUs, $K(s) \subset K$, that flow s travels through. Let $S(l) = \{s \in S | l \in L(s)\}$ be the set of flows that travel through link l and let $S(k) = \{s \in S | k \in K(s)\}$ be the set of flows that travel through CPU k . Note that this model is general enough to include various types of router architecture and network element with multiple CPUs and transmission links.

Flows can have different CPU demands. We represent this notion by *processing density* w_s^k , $k \in K$, of each flow s , which is defined to be the average number of CPU cycles required per bit when flow s is processed by CPU k . w_s^k depends on k since different processing platforms (CPU, OS, and software) would require a different number of CPU cycles to process the same flow s . The processing demand of flow s at CPU k is then $w_s^k r_s$ (cycles/sec).

Since there are limits on CPU and bandwidth capacities, the amount of processing and bandwidth usage by all flows sharing these resources must be less than or equal to the capacities at anytime. We represent this notion by the following two constraints: for each CPU $k \in K$, $\sum_{s \in S(k)} w_s^k r_s \leq C_k$ (*processing constraint*) and for each link $l \in L$, $\sum_{s \in S(l)} r_s \leq B_l$ (*bandwidth constraint*). These constraints are called *dual-resource*

constraints and a nonnegative rate vector $r = [r_1, \dots, r_S]^T$ satisfying these dual constraints for all CPUs $k \in K$ and all links $l \in L$ is said to be *feasible*.

A1: We assume that each CPU $k \in K$ knows the processing densities w_s^k 's for all the flows $s \in S(k)$.

This assumption is reasonable because a majority of Internet applications are known and their processing requirements can be measured either off-line or on-line as discussed below. In practice, network flows could be readily classified into a small number of application types [15], [17]–[19]. That is, there is a finite set of application types, a flow is an instance of an application type, and flows will have different processing densities only if they belong to different application types. In [17], applications have been divided into two categories: header-processing applications and payload-processing applications, and each category has been further divided into a set of benchmark applications. In particular, authors in [15] experimentally measure the per-packet processing times for several benchmark applications such as encryption, compression, and forward error correction. The measurement results find the network processing workloads to be highly regular and predictable. Based on the results, they propose an empirical model for the per-packet processing time of these applications for a given processing platform. Interestingly, it is a simple affine function of packet size M , i.e., $\mu_a^k + \nu_a^k M$ where μ_a^k and ν_a^k are the parameters specific to each benchmark application a for a given processing platform k . Thus, the processing density (in cycles/bit) of a packet of size M from application a at platform k can be modelled as $(\mu_a^k/M) + \nu_a^k$. Therefore, the average processing density w_a^k of application a at platform k can be computed upon arrival of a packet using an exponentially weighted moving average (EWMA) filter:

$$w_a^k \leftarrow (1 - \lambda)w_a^k + \lambda \left(\frac{\mu_a^k}{M} + \nu_a^k \right), \quad 0 < \lambda < 1. \quad (1)$$

One could also directly measure the quantity $(\mu_a^k/M) + \nu_a^k$ in (1) as a whole instead of relying on the empirical model by counting the number of CPU cycles actually consumed by a packet while the packet is being processed. Lastly, determining the application type an arriving packet belongs to is an easy task in many commercial routers today since L3/L4 packet classification is a default functionality.

B. Proportional Fairness in the Dual-Resource Environment

Fairness and efficiency are two main objectives in resource allocation. The notion of fairness and efficiency has been extensively studied and well understood with respect to bandwidth sharing. In particular, proportionally fair (PF) rate allocation has been considered as the bandwidth sharing strategy that can provide a good balance between fairness and efficiency [20], [21].

In our recent work [7], we extended the notion of proportional fairness to the dual-resource environment where processing and bandwidth resources are jointly constrained. In the following, we present this notion and its potential advantages for the dual-resource environment to define our goal for our main study of this paper on TCP/AQM.

Consider an aggregate log utility maximization problem (P) with dual constraints

$$\mathbf{P} : \max_r \sum_{s \in S} \alpha_s \log r_s \quad (2)$$

$$\text{subject to } \sum_{s \in S(k)} w_s^k r_s \leq C_k, \quad \forall k \in K \quad (3)$$

$$\sum_{s \in S(l)} r_s \leq B_l, \quad \forall l \in L \quad (4)$$

$$r_s \geq 0, \quad \forall s \in S \quad (5)$$

where α_s is the weight (or willingness to pay) of flow s . The solution r^* of this problem is unique since it is a strictly concave maximization problem over a convex set [22]. Furthermore, r^* is weighted proportionally fair since $\sum_{s \in S} \alpha_s ((r_s - r_s^*)/r_s^*) \leq 0$ holds for all feasible rate vectors r by the optimality condition of the problem. We define this allocation to be (*dual-resource*) *PF rate allocation*. Note that this allocation can be different from Kelly's PF allocation [20] since the set of feasible rate vectors can be different from that of Kelly's formulation due to the extra processing constraint (3).

From the duality theory [22], r^* satisfies that

$$r_s^* = \frac{\alpha_s}{\sum_{k \in K(s)} w_s^k \theta_k^* + \sum_{l \in L(s)} \pi_l^*}, \quad \forall s \in S \quad (6)$$

where $\theta^* = [\theta_1^*, \dots, \theta_K^*]^T$ and $\pi^* = [\pi_1^*, \dots, \pi_L^*]^T$ are Lagrange multiplier vectors for (3) and (4), respectively, and θ_k^* and π_l^* can be interpreted as congestion prices of CPU k and link l , respectively. Eq. (6) reveals an interesting property that the PF rate of each flow is inversely proportional to the aggregate congestion price of its route with the contribution of each θ_k^* being weighted by w_s^k . The congestion price θ_k^* or π_l^* is positive only when the corresponding resource becomes a bottleneck, and is zero, otherwise.

To illustrate the characteristics of PF rate allocation in the dual-resource environment, let us consider a limited case where there are only one CPU and one link in the network, as shown in Fig. 1. For now, we drop k and l in the notation for simplicity. Let \bar{w}_a and \bar{w}_h be the weighted arithmetic and harmonic means of the processing densities of flows sharing the CPU and link, respectively. So, $\bar{w}_a = \sum_{s \in S} (w_s \alpha_s / \sum_{s \in S} \alpha_s)$ and $\bar{w}_h = (\sum_{s \in S} (\alpha_s / w_s \sum_{s \in S} \alpha_s))^{-1}$. There exist three cases as below.

- *CPU-limited case* ($\theta^* > 0$ and $\pi^* = 0$): $r_s^* = \alpha_s / w_s \theta^*$, $\forall s \in S$, $\sum_{s \in S} w_s r_s^* = C$ and $\sum_{s \in S} r_s^* \leq B$. From these, we know that this case occurs when $C/B \leq \bar{w}_h$ and PF rate allocation becomes $r_s^* = \alpha_s C / w_s \sum_{s \in S} \alpha_s$, $\forall s \in S$.
- *Bandwidth(BW)-limited case* ($\theta^* = 0$ and $\pi^* > 0$): $r_s^* = \alpha_s / \pi^*$, $\forall s \in S$, $\sum_{s \in S} w_s r_s^* \leq C$ and $\sum_{s \in S} r_s^* = B$. From these, we know that this case occurs when $C/B \geq \bar{w}_a$ and PF rate allocation becomes $r_s^* = \alpha_s B / \sum_{s \in S} \alpha_s$, $\forall s \in S$.
- *Jointly-limited case* ($\theta^* > 0$ and $\pi^* > 0$): This case occurs when $\bar{w}_h < (C/B) < \bar{w}_a$. By plugging $r_s^* = \alpha_s / (w_s \theta^* +$




Fig. 1. Single-CPU and single-link network.

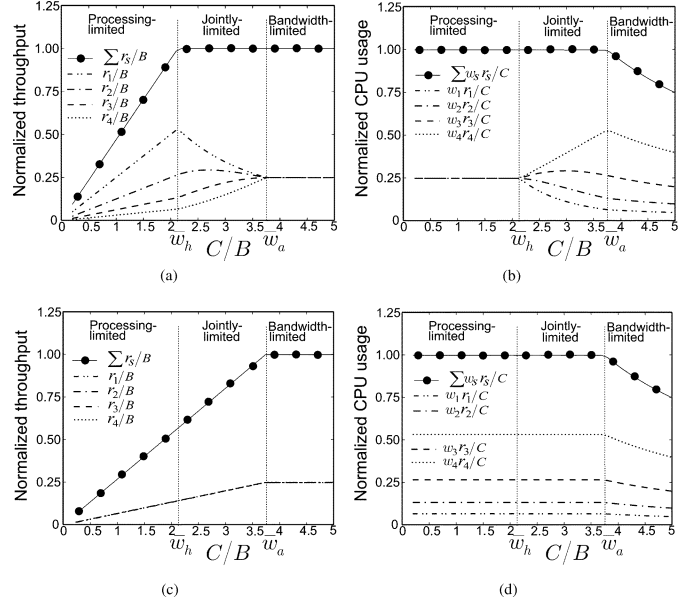


Fig. 2. Fairness and efficiency in the dual-resource environment (single-CPU and single-link network): (a) and (b) respectively show the normalized bandwidth and CPU allocations enforced by PF rate allocation, and (c) and (d) respectively show the normalized bandwidth and CPU allocations enforced by TCP-like rate allocation. When $C/B < \bar{w}_a$, TCP-like rate allocation gives lower bandwidth utilization than PF rate allocation (shown in (a) and (c)) and has an unfair allocation of CPU cycles (shown in (d)).

$\pi^*)$, $\forall s \in S$, into $\sum_{s \in S} w_s r_s^* = C$ and $\sum_{s \in S} r_s^* = B$, we can obtain θ^* , π^* and consequently r_s^* , $\forall s \in S$.

We can apply other increasing and concave utility functions (including the one from TCP itself [23]) in the dual-resource problem in (2)–(5). The reason why we give a special attention to proportional fairness by choosing log utility function is that it automatically yields weighted fair CPU sharing ($w_s r_s^* = \alpha_s C / \sum_{s \in S} \alpha_s$, $\forall s \in S$) if CPU is limited, and weighted fair bandwidth sharing ($r_s^* = \alpha_s B / \sum_{s \in S} \alpha_s$, $\forall s \in S$) if bandwidth is limited, as illustrated in the example of Fig. 1. This property is obviously what is desirable and a direct consequence of the particular form of rate-price relationship given in (6). Thus, this property is not achievable when other utility functions are used.

Fig. 2(a) and (b) illustrates the bandwidth and CPU allocations enforced by PF rate allocation in the single-CPU and single-link case using an example of four flows with identical weights ($\alpha_s = 1$, $\forall s$) and different processing densities (w_1, w_2, w_3, w_4) = (1, 2, 4, 8) where $\bar{w}_h = 2.13$ and $\bar{w}_a = 3.75$. For comparison, we also consider a rate allocation in which flows with an identical end-to-end path get an equal share of the maximally achievable throughput of the path and call it *TCP-like rate allocation*. That is, if TCP flows run on the example network in Fig. 1 with ordinary AQM schemes such as RED on both CPU and link queues, they would have the same long-term throughput. Thus, in our example, TCP-like

rate allocation is defined to be the maximum equal rate vector satisfying the dual constraints, which is $r_s = B/S$, $\forall s$, if $C/B \geq \bar{w}_a$, and $r_s = C/\bar{w}_a S$, $\forall s$, otherwise. The bandwidth and CPU allocations enforced by TCP-like rate allocation are shown in Figs. 2(c) and (d).

