Technical Report: Achieving Quality of Service via Packet Distribution Shaping

Duong Nguyen-Huu, Thai Duong, Thinh Nguyen
June 15th, 2012

***Abstract*—Conventional Quality of Service (QoS) for multimedia networking applications are typically specified by a certain set of requirements on latency, jitter, bandwidth, and packet loss rate. In this paper, we introduce a novel approach to QoS via the notion of distribution shaping in which a pre-specified distribution of packets in a queue is achieved via queuing policies. In a way, the distribution-based QoS is more general since the distribution of packets in the queue captures all the statistical information regarding the throughput, latency, delay jitter, and packet loss rate. We present a convex optimization framework for obtaining the optimal queueing policy that drives any initial distribution of packets in the queue to the desired distribution in the fastest time. We then augment the proposed framework to obtain a queueing policy that produces ϵ-approximation to the target distribution with even faster convergence time. The augmented framework is useful in dynamic settings where traffic statistics change frequently, and thus fast adaptation is preferable. Both simulation and theoretical results are provided to verify our approach.**

I. Introduction

Guaranteeing end-to-end Quality of Service (QoS) for multimedia applications over a best-effort network such as the Internet is challenging. The difficulty is due to a number of factors including the time-varying nature of Internet traffic as well as the heterogeneity of network architectures and policies across different autonomous domains (AS). A major effort has been focused on DiffServ architecture [1] in which packets of different flows are classified and marked at the ingress routers. The markings are then used by the intermediate routers to determine their forwarding/queueing policies. For example, packets with Expedited Forwarding (EF) marking are intended for flows/applications with low-loss, low-latency such as video conference traffic. The intermediate routers then implement certain queuing policies that ensure the EF packets have higher forwarding priority than other best effort packets. In a way, this is an attempt to provide scalable end-to-end QoS by enforcing differentiated service of flows on a per-hop behavior basis.

Achieving QoS via enforcing per-hop behavior for a flow is also a common approach in wireless networks to map the end-to-end QoS requirements into packet transmission policies at each hop. For example, in a local wireless area network (WLAN), using the MAC protocol 802.11e in the Enhanced Distributed Channel Access (EDCA) mode [2], packets are classified into different types: Background (AC_BK), Best Effort (AC_BE), Video (AC_VI), Voice (AC_VO). The minimum and maximum contention window (CWmin, CWmax) and Arbitration Inter-Frame Space (AIFS) are primary parameters to control the priorities for different packet types. A flow using small contention windows and AIFS will have higher chance to access the wireless medium. For example, CWmax for best-effort packets is set to 1023 while it is set to 32 for video packets.

Another approach to provisioning flows of different priorities is to employ multiple physical or virtual queues at a router. Each queue consists of packets of the same type. A queueing policy is used at each transmission opportunity, to decide which of the queues whose a packet should be transmitted. A simple fair queuing policy will transmit packets from each non-empty queue in a round robin fashion [3]. On the other hand, a priority or weighted queueing policy give preference for transmitting packets from higher priority queues [4].

All the aforementioned techniques aim to achieve QoS under resource constraints. These queueing policies are often not designed to provide statistical guarantee for bounding the maximum packet delay or loss. Rather, they are used to provide differentiated services among the flows. Under well-specified network traffic conditions, it is possible to derive the average and variance of packet latencies and loss rates for different flows when a particular queuing policy is used. Although these queuing polices are primarily heuristic driven. That said, queuing policy plays a critical role in providing QoS.

Contribution. As discussed, queuing policies play a critical role in providing QoS for multimedia applications. Thus, in this paper, we introduce a novel approach to QoS via the notion of distribution shaping in which a pre-specified distribution of packets in a queue is achieved via appropriate queuing policies. In a way, the distribution-based QoS is more general and precise than the previous approaches since the distribution of packets in the queue captures all the statistical information regarding the throughput, latency, delay jitter, and packet loss rate. In particular, given a distribution it is theoretically possible to compute its moments of any order. We present a convex optimization framework for obtaining the optimal queueing policy that drives any initial distribution of packets in the queue to the target distribution in the fastest time. While there are many queueing policies that can produce the target distribution, the fastest policy is preferable since it allows fast adaptation to time-varying network traffic. We then extend the proposed framework for finding a queueing policy that produces ϵ-approximation to the given target distribution with even faster convergence time. This extended framework allows for a trade-off between how close the obtained policy to the desired policy and how fast it can be obtained, which can be useful in fast changing network environments.

Our paper is organized as follows. In Section II, we discuss the approach to QoS via distribution shaping. In Section III, we provide necessary mathematical notations and background. Section IV is devoted to convex optimization formulation for finding the optimal policy and some theoretical results. We present simulation results for the proposed approach in Section V. Finally, we provide a few concluding remarks in Section VI.

II. Achieving QoS via Distribution Shaping

A Simple Example. We first illustrate the distribution-based approach to QoS with an example of a time-discrete version of the classical $M/M/1/k$ queueing model. In this time-discrete model, time is divided into time steps of equal duration. At the beginning of each time step, exactly one packet arrives at the queue with probability $p = 0.4$. Otherwise, with probability $1 - p = 0.6$, no packet arrives during that entire time step. Assume a queueing policy is used such that at the beginning of each time step, exactly one packet is dequeued with probability $q = 0.6$. Otherwise, with probability $1 - q = 0.4$, no packet is dequeued during that entire time step. Furthermore, let $k = 2$ be the maximum queue size, and a newly arrived packet is dropped if the queue is full. The dynamic of the number of packets in the queue can be shown to be governed by the following transition probability matrix:

$$P = \begin{pmatrix} 0.84 & 0.16 & 0 \\ 0.36 & 0.48 & 0.16 \\ 0 & 0.36 & 0.64 \end{pmatrix},$$

where P_{ij} denotes the probability that the queue will have j packets in the next time step, given that it currently has i packets with $i, j \in \{0, 1, 2\}$. For each aperiodic and irreducible P, there exists a unique corresponding stationary distribution π such that $\pi^T P = \pi^T$. In this particular case,

$$\pi = \begin{pmatrix} 0.61 \\ 0.27 \\ 0.12 \end{pmatrix}.$$

The stationary distribution π characterizes the long term or stationary probability of the queue occupancy. In this case, out of all the observed time slots, 61% of time the queue is empty, 27% of the time the queue has exactly one packet, and 12% of the time the queue has two packets. Knowing exactly this distribution, the average queuing delay can be precisely calculated. One can also immediately bound the probability of dropped packets to no more than 0.12. In fact, any statistical measure, e.g., moments of any order can be theoretically calculated for a given distribution.

Transition Probability is induced by Queuing Policy. Suppose the QoS requirements are given in terms of maximum average packet latency and minimum packet drop rate, then one can find a stationary distribution π that satisfies such requirements. However, there are many transition probability matrices P that have the same stationary distribution π. It is important to note that each transition probability P is a result of applying a certain queuing policy. For the example above, the associated queueing policy is to send packets with probability of 0.6. One can easily implement another policy that sends packets with a different probability which results in a different transition probability. Moreover, we need not restrict ourselves to the class of policies that sends packets with a fixed probability. Rather, one can design a policy that sends packets with different probabilities based on the number of packets presently in the queue.

Constraints on Queuing Policy. Intuitively, for a high priority flow $\pi = [1, 0, 0]^T$ seems to be the best stationary distribution since the queue is always empty. However, this implies that a packet is always dequeued at every time slot. This policy might not be possible or optimal due to several reasons. For example, let us consider a wireless network consisting of multiple nodes. First, if an application does not require much throughput, then sending packets all the time consumes more power than necessary. Second, if every node in the wireless network implements the same greedy queuing policy, then collisions will happen all the time, resulting in low overall throughput. Thus, the transition probability matrix (hence the queuing policy) must be selected from a pre-specified class of transition probability matrices that gives rise to reasonable queuing policies for the given settings. This constraint is an input to our convex optimization framework to be described shortly.

Fastest Queuing Policy. We noted above that there are many transition probability matrices P (equivalently many queuing policies) that have the same given stationary distribution π, and all satisfy the pre-specified QoS requirements. So which transition probability matrix should one choose? The theory of Markov chain shows that if we apply the same queuing policy over many time steps, the distribution of packets in the queue will converge to a unique stationary distribution corresponding to a stochastic, aperiodic and irreducible matrix P, regardless of the initial distribution of packets in the queue. Mathematically, let ν be any initial distribution, then

$$\lim_{n \to \infty} \nu^T P^n = \pi^T, \tag{1}$$

where n is the number of time steps.