From Fig. 2, we observe that TCP-like rate allocation yields far less aggregate throughput than PF rate allocation when $C/B < \bar{w}_a$, i.e., in both CPU-limited and jointly-limited cases. Intuitively, this is because TCP-like allocation which finds an equal rate allocation yields unfair sharing of CPU cycles as CPU becomes a bottleneck (see Fig. 2(d)), which causes the severe aggregate throughput drop. In contrast, PF allocation yields equal sharing of CPU cycles, i.e., $w_s r_s$ become equal for all $s \in S$, as CPU becomes a bottleneck (see Fig. 2(b)), which mitigates the aggregate throughput drop. This problem in TCP-like allocation would get more severe when the processing densities of flows have a more skewed distribution.

In summary, in a single-CPU and single-link network, PF rate allocation achieves equal bandwidth sharing when bandwidth is a bottleneck, equal CPU sharing when CPU is a bottleneck, and a good balance between equal bandwidth sharing and equal CPU sharing when bandwidth and CPU form a joint bottleneck. Moreover, in comparison to TCP-like rate allocation, such consideration of CPU fairness in PF rate allocation can increase aggregate throughput significantly when CPU forms a bottleneck either alone or jointly with bandwidth.

III. MAIN RESULT: SCALABLE TCP/AQM ALGORITHM

In this section, we present a scalable AQM scheme, called Dual-Resource Queue (DRQ), that can approximately implement dual-resource PF rate allocation described in Section II for TCP-Reno flows. DRQ modifies RED [10] to achieve PF allocation without incurring per-flow operations (queueing or state management). DRQ does not require any change in TCP stacks.

A. DRQ Objective and Optimality

We describe a TCP/AQM network using the fluid model as in the literature [23]–[28]. In the fluid model, the dynamics whose timescale is shorter than several tens (or hundreds) of round-trip times (RTTs) are neglected. Instead, it is convenient to study the longer timescale dynamics and so adequate to model the macroscopic dynamics of long-lived TCP flows that we are concerning.

Let $x_s(t)$ (bits/sec) be the average data rate of TCP source s at time t where the average is taken over the time interval Δ (seconds) and Δ is assumed to be on the order of tens (or hundreds) of RTTs, i.e., large enough to average out the additive-increase and multiplicative decrease (AIMD) oscillation of TCP. Define the RTT τ_s of source s by $\tau_s = \tau_{si} + \tau_{is}$ where τ_{si} denotes forward-path delay from source s to resource i and τ_{is} denotes backward-path delay from resource i to source s .

A2: We assume that each TCP flow s has a constant RTT τ_s , as customary in the fluid modeling of TCP dynamics [23]–[28].

Let $y_l(t)$ be the average queue length at link l at time t , measured in bits. Then

$$\dot{y}_l(t) = \begin{cases} \sum_{s \in S(l)} x_s(t - \tau_{sl}) - B_l & y_l(t) > 0 \\ \left[\sum_{s \in S(l)} x_s(t - \tau_{sl}) - B_l \right]^+ & y_l(t) = 0 \end{cases} \quad (7)$$

Similarly, let $z_k(t)$ be the average queue length at CPU k at time t , measured in CPU cycles. Then

$$\dot{z}_k(t) = \begin{cases} \sum_{s \in S(k)} w_s^k x_s(t - \tau_{sk}) - C_k & z_k(t) > 0 \\ \left[\sum_{s \in S(k)} w_s^k x_s(t - \tau_{sk}) - C_k \right]^+ & z_k(t) = 0 \end{cases} \quad (8)$$

Let $p_s(t)$ be the end-to-end marking (or loss) probability at time t to which TCP source s reacts. Then, the rate-adaptation dynamics of TCP Reno or its variants, particularly in the timescale of tens (or hundreds) of RTTs, can be readily described by [23]

$$\dot{x}_s(t) = \begin{cases} \frac{M_s(1-p_s(t))}{N_s \tau_s^2} - \frac{2}{3} \frac{x_s^2(t)p_s(t)}{N_s M_s} & x_s(t) > 0 \\ \left[\frac{M_s(1-p_s(t))}{N_s \tau_s^2} - \frac{2}{3} \frac{x_s^2(t)p_s(t)}{N_s M_s} \right]^+ & x_s(t) = 0 \end{cases} \quad (9)$$

where M_s is the average packet size in bits of TCP flow s and N_s is the number of consecutive data packets that are acknowledged by an ACK packet in TCP flow s (N_s is typically 2).

In DRQ, we employ one RED queue per one resource. Each RED queue computes a probability (we refer to it as *pre-marking probability*) in the same way as an ordinary RED queue computes its marking probability.

That is, the RED queue at link l computes a pre-marking probability $\rho_l(t)$ at time t by

$$\rho_l(t) = \begin{cases} 0 & \hat{y}_l(t) \leq \underline{b}_l \\ \frac{m_l}{\bar{b}_l - \underline{b}_l} (\hat{y}_l(t) - \underline{b}_l) & \underline{b}_l \leq \hat{y}_l(t) \leq \bar{b}_l \\ \frac{1-m_l}{\bar{b}_l} (\hat{y}_l(t) - \bar{b}_l) + m_l & \bar{b}_l \leq \hat{y}_l(t) \leq 2\bar{b}_l \\ 1 & \hat{y}_l(t) \geq 2\bar{b}_l \end{cases} \quad (10)$$

$$\dot{\hat{y}}_l(t) = \frac{\log_e(1 - \lambda_l)}{\eta_l} \hat{y}_l(t) - \frac{\log_e(1 - \lambda_l)}{\eta_l} y_l(t) \quad (11)$$

where $m_l \in (0, 1]$, $0 \leq \underline{b}_l < \bar{b}_l$ and (11) is the continuous-time representation of the EWMA filter [25] used by the RED, i.e.,

$$\hat{y}_l((k+1)\eta_l) = (1 - \lambda_l)\hat{y}_l(k\eta_l) + \lambda_l y_l(k\eta_l), \quad \lambda_l \in (0, 1). \quad (12)$$

Eq. (11) does not model the case where the averaging timescale of the EWMA filter is smaller than the averaging timescale Δ on which $y_l(t)$ is defined. In this case, (11) must be replaced by $\dot{\hat{y}}_l(t) = y_l(t)$.

Similarly, the RED queue at CPU k computes a pre-marking probability $\sigma_k(t)$ at time t by

$$\sigma_k(t) = \begin{cases} 0 & \hat{v}_k(t) \leq \underline{b}_k \\ \frac{m_k}{\bar{b}_k - \underline{b}_k} (\hat{v}_k(t) - \underline{b}_k) & \underline{b}_k \leq \hat{v}_k(t) \leq \bar{b}_k \\ \frac{1-m_k}{\bar{b}_k} (\hat{v}_k(t) - \bar{b}_k) + m_k & \bar{b}_k \leq \hat{v}_k(t) \leq 2\bar{b}_k \\ 1 & \hat{v}_k(t) \geq 2\bar{b}_k \end{cases} \quad (13)$$

$$\dot{\hat{v}}_k(t) = \frac{\log_e(1-\lambda_k)}{\eta_k} \hat{v}_k(t) - \frac{\log_e(1-\lambda_k)}{\eta_k} v_k(t) \quad (14)$$

where $v_k(t)$ is the translation of $z_k(t)$ in bits.

Given these pre-marking probabilities, the objective of DRQ is to mark (or discard) packets in such a way that the end-to-end marking (or loss) probability $p_s(t)$ seen by each TCP flow s at time t becomes

$$p_s(t) = \frac{\left(\sum_{k \in K(s)} w_s^k \sigma_k(t - \tau_{ks}) + \sum_{l \in L(s)} \rho_l(t - \tau_{ls}) \right)^2}{1 + \left(\sum_{k \in K(s)} w_s^k \sigma_k(t - \tau_{ks}) + \sum_{l \in L(s)} \rho_l(t - \tau_{ls}) \right)^2}. \quad (15)$$

The actual marking scheme that can closely approximate this objective function will be given in Section III-B.

The Reno/DRQ network model given by (7)–(15) is called *average Reno/DRQ network* as the model describes the interaction between DRQ and Reno dynamics in long-term average rates rather than explicitly capturing instantaneous TCP rates in the AIMD form. This average network model enables us to study fixed-valued equilibrium and consequently establish *in an average sense the equilibrium equivalence* of a Reno/DRQ network and a network with the same configuration but under dual-resource PF congestion control.

Let $x = [x_1, \dots, x_S]^T$, $\sigma = [\sigma_1, \dots, \sigma_K]^T$, $\rho = [\rho_1, \dots, \rho_L]^T$, $p = [p_1, \dots, p_S]^T$, $y = [y_1, \dots, y_L]^T$, $z = [z_1, \dots, z_K]^T$, $v = [v_1, \dots, v_K]^T$, $\hat{y} = [\hat{y}_1, \dots, \hat{y}_L]^T$ and $\hat{v} = [\hat{v}_1, \dots, \hat{v}_K]^T$.

Proposition 1: Consider an average Reno/DRQ network given by (7)–(15) and formulate the corresponding aggregate log utility maximization problem (Problem **P**) as in (2)–(5) with $\alpha_s = \sqrt{3/2} M_s / \tau_s$. If the Lagrange multiplier vectors, θ^* and π^* , of this corresponding Problem **P** satisfy the following conditions:

$$\mathbf{C1} : \theta_k^* < 1, \forall k \in K(s), \forall s \in S \quad (16)$$

$$\mathbf{C2} : \pi_l^* < 1, \forall l \in L(s), \forall s \in S \quad (17)$$

then, the average Reno/DRQ network has a unique equilibrium point $(x^*, \sigma^*, \rho^*, p^*, y^*, z^*, v^*, \hat{y}^*, \hat{v}^*)$ and (x^*, σ^*, ρ^*) is the primal-dual optimal solution of the corresponding Problem **P**. In addition, $v_k^* > \underline{b}_k$ if $\sigma_k^* > 0$ and $0 \leq v_k^* \leq \underline{b}_k$ otherwise, and $y_l^* > \underline{b}_l$ if $\rho_l^* > 0$ and $0 \leq y_l^* \leq \underline{b}_l$ otherwise, for all $k \in K$ and $l \in L$.

Proof: The proof is given in the Appendix. ■

Proposition 1 implies that once the Reno/DRQ network reaches its steady state (i.e., equilibrium), the average data rates of Reno sources satisfy weighted proportional fairness with weights $\alpha_s = \sqrt{3/2} M_s / \tau_s$. In addition, if a CPU k is a

bottleneck (i.e., $\sigma_k^* > 0$), its average equilibrium queue length v_k^* stays at a constant value greater than \underline{b}_k , and if not, it stays at a constant value between 0 and \underline{b}_k . The same is true for link congestion.

The existence and uniqueness of such an equilibrium point in the Reno/DRQ network is guaranteed if conditions **C1** and **C2** hold in the corresponding Problem **P**. Otherwise, the Reno/DRQ networks do not have an equilibrium point.