If the network traffic is stationary, then π can be obtained approximately using the same queuing policy after some number of time steps. Ideally, we want the queuing policy that drives the distribution of packets in the queue to the desired stationary distribution fastest for any initial distribution. This is especially useful when the network conditions change and thus fast adapting queueing policy is preferable.

Another important point about this fast adapting principle is that if for some reason, the network traffic becomes bursty for a short while that fill up the queue, a fast queuing policy will drive the queue to the desired stationary distribution in the fastest time.

A Slightly More Sophisticated Example. Let us consider the following example of a wireless network of consisting of three nodes running three applications with different priorities, e.g., video, audio, and data. A queue can be used to model the

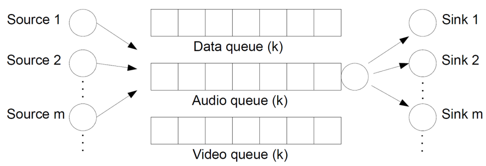

Fig. 1. A queuing model in a wireless network

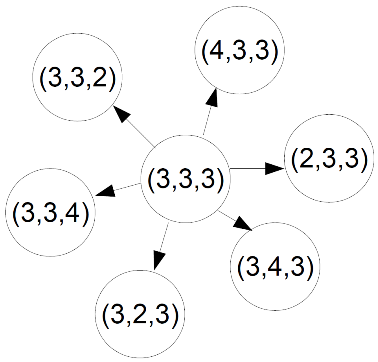

Fig. 2. The possible states at time step $i + 1$, given the state at time i is (3,3,3)

packet generation and processing at each node. Specifically, packets arriving at the queue are generated by an application, and a (de)queuing policy determines the rate at which the packets are sent out. Furthermore, suppose three different distributions of packets in the queues are given as QoS requirements for these flows as shown in Fig. 1. Intuitively, the distribution for the audio flow should be skewed toward having high probability for smaller number of packets in the queue in order to ensure low queuing latency for its packets. On the other hand, the distribution for the data flow can be less skewed since the requirement on the latency of its packet is less stringent. In the simple single queue example with maximum queue size k, the number of possible states for the stationary distribution is $k + 1$, corresponding to the possible number of packets in the queue. In the three queue example, the corresponding stationary distribution π is now a joint distribution of the number of packets in each queue, and thus the number of possible states is $(k + 1)^3$. Then, given a queuing policy one can obtain the transition probability matrix $P^{(k+1)^3 \times (k+1)^3}$ that specifies the state of all three queues simultaneously.

It is convenient to represent the dynamics of packets in the multiple queues via a walk on a graph. Fig. 2 depicts the state transition using a graph. If at time step i, the system state is (3,3,3), corresponding to the number of packets in each queue is 3, then at time step $(i+1)$, the system can be in six possible states with different probabilities. This is due to the fact that the number of packets in each node (queue) can only increase or decrease by exactly one, or stays the same.

In general, for any system of m queues in a wireless network, the system dynamics can be represented by a finite undirected connected graph $G(V, E)$ with V and E denoting a set of vertices (set of states) and set of edges (transition probability), respectively. A weighted edge between two vertices v_i and v_j represents a non-zero probability for the system to transition from state i to state j in one time step with appropriate queueing policy. A multiple-queue system can thus be viewed as a random walk on weighted graph $G = (V, E)$. This walk forms a reversible Markov chain which is important in our later discussion.

Our framework. For a given setting, e.g., m queues in a wireless network, one can always immediate write down the structure of the corresponding graph. For a given queuing policy and network condition (SNR), we can also compute the weights on the edge of the graph which corresponds to the transition probability matrix P. For a given transition probability matrix P, a stationary distribution π can be found.

Our first problem is how to find a queueing policy that drives the system to a pre-specified target stationary distribution π in the fastest time. More generally, given a graph with its connectivity, how to put the weights on the graph, i.e., transition probability matrix P, so that the corresponding walk on that graph will reach the desired stationary distribution quickly. Our second problem is that rather than finding a queueing policy that is guaranteed to drive the system to the target stationary distribution, we find a queueing policy that drives the system ϵ-close to the target stationary distribution in much faster time.