In the current Internet environment, however, these conditions will hardly be violated particularly as the bandwidth-delay products of flows increase. By applying **C1** and **C2** to the Lagrangian optimality condition of Problem **P** in (6) with $\alpha_s = \sqrt{3/2} M_s / \tau_s$, we have

$$\frac{r_s^* \tau_s}{M_s} = \frac{\sqrt{3/2}}{\sum_{k \in K(s)} w_s^k \theta_k^* + \sum_{l \in L(s)} \pi_l^*} \quad (18)$$

$$> \frac{\sqrt{3/2}}{\sum_{k \in K(s)} w_s^k + |L(s)|} \quad (19)$$

where $r_s^* \tau_s / M_s$ is the bandwidth-delay product (or window size) of flow s , measured in packets. The maximum packet size in the Internet is $M_s = 1,536$ bytes (i.e., maximum Ethernet packet size). Flows that have the minimum processing density are IP forwarding applications with maximum packet size [17]. For instance, a measurement study in [15] showed that per-packet processing time required for NetBSD radix-tree routing table lookup on a Pentium 167 MHz processor is $51 \mu\text{s}$ (for a faster CPU, the processing time reduces; so as what matters is the number of cycles per bit, this estimate applies to the other CPUs). Thus, the processing density for this application flow is about $w_s^k = 51 (\mu\text{sec}) \times 167 (\text{MHz}) / 1,536 (\text{bytes}) = 0.69$ (cycles/bit). Therefore, from (19), the worst-case lower bound on the window size becomes $(r_s^* \tau_s / M_s) > 1.77$ (packets), which occurs when the flow traverses a CPU only in the path (i.e., $|K(s)| = 1$ and $|L(s)| = 0$). This concludes that the conditions **C1** and **C2** will never be violated as long as the steady-state average TCP window size is sustainable at a value greater than or equal to 2 packets, even in the worst case.

B. DRQ Implementation

In this section, we present a simple scalable packet marking (or discarding) scheme that closely approximates the DRQ objective function we laid out in (15).

A3: We assume that for all times

$$\left(\sum_{k \in K(s)} w_s^k \sigma_k(t - \tau_{ks}) + \sum_{l \in L(s)} \rho_l(t - \tau_{ls}) \right)^2 \ll 1, \forall s \in S. \quad (20)$$

This assumption implies that $(w_s^k \sigma_k(t))^2 \ll 1, \forall k \in K$, $\rho_l(t)^2 \ll 1, \forall l \in L$, and any product of $w_s^k \sigma_k(t)$ and $\rho_l(t)$ is also much smaller than 1. Note that our analysis is based on long-term average values of $\sigma_k(t)$ and $\rho_l(t)$. The typical operating points of TCP in the Internet during steady state where TCP shows a reasonable performance are under low end-to-end loss probabilities (less than 1%) [29]. Since the end-to-end

When a packet arrives at resource i at time t :

- if (ECN \neq 11)
 - set ECN to 11 with probability $\delta_i(t)$;
- if (ECN == 00)
 - set ECN to 10 with probability $\epsilon_i(t)$;
- else if (ECN == 10)
 - set ECN to 11 with probability $\epsilon_i(t)$;

Fig. 3. DRQ's ECN marking algorithm.

average probabilities are low, the marking probabilities at individual links and CPUs can be much lower.

Let R be the set of all the resources (including CPUs and links) in the network. Also, for each flow s , let $R(s) = \{1, \dots, |R(s)|\} \subset R$ be the set of all the resources that it traverses along its path and let $i \in R(s)$ denote the i -th resource along its path and indicate whether it is a CPU or a link. Then, some manipulation after applying Assumption **A3** to (15) gives

$$p_s(t) \approx \left(\sum_{k \in K(s)} w_s^k \sigma_k(t - \tau_{ks}) + \sum_{l \in L(s)} \rho_l(t - \tau_{ls}) \right)^2$$

$$= \sum_{i=1}^{|R(s)|} \delta_i(t - \tau_{is}) + \sum_{i=2}^{|R(s)|} \sum_{i'=1}^{i-1} \epsilon_i(t - \tau_{is}) \epsilon_{i'}(t - \tau_{i's}) \quad (21)$$

where

$$\delta_i(t) = \begin{cases} (w_s^i \sigma_i(t))^2, & \text{if } i \text{ indicates CPU} \\ \rho_i(t)^2, & \text{if } i \text{ indicates link} \end{cases} \quad (22)$$

and

$$\epsilon_i(t) = \begin{cases} \sqrt{2} w_s^i \sigma_i(t), & \text{if } i \text{ indicates CPU} \\ \sqrt{2} \rho_i(t), & \text{if } i \text{ indicates link.} \end{cases} \quad (23)$$

Eq. (21) tells that each resource $i \in R(s)$ (except the first resource in $R(s)$, i.e., $i = 1$) contributes to $p_s(t)$ with two quantities, $\delta_i(t - \tau_{is})$ and $\sum_{i'=1}^{i-1} \epsilon_i(t - \tau_{is}) \epsilon_{i'}(t - \tau_{i's})$. Moreover, resource i can compute the former using its own congestion information, i.e., $\sigma_i(t)$ if it is a CPU or $\rho_i(t)$ if it is a link, whereas it cannot compute the latter without knowing the congestion information of its upstream resources on its path ($\forall l' < l$). That is, the latter requires an inter-resource signaling to exchange the congestion information. For this reason, we refer to $\delta_i(t)$ as *intra-resource marking probability* of resource i at time t and $\epsilon_i(t)$ as *inter-resource marking probability* of resource i at time t . We solve this intra- and inter-resource marking problem using two-bit ECN flags without explicit communication between resources.

Consider the two-bit ECN field in the IP header [30]. Among the four possible values of ECN bits, we use three values to indicate three cases: *initial state* (ECN = 00), *signaling-marked* (ECN = 10) and *congestion-marked* (ECN = 11). When a packet is congestion-marked (ECN = 11), the packet is either marked (if TCP supports ECN) or discarded (if not). DRQ sets the ECN bits as shown in Fig. 3.

Below, we verify that the ECN marking scheme in Fig. 3 approximately implements the objective function in (21). Consider

a flow s with path $R(s)$. For now, we drop the time index t to simplify the notation. Let $P_{00}^i, P_{10}^i, P_{11}^i$ respectively denote the probabilities that packets of flow s will have ECN = 00, ECN = 10, ECN = 11, upon departure from resource i in $R(s)$. Then, the proposed ECN marking scheme can be expressed by the following recursion. For $i = 1, 2, \dots, |R(s)|$,

$$P_{11}^i = P_{11}^{i-1} + (1 - P_{11}^{i-1}) \delta_i + P_{10}^{i-1} (1 - \delta_i) \epsilon_i \quad (24)$$

$$= 1 - (1 - \delta_i) (1 - P_{11}^{i-1} - P_{10}^{i-1} \epsilon_i) \quad (25)$$

$$P_{10}^i = P_{10}^{i-1} (1 - \delta_i) (1 - \epsilon_i) + P_{00}^{i-1} (1 - \delta_i) \epsilon_i \quad (26)$$

$$P_{00}^i = P_{00}^{i-1} (1 - \delta_i) (1 - \epsilon_i) \quad (27)$$

with the initial condition that $P_{00}^0 = 1, P_{10}^0 = 0, P_{11}^0 = 0$.

Evolving i from 0 to $|R(s)|$, we obtain

$$P_{11}^{|R(s)|} = 1 - \prod_{i=1}^{|R(s)|} (1 - \delta_i) \left(1 - \sum_{i=2}^{|R(s)|} \sum_{i'=1}^{i-1} \epsilon_i \epsilon_{i'} + \Theta \right) \quad (28)$$

where Θ is the higher-order terms (order ≥ 3) of ϵ_i 's. By Assumption **A3**, we have

$$P_{11}^{|R(s)|} \approx 1 - \prod_{i=1}^{|R(s)|} (1 - \delta_i) \left(1 - \sum_{i=2}^{|R(s)|} \sum_{i'=1}^{i-1} \epsilon_i \epsilon_{i'} \right) \quad (29)$$

$$\approx \sum_{i=1}^{|R(s)|} \delta_i + \sum_{i=2}^{|R(s)|} \sum_{i'=1}^{i-1} \epsilon_i \epsilon_{i'} \quad (30)$$

which concludes that the proposed ECN marking scheme approximately implements the DRQ objective function in (21) since $P_{11}^{|R(s)|} = p_s$.

Disclaimer: DRQ requires alternative semantics for the ECN field in the IP header, which are different from the default semantics defined in RFC 3168 [31]. What we have shown here is that DRQ can be implemented using two-bit signaling such as ECN. The coexistence of the default semantics and the alternative semantics required by DRQ needs further study.

C. DRQ Stability

In this section, we explore the stability of Reno/DRQ networks. Unfortunately, analyzing its global stability is an extremely difficult task since the dynamics involved are nonlinear and retarded. Here, we present a partial result concerning local stability, i.e., stability around the equilibrium point.

Define $|R| \times |S|$ matrix $\Gamma(z)$ whose (i, s) element is given by

$$\Gamma_{is}(z) = \begin{cases} w_s^i e^{-z\tau_{is}}, & \text{if } s \in S(i) \text{ and } i \text{ indicates CPU} \\ e^{-z\tau_{is}}, & \text{if } s \in S(i) \text{ and } i \text{ indicates link} \\ 0, & \text{otherwise} \end{cases} \quad (31)$$

Proposition 2: An average Reno/DRQ network is locally stable if we choose the RED parameters in DRQ such that $\max\{(m_k/(\bar{b}_k - \underline{b}_k)), ((1 - m_k)/\bar{b}_k)\} C_k \in (0, \psi)$, $\forall k \in K$, and $\max\{(m_l/(\bar{b}_l - \underline{b}_l)), ((1 - m_l)/\bar{b}_l)\} B_l \in (0, \psi)$, $\forall l \in L$, and

$$\psi \leq \frac{\sqrt{3/2} \phi_{\min}^2 \Gamma(0)}{|R| \Lambda_{\max}^2 \tau_{\max}^3 w_{\max}} \quad (32)$$

where $\Lambda_{max} = \max\{\max\{C_k\}/\min\{w_s^k\} \min\{M_s\}, \max\{B_l\} \min\{M_s\}\}$, $\tau_{max} = \max\{\tau_s\}$, $w_{max} = \max\{w_s^k, 1\}$ and $\phi_{min}[\Gamma(0)]$ denotes the smallest singular values of the matrix $\Gamma(z)$ evaluated at $z = 0$.

Proof: The proof is given in the Appendix and it is a straightforward application of the TCP/RED stability result in [32]. ■

IV. PERFORMANCE

A. Simulation Setup

In this section, we use simulation to verify the performance of DRQ in the dual-resource environment with TCP Reno sources. We compare the performance of DRQ with that of the two other AQM schemes that we discussed in the introduction. One scheme is to use the simplest approach where both CPU and link queues use RED and the other is to use DRR (a variant of WFQ) to schedule CPU usage among competing flows according to the processing density of each flow. DRR maintains per flow queues, and equalizes the CPU usage in a round robin fashion when the processing demand is higher than the CPU capacity (i.e., CPU-limited). In some sense, these choices of AQM are two extreme; one is simple, but less fair in use of CPU as RED is oblivious to differing CPU demands of flows and the other is complex, but fair in use of CPU as DRR installs equal shares of CPU among these flows. Our goal is to demonstrate through simulation that DRQ using two FIFO queues always offers provable fairness and efficiency, which is defined as the dual-resource PF allocation. Note that all three schemes use RED for link queues, but DRQ uses its own marking algorithm for link queues as shown in Fig. 3 which uses the marking probability obtained from the underlying RED queue for link queues. We call the scheme with DRR for CPU queues and RED for link queues, *DRR-RED*, the scheme with RED for CPU queues and RED for link queues, *RED-RED*.