Our paper is focused on finding the transition probability P within a set of specified class of transition probabilities, specifically tridiagonal matrices. This restriction is imposed by the real-world constraints on the possible queuing policies as discussed previously. We note that having P, on can then find the queuing policy appropriately. For example, in the single queue example, knowing the entries of the matrix P and the network environments (p), the queuing policy, i.e., sending rate (q) can be obtained without much difficulty. Next, we will present the necessary background for formulating and solving the proposed problems.

III. Mathematical Notations and Preliminaries

In this section, we provide notations, definitions, and a few well-known results to be used in our convex optimization framework.

Proposition 1: For an irreducible, aperiodic, finite and discrete Markov chain with a transition probability matrix P, there exists a unique stationary distribution π such that

$$\lim_{n\to\infty} \nu^T P^n = \pi^T. \tag{2}$$

Definition 2 (Total variation distance): For any two probability distributions ν and π on a finite state space Ω, we define the total variation distance as:

$$\|\nu - \pi\|_{TV} = \frac{1}{2} \sum_{i\in\Omega} |\nu(i) - \pi(i)| .$$

Total variance is the common metric to measure the distance between two probability distributions.

Definition 3 (Mixing time): For a discrete, aperiodic and irreducible Markov chain with transition probability P and stationary distribution π, given an $\epsilon > 0$, the mixing time $t_{mix}(\epsilon)$ is defined as

$$t_{mix}(\epsilon) = \inf\{n : \|\nu^T P^n - \pi^T\|_{TV} \le \epsilon, \text{for all probability distributions } \nu\}.$$

The mixing time of a Markov chain measures the time needed for the Markov chain to converge to within an ϵ of its stationary distribution. The notion of mixing time is critical to our framework to characterize the convergence rate of a queuing policy.

A non-zero vector v_i is called a right (left) eigenvector of a square matrix P if there is a scalar λ_i such that: $Pv_i = \lambda_i v_i$ or ($v_i^T P = \lambda v_i^T$). The scalar λ_i is said to be an eigenvalue of P. If P is a stochastic matrix, then $|\lambda_i| \le 1, \forall i$.

Denote the set of eigenvalues in non-increasing order:

$$1 = \lambda_1(P) \ge \lambda_2(P) \ge \cdots \ge \lambda_{|\Omega|}(P) \ge -1$$

Definition 4 (Second largest eigenvalue modulus): The second largest eigenvalue modulus (SLEM) of a matrix P is defined as:

$$\mu(P) = \max_{i=2,\ldots,|\Omega|} |\lambda_i(P)| = \max\{\lambda_2(P), -\lambda_{|\Omega|}(P)\} \quad (3)$$

Definition 5 (Reversible Markov Chain): A discrete Markov chain with a transition probability P is said to be reversible if

$$P_{ij}\pi(i) = P_{ji}\pi(j) \quad (4)$$

Theorem 1 (Bound on mixing time): [5]. Let P be the transition matrix of a reversible, irreducible and aperiodic Markov chain with state space Ω, and let $\pi_{min} := min_{x\in\Omega}\pi(x)$. Then

$$t_{mix}(\epsilon) \le \frac{1}{1-\mu(P)} \log\big(\frac{1}{\epsilon\pi_{min}}\big). \quad (5)$$

Theorem 1 shows that a transition matrix P with a smaller $\mu(P)$ would have faster convergence rate to the stationary distribution. It is not difficult to see that from Theorem 1, the error ϵ reduces over time at a rate of no greater than $\frac{e^{-(1-\mu(P))t}}{\pi_{min}}$. Thus, finding the matrix P with minimum $\mu(P)$ would result in the fastest convergence time which will be the topic in the next section.

IV. Convex Optimization Formulation

It was shown in [6] that

$$\mu(P) = ||D_\pi^{1/2} P D_\pi^{-1/2} - \sqrt{\pi}(\sqrt{\pi})^T||_2, \quad (6)$$

where π denotes the stationary distribution of P, D_π denotes the square diagonal matrix whose diagonal entries are taken from each elements of π, and $||.||$ denote l_2-induced matrix norm. Furthermore, $\mu(P)$ is a convex function in P.