The simulation is performed in the NS-2 [33] environment. We modified NS-2 to emulate the CPU capacity by simply holding a packet for its processing time duration. In the simulation, TCP-NewReno sources are used at end hosts and RED queues are implemented using its default setting for the “gentle” RED mode [34] ($m_i = 0.1$, $\bar{b}_i = 50$ pkts, $\bar{b}_i = 550$ pkts and $\lambda_i = 10^{-4}$). The packet size is fixed at 500 Bytes). The same RED setting is used for the link queues of DRR-RED and RED-RED, and also for both CPU and link queues of DRQ (DRQ uses a function of the marking probabilities to mark or drop packets for both queues). In our analytical model, we separate CPU and link. To simplify the simulation setup and its description, when we refer to a “link” for the simulation setup, we assume that each link l consists of one CPU and one Tx link (i.e., bandwidth).

By adjusting CPU capacity C_l , link bandwidth B_l , and the amount of background traffic, we can control the bottleneck conditions. Our simulation topologies are chosen from a various set of Internet topologies from simple dumbbell topologies to more complex WAN topologies. Below we discuss these setups and simulation scenarios in detail and their corresponding results for the three schemes we discussed above.




Fig. 4. Single link scenario in dumbbell topology.

B. Dumbell With Long-Lived TCP Flows

To confirm our analysis in Section II-B, we run a single link bottleneck case. Fig. 4 shows an instance of the dumbbell topology commonly used in congestion control research. We fix the bandwidth of the bottleneck link to 40 Mbps and vary its CPU capacity from 5 Mcycles/s to 55 Mcycles/s. This variation allows the bottleneck to move from the CPU-limited region to the BW-limited region. Four classes of long-lived TCP flows are added for simulation whose processing densities are 0.25, 0.5, 1.0 and 2.0 respectively. We simulate ten TCP Reno flows for each class. All the flows have the same RTT of 40 ms.

In presenting our results, we take the average throughput of TCP flows that belong to the same class. Fig. 5 plots the average throughput of each class. To see whether DRQ achieves PF allocation, we also plot the ideal proportional fair rate for each class (which is shown in a dotted line). As shown in Fig. 5(a), where we use typical RED schemes at both queues, all TCP flows achieve the same throughput regardless of the CPU capacity of the link and their processing densities. Figs. 5(b) and (c) show that the average throughput curves of DRR-RED and DRQ follow the ideal PF rates reasonably well. When CPU is only a bottleneck resource, the PF rate of each flow must be inversely proportional to its processing density w_s , in order to share CPU equally. Under the BW-limited region, the proportionally-fair rate of each flow is identical to the equal share of the bandwidth. Under the jointly-limited region, flows maintain the PF rates while fully utilizing both resources. Although DRQ does not employ the per-flow queue structure as DRR, its performance is comparable to that of DRR-RED.

Fig. 6 shows that the aggregate throughput achieved by each scheme. It shows that RED-RED has much lower bandwidth utilization than the two other schemes. This is because, as discussed in Section II-B, when CPU is a bottleneck resource, the equilibrium operating points of TCP flows over the RED CPU queue that achieve the equal bandwidth usage while keeping the total CPU usage below the CPU capacity are much lower than those of the other schemes that need to ensure the equal sharing of CPU (not the bandwidth) under the CPU-limited region.

C. Impact of Flows With High Processing Demands

In the introduction, we indicated that RED-RED can cause extreme unfairness in use of resources. To show this by experiment, we construct a simulation run where we fix the CPU capacity to 40 Mcycles/s and add an increasing number of flows with a high CPU demand ($w_s = 10$) in the same setup as the dumbbell sink bottleneck environment in Section IV-B. We call these flows *high processing flows*. From Fig. 5, at 40 Mcycles/s, when no high processing flows are added, CPU is not a bottleneck. But as the number of high processing flows increases, the network moves into the CPU-limited region. Fig. 7 shows the results of this simulation run. In Fig. 7(a), as we increase the number of high processing flows, the aggregate CPU share

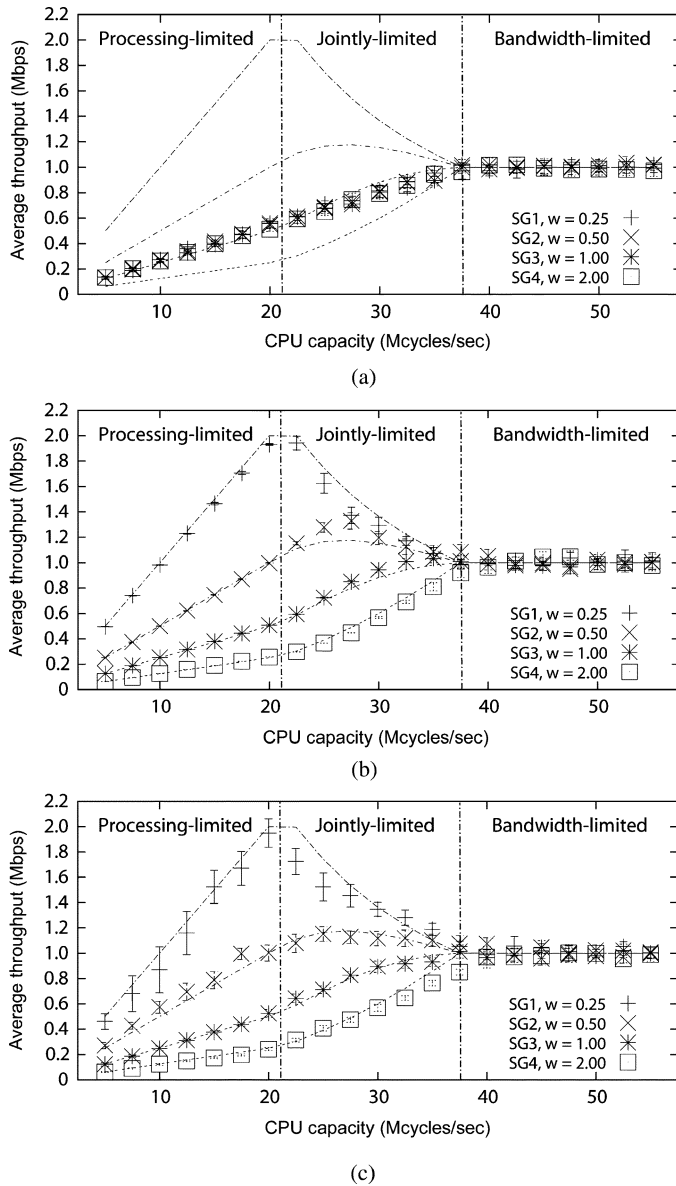


Fig. 5. Average throughput of four different classes of long-lived TCP flows in the dumbbell topology. Each class has a different CPU demand per bit (w). No other background traffic is added. The dotted lines indicate the ideal PF rate allocation for each class. In the figure, we find that DRQ and DRR-RED show good fairness under the CPU-limited region while RED-RED does not. Vertical bars indicate 95% confidence intervals. (a) RED-RED; (b) DRR-RED; (c) DRQ.

of high processing flows deviates significantly from the equal CPU share; under a larger number of high processing flows (e.g., 10 flows), these flows dominate the CPU usage over the other lower processing density flows, driving them to starvation. In contrast, DRQ and DRR approximately implement the equal CPU sharing policy. Even though the number of high processing flows increases, the bandwidth remains a bottleneck resource as before, so the link is in the jointly-limited region which is the reason why the CPU share of high-processing flows go beyond 20%.

D. Dumbell With Background Internet Traffic

No Internet links are without cross traffic. In order to emulate more realistic Internet environments, we add cross traffic modelled from various observations on RTT distribution [35], flow

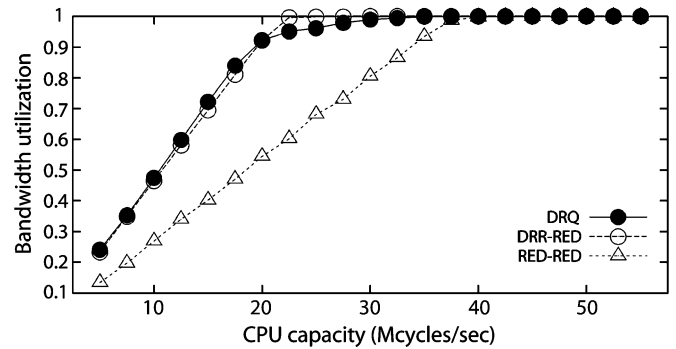


Fig. 6. Comparison of bandwidth utilization in the Dumbbell single bottleneck topology. RED-RED achieves far less bandwidth utilization than DRR-RED and DRQ when CPU becomes a bottleneck.

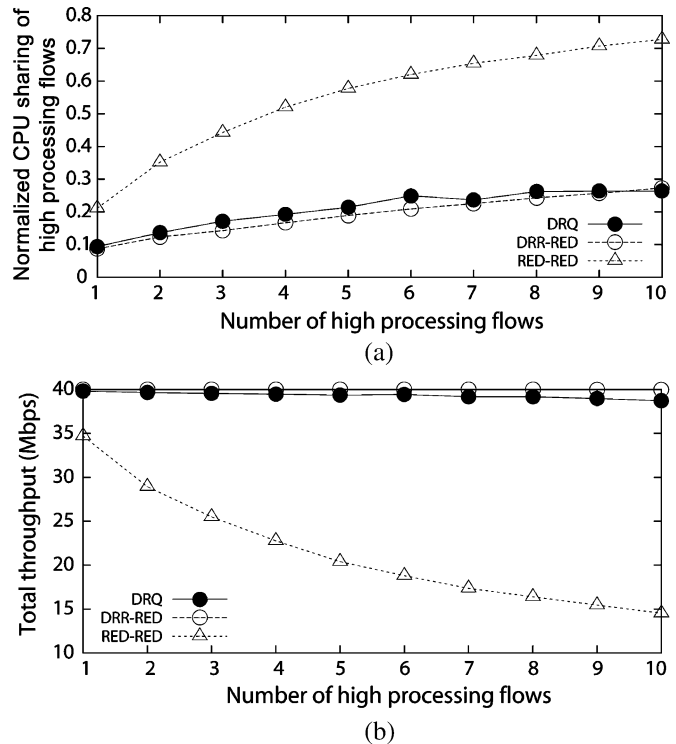


Fig. 7. Impact of high processing flows. As the number of high processing flows increase, the network becomes more CPU-bound. Under RED-RED, these flows can dominate the use of CPU, reaching about 80% CPU usage with only 10 flows, starving 40 competing, but low processing flows. (a) CPU sharing; (b) total throughput.