For a queueing system, the corresponding transition probability matrix is tridiagonal since the number of packets in the queue can only increase, decrease, or remain the same in the next time step. Also, all tridiagonal matrices are reversible. Therefore, our first convex optimization is: Given the QoS requirements, i.e., a desired stationary distribution of packets in the queue, find the fastest queuing policy (P) that drives the queue from any state to the desired stationary distribution. It was first formulated broadly in [6] as:

Problem 1.

$$\text{Minimize } ||D_{\pi^*}^{1/2} P D_{\pi^*}^{-1/2} - \sqrt{\pi^*}(\sqrt{\pi^*})^T||_2$$
$$\text{Subject to : } \begin{cases} P\mathbf{1} = \mathbf{1} \\ D_{\pi^*}P = P^T D_{\pi^*} \\ \text{other convex constraints on } P. \end{cases} \quad (7)$$

The objective function is SLEM. The first constraint ensures P is a stochastic matrix. The second constraint is for reversibility. The third constraint is imposed by limitations of certain settings, e.g., interference, power consumption, etc, as discussed in Section II. The solution of the problem (if exists) is a transition matrix P_{opt} which has the smallest SLEM, resulting fastest convergence time to the given target distribution π^*. However, these constraints, especially the third constraint, can be restricted that given a stationary distribution π^*, there might not be a P that simultaneously satisfies all the constraints and produces the desired stationary distribution. For example, in the single queue example, if one restricts the queuing policy to always send packets at some constant rate (q) regardless of how many packets in the queue, then there is less flexibility in producing the desired π^*. In addition, in many settings, finding a queueing policy that produces a stationary distribution that is within some small ϵ of the target stationary distribution, but has faster convergence time might be preferable. This is especially useful when network conditions changes quickly. On the other hand, a slow adapting queueing policy is optimal for the past rather than the present network conditions. Based on this, we propose the following optimization problem (P2):

Problem 2.

$$\text{Minimize } ||D_{\pi}^{1/2} P D_{\pi}^{-1/2} - \sqrt{\pi}(\sqrt{\pi})^T||_2$$
$$\text{Subject to : } \begin{cases} P\mathbf{1} = \mathbf{1} \\ D_{\pi}P = P^T D_{\pi} \\ \text{other convex constraints on } P. \\ ||\pi^* - \pi|| \le \epsilon \end{cases} \quad (8)$$

The optimization variables in (P2) are both P and π. Unfortunately, (P2) is non-convex as evident by the constraint $||\pi^* - \pi|| \le \epsilon$. Therefore, we propose the following convex problem (P3) to find the approximate solution for (P2).

Problem 3.

$$\text{Minimize } ||D_{\pi^*}^{1/2} P D_{\pi^*}^{-1/2} - \sqrt{\pi^*}(\sqrt{\pi^*})^T||_2$$
$$\text{Subject to : } \begin{cases} P\mathbf{1} = \mathbf{1} \\ ||\pi^T * P - \pi^{*T}||_2 \le \delta \\ \text{Other convex constraints on } P. \end{cases} \quad (9)$$

Unlike (P2), P is the only optimization variable in (P3). It is not difficult to see that (P3) is convex. One issue to consider is how to pick δ in the constraint $||\pi^T * P - \pi^*||_2 \le \delta$, so that the solution to (P3) indeed satisfies all the constraints in (P2). Specifically, we want to determine the bound on the value of

δ to guarantee that the constraint $||\pi^* - \pi|| \leq \epsilon$ in problem (P2) is satisfied. We have the following proposition.

Proposition 6: For any irreducible aperiodic reversible P, we have:

$$||\pi^* - \pi||_2 \leq \frac{\pi_{max}^{1/2}}{\pi_{min}^{1/2}} \frac{||\pi^* P - \pi^*||_2}{1 - \lambda_2}. \quad (10)$$

Proof: See appendix. ■

From Proposition 6, it is straightforward to see that if we pick $\delta \geq \epsilon\sqrt{\frac{\pi_{min}}{\pi_{max}}}(1-\lambda_2)$, then $||\pi^* - \pi||_2 \leq \epsilon$. On the other hand, we cannot possibly know π_{min}, π_{max}, and λ_2 without knowing P first. However, one often can put upper and lower bounds on these quantities by looking the structure of the class of the transition matrix. For example, one can bound λ_2 via the conductance obtained by examining the corresponding graph $G(V, E)$ [5]. Specifically, we have the following results on the upper and lower bounds of the quantities above for a class of tridiagonal transition probability matrices that represents a class of queueing policies.