sizes [36] and flow arrival [37]. As modelling the Internet traffic in itself is a topic of research, we do not dwell on which model is more realistic. In this paper, we present one model that contains the statistical characteristics that are commonly assumed or confirmed by researchers. These characteristics include that the distribution of flow sizes has a long-range dependency [36], [38], the RTTs of flows is rather exponentially distributed [39] and the arrivals of flows are exponentially distributed [37]. Following these characteristics, our cross traffic consists of a number of short and medium-lived TCP flows that follow a Poisson arrival process and send a random number of packets derived from a hybrid distribution of Lognormal (body) and Pareto (tail) distributions with cutoff 133 KB (55% of packets are from flows larger than the cutoff size). We set the parameters of flow sizes identical to those from Internet traffic characteristics in [36] (Lognormal: $\mu = 9.357$, $\sigma = 1.318$, Pareto: $\alpha = 1.1$), so

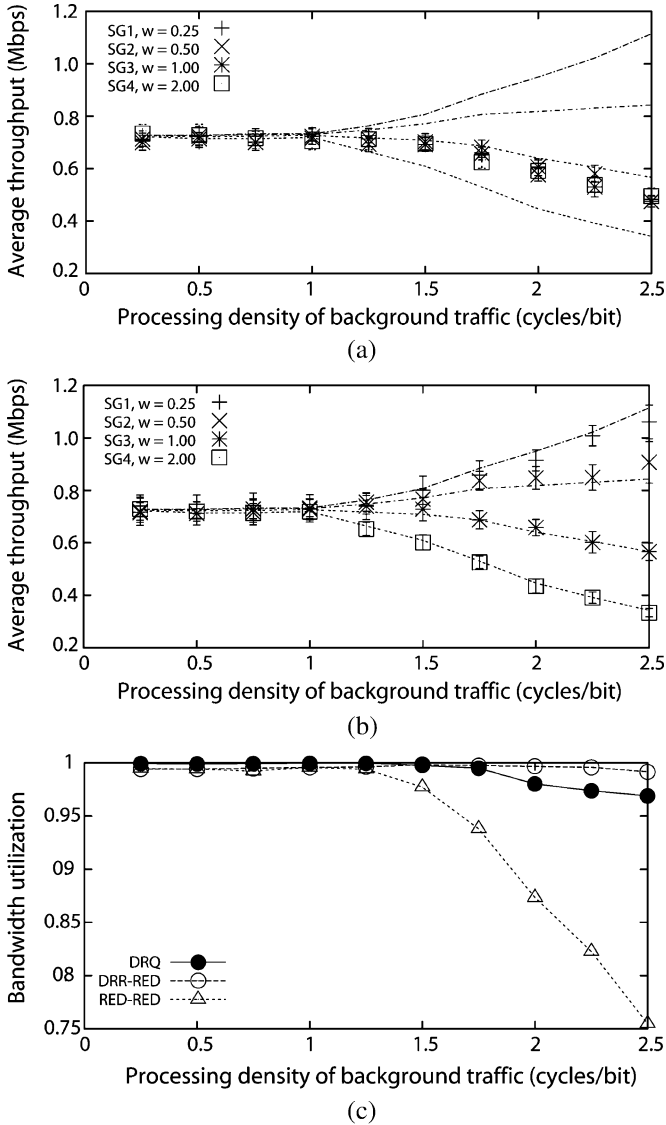


Fig. 8. Simulation result in dumbbell topology with background traffic. Dashed line implies proportional fair rate for each source group. Even with cross traffic, the result is similar to that from the simple model without background traffic. Vertical bars indicate 95% confidence intervals. (a) RED-RED; (b) DRQ; (c) bandwidth utilization.

that a continuous distribution of flow sizes is included in the background traffic. Furthermore, we also generate reverse-path traffic consisting of 10 long-lived TCP flows and a number of short and medium-lived flows to increase realism and also to reduce the phase effect. The RTT of each cross traffic flow is randomly selected from a range of 20 to 60 ms. We fix the bottleneck bandwidth and CPU capacities to 40 Mbps and 40 Mcycles/s, respectively, and generate the same number of long-lived TCP flows as in the experiment for Fig. 5. With these parameters, the cross traffic consumes about 30% of the link bandwidth capacity. In this experiment, we vary the CPU demand of cross traffic to create various bottleneck conditions. We make no claims about how realistically our model characterizes the Internet cross traffic, but we believe that the simulation in this model more realistically reflects real network situations than the one with no background traffic.

Fig. 8 plots the average throughput of long-lived TCP flows, and the total bandwidth utilization of the bottleneck link as we




Fig. 9. Multiple link simulation scenario in parking lot topology.

increase the processing densities of short and medium-lived TCP flows used to generate cross traffic. The results are very similar to those from the simulation without cross traffic; in all situations, RED-RED ensures the equal bandwidth sharing while DRQ and DRR-RED ensure equal sharing of CPU usages under the CPU-limited region and equal sharing of bandwidth under the BW-limited region. In this experiment, we can also find severe under-utilization of bandwidth by RED-RED under the CPU-limited region. As cross traffic helps increase the utilization of bandwidth, we find DRQ and DRR-RED to encourage higher utilization of the bottleneck bandwidth.

E. Parking Lot Simulation With Multiple Dual-Resource Bottleneck Links

To increase realism in our simulation, we simulate the environment where multiple links can be dual-resource constrained. Fig. 9 shows a parking lot topology with following link capacities $B_1 = B_2 = 50$ Mbps, $C_1 = C_2 = C_3 = 40$ Mcycles/s. We vary the bandwidth capacity of link L_3 from 20 to 80 Mbps to create varying bottleneck conditions. Four class of long-lived TCP flows (denoted as $SG1$ to $SG4$) traverse links $L1$, $L2$ and $L3$ from sources 1-4 to sinks 1-4, respectively. One class of long-lived TCP flows from source 5 to sink 5 (denoted as $SG5$) traverses link $L1$ only and another class from source 6 to sink 6 (denoted as $SG6$) traverses link $L2$ only. Additional four sets of long-lived TCP flows traverse link $L3$ only (denoted as $SG7$ to $SG10$). We report the average throughput of long-lived flows ($SG1$ to $SG4$, and $SG7$ to $SG10$) that pass through link $L3$ as we vary the link bandwidth of $L3$. On each path that long-lived TCP flows travel through, we add short and medium-lived background traffic as well as reverse-path traffic to increase dynamics in the network traffic patterns. In this setup, each link runs one of the AQM schemes being evaluated.

Fig. 10(a) and (b) shows the average throughput of the long-lived TCP flows that go through link $L3$. The average throughput values are shown with their corresponding 95% confidence intervals. We also plot with a dotted line the ideal PF rate for each class. The average throughput values closely follow the analytical PF rates.

Fig. 11 shows the total throughput of all flows from $SG1$ to $SG10$ including those that traverse links $L1$ and $L2$. In this figure, we find that DRQ achieves a much higher total throughput than DRR-RED and RED-RED. Although we have consistently seen that DRR-RED and DRQ achieve higher utilization of bandwidth than RED-RED, this is the first time we find DRQ gets more throughput than DRR-RED. This happens because though DRR-RED performs similar to DRQ

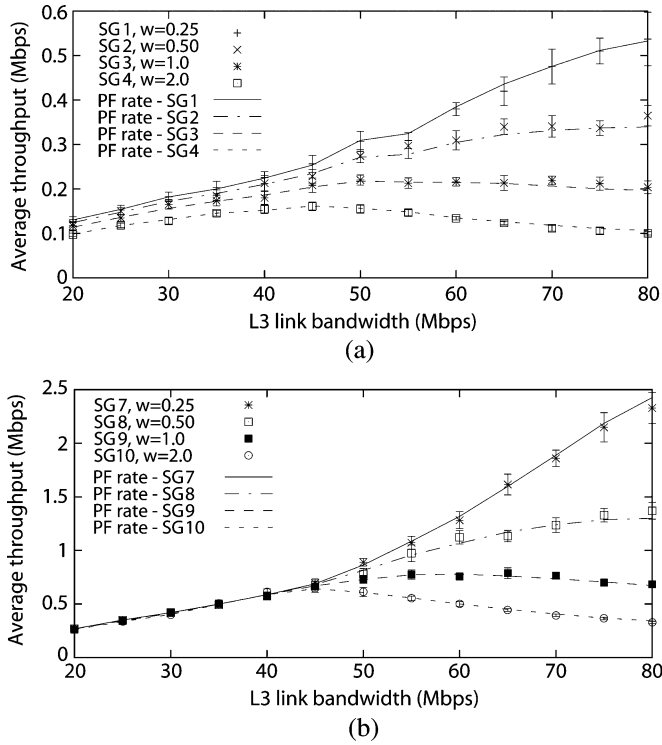


Fig. 10. Throughput of TCP/DRQ in multiple link simulations: (a) SG1-SG4; (b) SG7-SG10. Even in cases where multiple links are dual-resource constrained, DRQ achieves proportional fairness. Vertical bars indicate 95% confidence intervals.




Fig. 11. Total throughput comparison between DRQ, DRR-RED, and RED-RED. It shows the sum of throughput for all flows from SG1 to SG10. In this setup, DRQ achieves the best throughput over the other schemes. This is because the PF rate allocation of DRQ installs fairness among flows that traverse different numbers of hops. RED-RED still consistently shows lower throughput.

when a single resource is congested in the network, in principle, DRR-RED and DRQ have different fairness notions especially when multiple resources are simultaneously congested and flows have different number of hops. DRR has a goal to equally share the resource irrespective of path routes (which is the same goal as max-min fairness), but DRQ follows the proportional fairness criteria. In this topology, SG1-SG4 have a longer route and use more network resources than SG5, SG6, and SG7-SG10. So the total throughput difference comes from DRQ's use of proportional fairness that discriminates SG1-SG4 over the other flows because it uses more hops. The comparison of L3 link utilization shows that DRQ and DRR-RED use nearly the same amount, in Fig. 12. This is because L3 lies in the CPU-limited region so the bandwidth usage of the flows is governed by the fair usage of CPU. This




Fig. 12. L3 link bandwidth utilization comparison between DRQ, DRR-RED, and RED-RED.

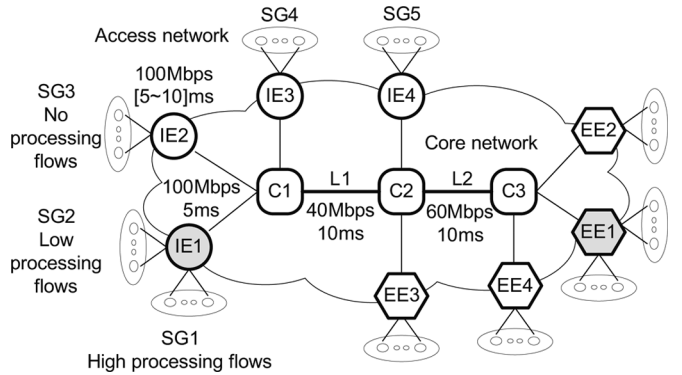


Fig. 13. Simulation topology for edge deployment.

means that the bandwidth usage difference between SG1-SG4, and the other flows occur because of SG5 and SG6 flows that traverse only one link creating a bandwidth bottleneck on their corresponding link. However, it is incorrect to say that the fairness notion of DRQ always guarantees higher throughput than the fairness notion that DRR-RED follows as it is quite possible that there are other unique situations where DRR-RED gets more throughput (which is indeed shown in the next simulation). We leave as future study, studying the exact conditions where DRR-RED can have better throughput than DRQ and vice versa.

F. Impact of Incremental Deployment

In this section, we examine the performance impact of incrementally deploying DRQ in the current Internet. The natural places where DRQ can be initially deployed are likely to be edges. This is because the current Internet trend is to keep the “middle” slim by pushing complicated tasks to the edges of the Internet. Thus, while core routers focus on moving packets as fast as possible, routers, proxies and gateways located at edges of the Internet perform various in-network processing on packets.

We consider an ISP environment where core routers and most edge routers are free from processing constraint but a small number of designated edge gateways handle TCP flows with in-network processing requirement. Fig. 13 models one example of such environments. Our goal is to assess whether a small incremental deployment of DRQ gives any performance advantage while keeping the rest of the network intact. In the figure, SG's denote TCP source groups, each with ten TCP sources, and C1 to C3 denote core routers, IE1 to IE4 denote ingress edge routers and EE1 to EE4 denote egress edge routers, respectively.

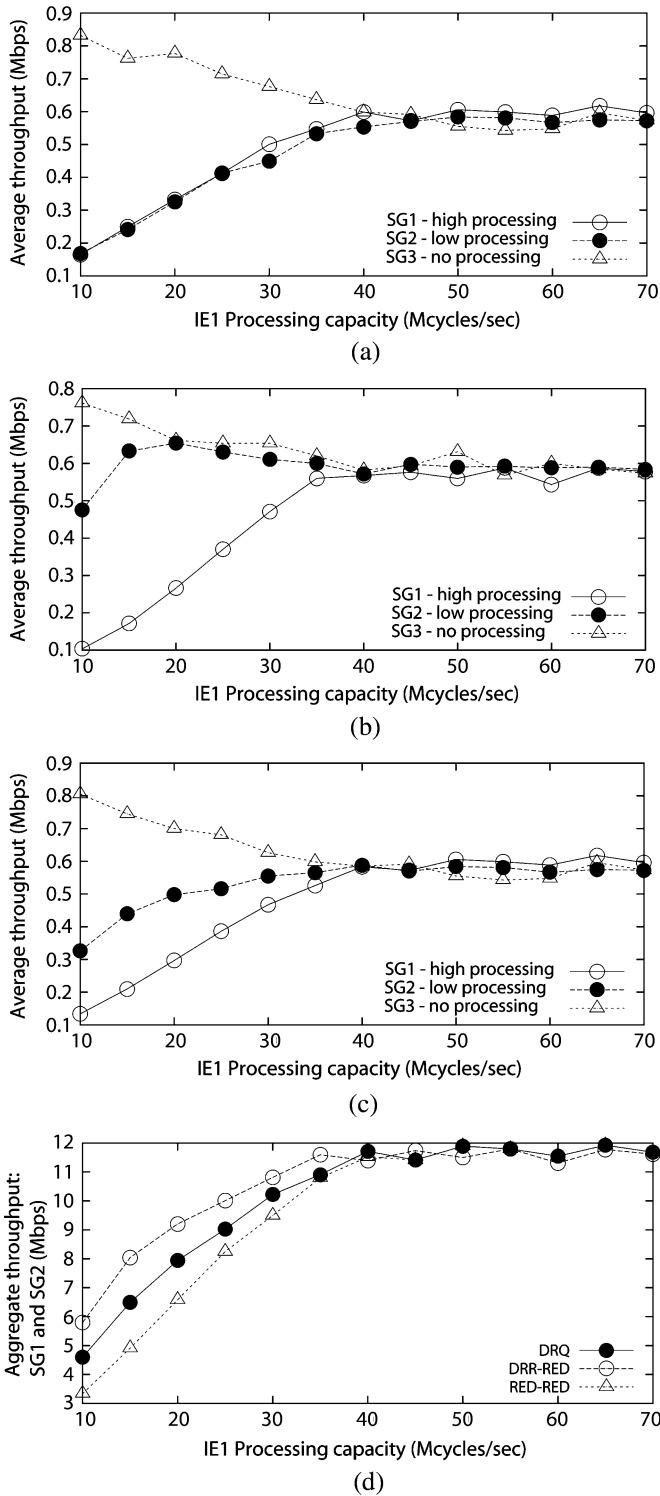


Fig. 14. Average and aggregate throughput of different schemes when some of the edge routers are under the CPU constraint. Even with a small number of edge routers employing DRQ, DRQ can provide higher throughput for the TCP flows with in-network processing requirements. (a) RED-RED; (b) DRR-RED; (c) DRQ; (d) aggregate throughput of flows using IE1.

Flows from IE1 are routed through the shortest path to EE1 etc. In this setup, all the routers except IE1 and EE1 are conventional packet routers with no CPU constraint and IE1 and EE1 are the media gateways/routers that may run in the CPU-limited region.

In this simulation, we consider two different source groups, $SG1$ and $SG2$, that traverse the same end-to-end path from IE1

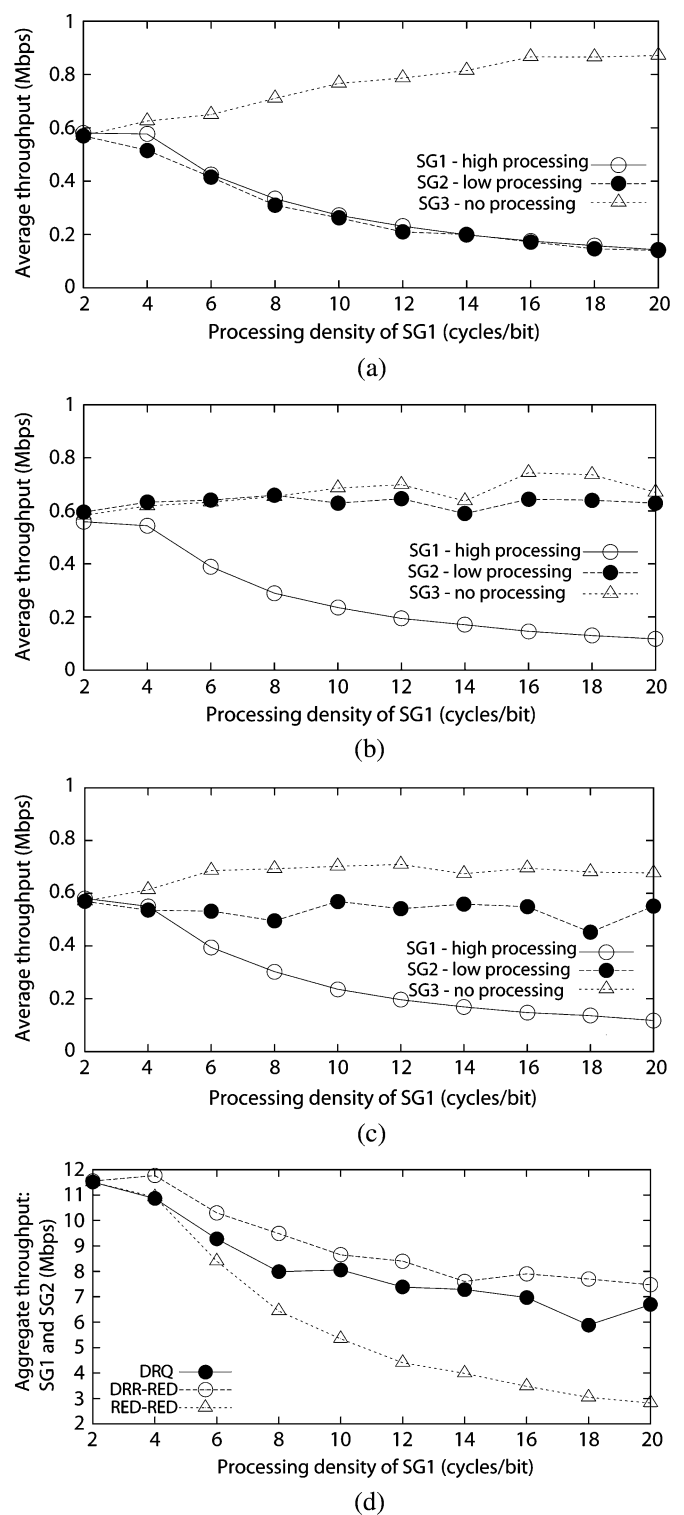


Fig. 15. Average and aggregate throughput when processing density values of neighboring flows ($SG1$) are increased. This shows that DRQ and DRR-RED, even in a partial deployment, achieve fairness in the CPU usage when some portion of flows has increasingly higher processing demands. (a) RED-RED; (b) DRR-RED; (c) DRQ; (d) aggregate throughput of flows using IE1.

to EE1. $SG1$ consists of flows with a high processing density ($w_1 = 2.5$) while $SG2$ consists of flows with a low processing density ($w_2 = 0.5$). Flows from $SG3$ to $SG5$ do not require any CPU processing. Note also that in this setup, $SG1$, $SG2$ and $SG3$ share the same bottleneck links ($L1$ and $L2$).

V. CONCLUSION

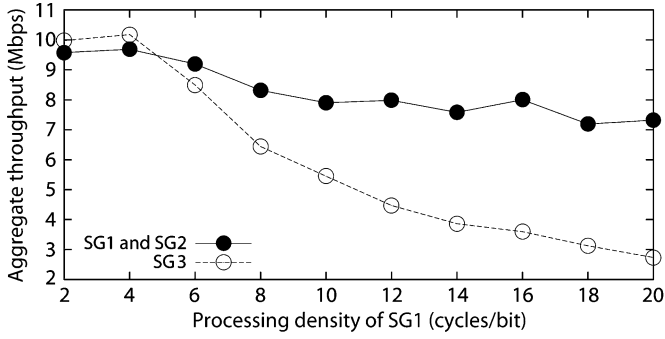


Fig. 16. Aggregate throughput comparison of flows going through DRQ deployed router and flows going through RED-RED deployed router.

In Fig. 14 we compare the average throughput of three source groups $SG1$, $SG2$ and $SG3$, and the aggregate throughput of $SG1$ and $SG2$ while varying the CPU capacity of IE1. As expected, with RED-RED, $SG1$ and $SG2$ achieve the equal bandwidth sharing irrespective of the CPU capacity, which give lower aggregate throughput. We can see that employing the DRQ (or DRR-RED) only at IE1 and EE1 serves more packets than RED-RED by preventing flows with high processing demands from starving those flows with lower processing demands. The aggregate throughput of DRR-RED is slightly higher than that of DRQ.

In the next simulation, we increase the processing density (w_1) of $SG1$ from 2 to 20 to see the impact of flows with high processing demands while fixing the processing density of $SG2$ to a relatively low value ($w_2 = 1.0$) in the same setup as the above and examine whether the fairness achieved by DRQ in a small number of routers contributes to increasing the total throughput of flows that pass through those routers. The CPU capacity of IE1 is fixed to 30 Mcycles/s.

In Fig. 15(a)–(c), we can verify that with DRQ or DRR-RED, the throughput of the high processing flows ($SG1$) is kept lower than that of the low processing flows ($SG2$) to balance its CPU usage with the low processing flows in IE1 and EE1 as we increase the processing density of the high processing flows. That results in much higher aggregate throughput for $SG1$ and $SG2$ than with RED-RED. In RED-RED, because the average throughput of both $SG1$ and $SG2$ flows is maintained to follow equal sharing, the CPU usage (not shown due to the space limitation) of $SG1$ is highly unfair to that of $SG2$ because $SG1$ will be using much CPU than $SG2$ when they have the same number of packets on the network. This phenomenon has a security implication where a malicious flow with a high processing density can easily starve out other competing flows from consuming CPU resources.