Proposition 7: Let P be a tridiagonal matrix with $\alpha \leq P_{ij} \leq \beta; (0 < \alpha < \beta)$ for all (i, j) in the off-diagonal line, we have

$$\begin{cases} \pi_{min} \geq \alpha^{|\Omega|-1} \\ \pi_{max} \leq \beta \\ \lambda_2 \leq 1 - 2\alpha^{|\Omega|} \end{cases}$$

Proof: We omit the proof due to limited space. ■

Using Proposition 7, it is not difficult to obtain the following corollary for selecting the right δ based on ϵ.

Corollary 1: For the class of tridiagonal matrices defined in Proposition 7, pick $\delta = \epsilon\frac{2\alpha^{(5|\Omega|-1)/2}}{\beta^{1/2}}$ will guarantee that

$$||\pi^* - \pi||_2 \leq \epsilon \quad (11)$$

We are ready to show the main result on bounding the optimal objective value of problem (P2) with that of problem (P3). We have the following proposition:

Proposition 8: Let the μ_2 and μ_3 be the optimal objective values of problems (P2) and P(3), respectively. Let $\Delta = \frac{\epsilon}{\sqrt{\pi^*_{min}}}$, π^*_{min} and π^*_{max} denote the maximum and minimum elements in π^*, respectively. Then,

$$\mu_2 \geq \mu_3 \geq \mu_2 - C, \quad (12)$$

where

$$\begin{aligned} C &= \frac{\Delta(2\sqrt{\pi^*_{min}} - \Delta)}{(\sqrt{\pi^*_{min}} - \Delta)^2} + (\sqrt{\pi^*_{max}} + 2\Delta)\frac{\Delta^2}{{\pi^*_{min}}^{3/2}} \\ &+ |\Omega|\Delta(2\sqrt{\pi^*_{max}} + 3\Delta) \end{aligned} \quad (13)$$

Proof: See Appendix. ■

Proposition 8 provides a bound on using solution to (P3) as an approximate solution for (P2).

V. Simulation Results

In this section, we present simulation results which agree with our intuitions and theoretical results. We use CVX [7] to solve all our convex problems. CVX routines implement subgradient methods for finding the optimal solutions. Since the all the matrices under consideration are relatively small, the time for CVX to obtain the solutions are negligible.

We assume a single discrete time queuing system in which at most one packet arrives or departs in a single time slot. This implies that the transition probability matrix is a tridiagonal matrix, and therefore is reversible as shown below.

$$P = \begin{pmatrix} r_1 & p_1 & & & \\ q_2 & r_2 & p_2 & & \\ & \ddots & \ddots & \ddots & \\ & & q_{|\Omega|-1} & r_{|\Omega|-1} & p_{|\Omega|-1} \\ & & & q_{|\Omega|} & r_{|\Omega|} \end{pmatrix} \quad (14)$$

To model the limitations on power consumption, inferences, etc., we further require that: $r_i, p_i, q_i \in (\alpha, \beta), \forall i$. $(\alpha, \beta) \subset (0, 1)$ are pre-specified that models certain limitations. Specifically, we set $(\alpha, \beta) = (0.05, 0.95)$, the maximum queue size $|\Omega| = 20$, and $\delta = 0.1$. The given target stationary distribution π^* is shown in Fig. 3.

First, in case (a), we consider a limited class of queuing policies where it can be modeled as a tridiagonal matrix with the following requirement:

$$r_i = 0 \text{ for } i = 2, \ldots, |\Omega| - 1$$

Given π^*, we solve problem (P1) to find the fastest policy that converges to π^*. Now in case (b), we enlarge the class of queueing policies by lifting the restriction on $r_i = 0$. Then, we solve problem (P1). Intuitively, the queuing policy found in case (b) should likely to have faster convergence time than that of case (a) since it is found from a larger class of policies. Indeed, this is the case. Fig. 4 shows the total variation distance between the target stationary distribution and the current distribution as a function of time steps. As seen, the curve for case (a) decreases slower than that of case (b). At the time step $n = 300$, the total variation distance for case (b) is almost zero while that of (a) is still around 0.08. Hence, queuing policy in (b) is more suitable for changing network conditions.