Next, we consider another scenario where $SG1$ and $SG2$ are going through DRQ deployed router IE1 but $SG3$ now consists of flows identical to aggregation of $SG1$ and $SG2$ and is going through RED-RED deployed IE2 router. Fig. 16 shows that the flows going through IE1 get much more throughput in total than those going through IE2, particularly when the processing densities of high processing flows ($SG1$ and $SG1$ portion of $SG3$) increase. This also confirms that even an incremental deployment of DRQ at the edge is beneficial.

We have shown that under DRQ the average equilibrium throughput for TCP Reno sources becomes proportionally fair in the dual-resource environment. Moreover, such an equilibrium is unique and almost always exists, more specifically, as long as the per-flow steady-state average TCP window size (or per-flow bandwidth-delay product) is sustainable at a value greater than or equal to 2 packets.

DRQ significantly outperforms RED-RED scheme while maintaining a certain level of fairness. DRQ is scalable to a large number of flows and incrementally deployable since it does not require any change in end-host TCP stacks and builds on an ordinary RED queue.

The throughput gain achieved by DRQ over RED-RED comes mostly from two features of its marking scheme. The weighted marking at a CPU queue (in proportion to processing densities of flows) yields fair sharing of CPU cycles whenever this resource is scarce, which consequently, results in overall throughput increase. Second, the inter-resource marking, which is a unique feature of DRQ, gives extra penalty to flows that traverses more number of resources, which also results in overall throughput increase.

We show by simulation that even a partial deployment of DRQ is beneficial to increasing performance when it is implemented on a few selected locations where special in-network processing services (e.g., media translation, protocol conversion, security and PEP) are enabled.

APPENDIX

Proof of Proposition 1: First, we show that the primal-optimal solution (r^*, θ^*, π^*) of Problem **P** with $\alpha_s = \sqrt{3/2}M_s/\tau_s$ forms an equilibrium point $(x^*, \sigma^*, \rho^*, p^*, y^*, z^*, v^*, \hat{y}^*, \hat{v}^*)$ in the average Reno/DRQ network where $x^* = r^*$, $\sigma^* = \theta^*$, $\rho^* = \pi^*$ if θ^* and π^* satisfy **C1** and **C2**. Since $\sum_{s \in S(l)} r_s^* \leq B_l$, $\forall l \in L$, and $\sum_{s \in S(k)} w_s^k r_s^* \leq C_k$, $\forall k \in K$, (7) and (8) imply that there exist $\hat{y}_l^* \geq 0$, $\forall l \in L$, and $\hat{z}_k^*, v_k^* \geq 0$, $\forall k \in K$. Thus, from (11) and (14), there exist $\hat{y}_l^* = y_l^* \geq 0$, $\forall l \in L$, and $\hat{v}_k^* = v_k^* \geq 0$, $\forall k \in K$, whose specific values can be obtained from (10) and (13) for $\rho_l^* = \pi_l^*$ and $\sigma_k^* = \theta_k^*$ since $0 \leq \pi_l^* < 1$ and $0 \leq \theta_k^* < 1$. Finally, for $x_s^* > 0$, (9) and (15) yield $x_s^* = (\sqrt{3/2}M_s/\tau_s)(\sum_{k \in K(s)} w_s^k \sigma_k^* + \sum_{l \in L(s)} \rho_l^*)^{-1}$, which also holds for (r^*, θ^*, π^*) by the Lagrangian optimality of (r^*, θ^*, π^*) in (6) with $\alpha_s = \sqrt{3/2}M_s/\tau_s$.

Next we prove the converse that if the average Reno/DRQ network has an equilibrium point $(x^*, \sigma^*, \rho^*, p^*, y^*, z^*, v^*, \hat{y}^*, \hat{v}^*)$, then (x^*, σ^*, ρ^*) is the primal-dual optimal solution of the corresponding Problem **P** with $\alpha_s = \sqrt{3/2}M_s/\tau_s$ and, moreover, no other equilibrium points can exist in the network. By the duality theory [22], the necessary and sufficient condition for (x^*, σ^*, ρ^*) to be the primal-dual optimal solution of Problem **P** is that x^* is primal feasible, (σ^*, ρ^*) is dual feasible, and (x^*, σ^*, ρ^*) satisfies Lagrangian optimality and complementary slackness. First, (7) and (8) imply that $\sum_{s \in S(l)} x_s^* \leq B_l$ for all $l \in L$

and $\sum_{s \in S(k)} w_s^k x_s^* \leq C_k$ for all $k \in K$, and (9) implies that $x^* \geq 0$. Thus, x^* is primal feasible. Second, the dual feasibility of (σ^*, ρ^*) is obvious since σ^* and ρ^* cannot be negative by the definition of REDs in (10) and (13). Third, since (6) is the Lagrangian optimality condition of Problem **P**, we need to show that

$$x_s^* = \frac{\alpha_s}{\sum_{k \in K(s)} w_s^k \sigma_k^* + \sum_{l \in L(s)} \rho_l^*}, \quad \forall s \in S \quad (33)$$

hold. Consider flow s . If we suppose that $\sigma_k^* = 0, \forall k \in K(s)$, and $\rho_l^* = 0, \forall l \in L(s)$, for this flow s , then $p_s^* = 0$ from (15), which contradicts (9) since it cannot have an equilibrium point satisfying $p_s^* = 0$. Thus, at least one of $\sigma_k^*, k \in K(s)$, and $\rho_l^*, l \in L(s)$, must be positive, which implies that $0 < p_s^* < 1$ from (15) and consequently $x_s^* > 0$ from (9). Therefore, from (9), we know that $M_s(1 - p_s^*)/\tau_s^2 = (2/3)(x_s^{*2} p_s^*/M_s)$ (N_s is cancelled out). Solving x_s^* from this equation by substituting (15) at equilibrium for p_s^* , we get (33) with $\alpha_s = \sqrt{3/2} M_s / \tau_s$, which concludes the Lagrangian optimality of (x^*, σ^*, ρ^*) . Lastly, we need to show that $\sigma_k^*(\sum_{s \in S(k)} w_s^k x_s^* - C_k) = 0$ and $\rho_l^*(\sum_{s \in S(l)} x_s^* - B_l) = 0$ for all $k \in K$ and $l \in L$ to check complementary slackness of (x^*, σ^*, ρ^*) . Consider arbitrary $k \in K$. If we suppose that $\sum_{s \in S(k)} w_s^k x_s^* - C_k < 0$, then $z_k^* = v_k^* = 0$ from (8). Thus, $\hat{v}_k^* = 0$ from (14) and consequently $\sigma_k^* = 0$ from (13), which concludes $\sigma_k^*(\sum_{s \in S(k)} w_s^k x_s^* - C_k) = 0$. Similarly, $\rho_l^*(\sum_{s \in S(l)} x_s^* - B_l) = 0$ for arbitrary $l \in L$. Therefore, we conclude that (x^*, σ^*, ρ^*) is the primal-dual optimal solution of Problem **P**. Moreover, x^* is unique since the optimal solution of Problem **P** is unique. From (10), it is obvious that $y_l^* > b_l$ if $\rho_l^* > 0$ and $0 \leq y_l^* \leq b_l$ otherwise. The same argument can be applied to CPU queues from (13).

Proof of Proposition 2: First, we derive a linearized model of Reno/DRQ network at the equilibrium point. In the DRQ network described by (15), the congestion window for the s th TCP source, ϖ_s , can be modelled by the nonlinear differential equation

$$\dot{\varpi}_s(t) = \frac{1 - p_s(t)}{N_s \tau_s} - \frac{2\varpi_s^2(t)}{3N_s \tau_s} p_s(t) \quad (34)$$

where

$$p_s(t) = \frac{\left(\sum_{i=1}^{|R|} [\Gamma(0)^T]_{si} \varphi_i(t - \tau_{is}) \right)^2}{1 + \left(\sum_{i=1}^{|R|} [\Gamma(0)^T]_{si} \varphi_i(t - \tau_{is}) \right)^2} \quad (35)$$

and $\varphi_i(t)$ is the i th element of unified pre-marking probability set as

$$\varphi_i = \begin{cases} \sigma_i, & \text{if } i \text{ indicates CPU} \\ \rho_i, & \text{if } i \text{ indicates link.} \end{cases}$$

Using (7), (8), we model the i th congested resource queue by

$$\dot{q}_i(t) = -\tilde{c}_i I_{q>0} + \sum_{s=1}^{|S|} [\Gamma_{is}(0)] \frac{M_s \varpi_s(t - \tau_{is})}{\tau_s} \quad (36)$$

where $q_i(t)$ is the i th element of the resource queue length set in bits (or cycles) as

$$q_i = \begin{cases} z_i, & \text{if } i \text{ indicates CPU} \\ y_i, & \text{if } i \text{ indicates link} \end{cases}$$

\tilde{c}_i is the i th element of resource capacity set as

$$\tilde{c}_i = \begin{cases} C_i, & \text{if } i \text{ indicates CPU} \\ B_i, & \text{if } i \text{ indicates link.} \end{cases}$$

M_s denotes the packet size in bits, and $I_{q>0}$ is the indicator function.

Linearizing (34)–(36) about equilibrium $(\varpi^*, q^*, \varphi^*)$ and taking Laplace transforms gives

$$\varpi(z) = -\tilde{F}(z)\Gamma(-z)^T \varphi(z) \quad (37)$$

$$q(z) = (zI + \Omega)^{-1} \Gamma(z) M T^{-1} \varpi(z) \quad (38)$$

where $\tilde{F}(z) = \text{diag}\{\tilde{f}_s(z)\}$

$$\tilde{f}_s(z) = \frac{e^{-z\tau_s}}{z + \frac{4\varpi_s^*}{3N_s\tau_s} \frac{3}{2\varpi_s^{*2}+3}} \frac{2}{N_s\tau_s u_s^* (1+u_s^*)^2}$$

$u^* = \Gamma(0)^T \varphi^*$, $\Omega = \Gamma(0) M W^* T^{-2} \Gamma(0)^T \tilde{C}^{-1}$, $T = \text{diag}\{\tau_s\}$, $M = \text{diag}\{M_s\}$, $W^* = \text{diag}\{\varpi_s^*\}$ and $\tilde{C} = \text{diag}\{\tilde{c}_i\}$.

Next, let $\kappa_i(z)$ denote the linearized RED dynamics which can be modelled as the low-pass filter [25]

$$\kappa_i(z) = \frac{\varphi_i(z)}{q_i(z)} = \frac{\psi_i}{z/\eta_i + 1} \quad (39)$$

where $\psi_i = m_i/(\bar{b}_i - b_i)$, and $\eta = \log_e(1 - \lambda)/h$, where h is sampling interval of RED. Using linearized (37)–(39), we can find the return ratio function of the DRQ network

$$\mathcal{L}(z) = \tilde{F}(z)\Gamma(-z)^T D(z)\tilde{C}^{-1}(zI + \Omega)^{-1}\Gamma(z)MT^{-1} \quad (40)$$

where $D(z) = \text{diag}\{\kappa_i(z)\tilde{c}_i\}$. To invoke the Generalized Nyquist Stability Criterion [32], we should show that the eigenvalues of $\mathcal{L}(jw)$ do not intersect $(-\infty, -1]$ for all $w \geq 0$.

From *Theorem 1* and *Proposition 1* in [32], these eigenvalues lie in

$$\frac{\max\{|f_s(jw)|\} \max\{|d_i(jw)|\}}{\min\{|g_i(jw)|\}} \quad (41)$$

where $f_s(jw) = (\mu_s/\varpi_s)\tilde{f}_s(jw)$, μ_s is the number of congested resources traversed by the s th source, $d_i(jw) = \kappa_i(jw) \cdot \tilde{c}_i$, $g_i(jw) = jw + \lambda_i$, and λ_i is an eigenvalue of $\tilde{C}^{-(1/2)}\Omega\tilde{C}^{-(1/2)}$. To be stable, (41) should be less than 1.

Because $\min\{|g_i(jw)|\}$ is an increasing function, $\max\{|f_s(jw)|\}$ and $\max\{|d_i(jw)|\}$ are decreasing, stability condition can be replaced by

$$\frac{\max\{|f_s(0)|\} \max\{|d_i(0)|\}}{\min\{|g_i(0)|\}} < 1. \quad (42)$$

We can find the following relations:

$$|g_i(0)| \geq \frac{\phi_{min}^2[\Gamma(0)]}{\Lambda_{max}\tau_{max}^2w_{max}} \quad (43)$$

$$|f_s(0)| \leq \frac{|R|\Lambda_{max}\tau_{max}}{\sqrt{3/2}} \quad (44)$$

where $\Lambda_{max} = \max\{\max\{C_k\}/\min\{w_s^k\} \min\{M_s\}, \max\{B_l\}/\min\{M_s\}\}$, $\tau_{max} = \max\{\tau_s\}$, $w_{max} = \max\{w_s^k, 1\}$, and $\phi_{min}[\Gamma(0)]$ denotes the smallest singular values of the matrix $\Gamma(z)$ evaluated at $z = 0$.

From (42)–(44)

$$\max\{|d_i(0)|\} \leq \frac{\sqrt{3/2}\phi_{min}^2[\Gamma(0)]}{|R|\Lambda_{max}^2\tau_{max}^3w_{max}}. \quad (45)$$

REFERENCES

- [1] R. Ramaswamy, N. Weng, and T. Wolf, "Analysis of network processing workloads," in *Proc. IEEE Int. Symp. Performance Analysis of Systems and Software (ISPASS)*, Austin, TX, Mar. 2005, pp. 226–235.
- [2] P. Crowley, M. E. Fiuczynski, J.-L. Baer, and B. N. Bershad, "Workloads for programmable network interfaces," in *Workload Characterization for Computer System Design*. Boston, MA: Kluwer Academic, 2000, ch. 7.
- [3] Y. Gottlieb and L. Peterson, "A comparative study of extensible routers," in *Proc. IEEE Open Architectures and Network Programming (OPENARCH)*, Jun. 2002, pp. 51–62.
- [4] A. T. Campbell, H. G. D. Meer, M. E. Kounavis, K. Miki, J. B. Vicente, and D. Vilella, "A survey of programmable networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 2, pp. 7–23, April 1999.
- [5] A. Barbir, R. Penno, R. Chen, M. Hofmann, and H. Orman, "An architecture for Open Pluggable Edge Services (OPES)," IETF RFC 3835, 2004.
- [6] B. Falchuk, J. Chiang, A. Hafid, Y.-H. Cheng, N. Natarjan, F. J. Lin, and H. Cheng, "An open service platform for deploying and managing services at network edges," in *Proc. IEEE Open Architectures and Network Programming (OPENARCH)*, San Francisco, CA, 2003.
- [7] S. Chong, M. Shin, J. Mo, and H.-W. Lee, "Flow control with processing constraint," *IEEE Commun. Lett.*, vol. 9, no. 10, pp. 957–959, Oct. 2005.
- [8] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulation of a fair queueing algorithm," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 19, no. 4, pp. 1–12, 1989.
- [9] M. Shreedhar and G. Varghese, "Efficient fair queueing using deficit round robin," in *Proc. ACM SIGCOMM*, Sep. 1995, pp. 231–242.
- [10] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.
- [11] S. Athuraliya, S. H. Low, V. H. Li, and Q. Yin, "REM: Active queue management," *IEEE Network*, vol. 15, pp. 48–53, May 2001.
- [12] C. V. Hollot, V. Misra, D. Towsley, and W. B. Gong, "On designing improved controllers for AQM routers supporting TCP flows," in *Proc. IEEE INFOCOM*, Anchorage, AK, Apr. 2001, pp. 1726–1734.
- [13] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless fair queueing: Achieving approximately fair bandwidth allocations in high speed networks," in *Proc. ACM SIGCOMM*, Sep. 1998, pp. 118–130.
- [14] Z. Cao, Z. Wang, and E. W. Zegura, "Rainbow fair queueing: Fair bandwidth sharing without per-flow state," in *Proc. IEEE INFOCOM*, Tel Aviv, Israel, Mar. 2000, pp. 922–931.
- [15] P. Pappu and T. Wolf, "Scheduling processing resources in programmable routers," in *Proc. IEEE INFOCOM*, New York, Jun. 2002, pp. 104–112.
- [16] F. Sabrina, C. D. Nguyen, S. K. Jha, D. Platt, and F. Safaei, "Processing resource scheduling in programmable networks," *Comput. Commun.*, vol. 28, no. 6, pp. 676–687, Apr. 2005.
- [17] T. Wolf and M. A. Franklin, "CommBench—A telecommunications benchmark for network processors," in *Proc. IEEE Int. Symp. Performance Analysis of Systems and Software (ISPASS)*, Austin, TX, Apr. 2000, pp. 154–162.
- [18] G. Memik, M. Smith, and W. Hu, "NetBench: A benchmarking suite for network processors," in *Proc. IEEE Int. Conf. Computer-Aided Design*, San Jose, CA, Nov. 2001, pp. 39–42.
- [19] M. Tsai, C. Kulkarni, C. Sauer, N. Shah, and K. Keutzer, "A benchmarking methodology for network processors," in *Proc. 1st Network Processor Workshop, 8th Int. Symp. High Performance Computer Architectures (HPCA)*, Cambridge, MA, Feb. 2002.
- [20] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, pp. 237–252, Apr. 1998.
- [21] S. H. Low and D. E. Lapsley, "Optimization flow control I: Basic algorithm and convergence," *IEEE/ACM Trans. Networking*, vol. 7, no. 6, pp. 861–875, Dec. 1999.
- [22] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.
- [23] S. H. Low, "A duality model of TCP and queue management algorithms," *IEEE/ACM Trans. Networking*, vol. 11, no. 4, pp. 525–536, Aug. 2003.
- [24] F. P. Kelly, "Mathematical modelling of the internet," in *Mathematics Unlimited—2001 and Beyond*, B. Engquist and W. Schmid, Eds. Berlin, Germany: Springer Verlag, 2001, pp. 685–702.
- [25] V. Misra, W. B. Gong, and D. Towsley, "Fluid-based analysis of a network of AQM routers supporting TCP flows with an application to RED," in *Proc. ACM SIGCOMM*, Sep. 2000, pp. 151–160.
- [26] Y. Liu, F. Presti, V. Misra, and D. Towsley, "Fluid models and solutions for large-scale IP networks," in *Proc. ACM SIGMETRICS*, San Diego, CA, Jun. 2003, pp. 151–160.
- [27] C. V. Hollot, V. Misra, D. Towsley, and W.-B. Gong, "A control theoretic analysis of RED," in *Proc. IEEE INFOCOM*, Anchorage, AK, Apr. 2001, pp. 1510–1519.
- [28] S. H. Low, F. Paganini, J. Wang, and J. C. Doyle, "Linear stability of TCP/RED and a scalable control," *Comput. Netw.*, vol. 43, pp. 633–647, Dec. 2003.
- [29] Internet End-to-End Performance Monitoring Group (IEPM) 2006 [Online]. Available: <http://www-iepm.slac.stanford.edu/>
- [30] K. Ramakrishnan and S. Floyd, "A proposal to add Explicit Congestion Notification (ECN) to IP," IETF RFC 2481, 1999.
- [31] K. Ramakrishnan, S. Floyd, and D. Black, "The addition of Explicit Congestion Notification (ECN) to IP," IETF RFC 3168, 2001.
- [32] H. Han, C. V. Hollot, Y. Chait, and V. Misra, "TCP networks stabilized by buffer-based AQMs," in *Proc. IEEE INFOCOM*, Hong Kong, Mar. 2004, pp. 964–974.
- [33] ns-2 Network Simulator. 2000 [Online]. Available: <http://www.isi.edu/nsnam/ns/>
- [34] S. Floyd, "Recommendations on using the Gentle variant of RED," Mar. 2000 [Online]. Available: <http://www.aciri.org/floyd/red/gentle.html>
- [35] J. Aikat, J. Kaur, F. D. Smith, and K. Jeffay, "Variability in TCP round-trip times," in *Proc. 3rd ACM SIGCOMM Conf. Internet Measurement Conference*, Oct. 2003, pp. 279–284.
- [36] P. Barford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in *Proc. ACM SIGMETRICS*, Jul. 1998, pp. 151–160.
- [37] V. Paxson and S. Floyd, "Wide-area traffic: The failure of Poisson modeling," *IEEE/ACM Trans. Networking*, vol. 3, pp. 226–244, Jun. 1995.
- [38] K. Park, *Self-Similar Network Traffic and Performance Evaluation*, W. Willinger, Ed. New York: Wiley, 2000.
- [39] H. Jiang and C. Dovrolis, "Passive estimation of TCP round-trip times," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 32, pp. 75–88, July 2002.



Minsu Shin (S'99) received the B.S., M.S., and Ph.D. degrees from the School of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, in 1998, 2000, and 2006, respectively.

Since July 2006, he has been a Visiting Postdoctoral Research Fellow at the Department of Computer Science, North Carolina State University, Raleigh. His research interests are in congestion control, programmable networks, and wireless networks.



Song Chong (M'93) received the B.S. and M.S. degrees in control and instrumentation engineering from Seoul National University, Seoul, Korea, in 1988 and 1990, respectively, and the Ph.D. degree in electrical and computer engineering from the University of Texas at Austin in 1995.

Since March 2000, he has been with the School of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, where he is an Associate Professor, leading the Network Systems Laboratory.

Prior to joining KAIST, he was with the Performance Analysis Department, AT&T Bell Laboratories, Holmdel, NJ, as a Member of Technical Staff. His research interests include high speed networks, wireless networks and performance evaluation. He has published more than 60 papers in international journals and conferences and holds three U.S. patents in these areas.

Dr. Chong is an Editor of the *Journal of Communications and Networks* and has served as a Technical Program Committee member of a number of international conferences.



Injong Rhee (SM'89) received the Ph.D. degree from the University of North Carolina at Chapel Hill.

He is an Associate Professor of computer science at North Carolina State University, Raleigh. His areas of research interests include computer networks, congestion control, wireless ad hoc networks and sensor networks. He works mainly on network protocol designs optimizing the transport performance of networks.