We now consider case (c). In this case, the class of queueing policies is the same at that of case (b). However, we solve problem (P3) in which, we intentionally find a queuing policy that might not produce exactly the target stationary distribution π^*, but close enough, i.e., $||\pi - \pi^*||_2 \leq \epsilon$. Intuitively, this policy should produce even faster adaptation than those of cases (a) and (b). In fact, this is the case. Fig. 4 shows the curve for case (c) which drops down quickly compared with the other two. At time $n = 50$, the total variation distance is 0.1284 for case (c) while they are more than 0.7 for the other two cases. The curve for case (c) however does not converge, i.e., decreases to zero, but stays around 0.12. This is intuitive since the solution to problem (P3) is not designed to

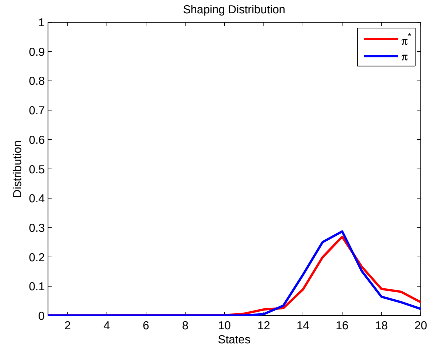

Fig. 3. Target and resulted distributions in case (c)

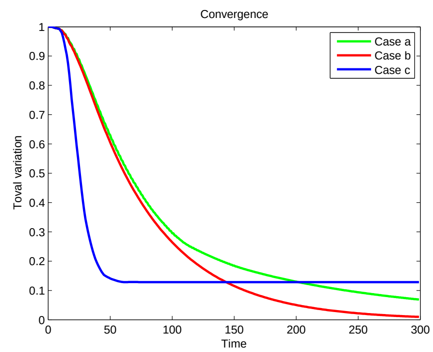

Fig. 4. Comparison of the convergence times in 3 cases

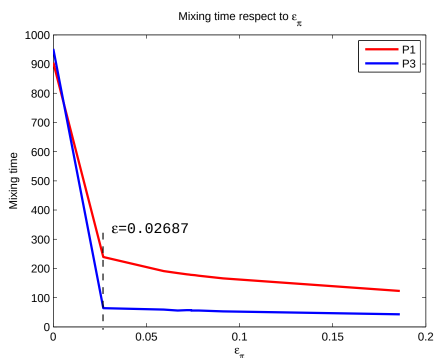

Fig. 5. Mixing time in case (b) and case (c) with respect to ϵ_π

obtain a queuing policy that converges to the target stationary distribution. It is important to point out that in a fast-changing network environment, it is preferable to use the queuing policy obtained by solving problem (P3) since it allows fast adaption at the expense of a bit less accurate. Fig. 3 shows that there are not much difference in the distribution π obtained by solving problem (P3) and the target distribution π^*. Thus, the QoS requirements would not be violated by using the queuing policy obtained from problem (P3).

We now study the trade-off between the accuracy of obtaining the target distribution and the convergence time. Fig. 5 shows the mixing times from problems (P3) and (P1) which decrease significantly when the allowable deviation (ϵ) from the stationary distribution increases. For the class of queuing policies in the simulation, setting $\epsilon = 0.02687$ seems to be the best as it reduces the mixing time significantly while keeping π close to π^*.

An interesting observation. We note that the resulted queuing polices above behave like a 802.11 protocol to some extent. To see this, assuming a model for a simple queue as in Section II, and the network condition is stationary, i.e., p is fixed. Then each different entries P_{ij} might correspond to a different q_i. This implies that, the optimal queuing polices will dynamically change the sending rates based on the current number of packets in the queue. For the 802.11 protocol, a node will change its sending rate based on the number of collisions. Under some settings, the number of queues in the packets can be highly related to the number of collisions. In short, both 802.11 and the proposed optimal queuing are similar in the way they adapt to the network conditions.

VI. Conclusion

In this paper, we introduce a novel approach to designing queueing policies that provides statistical guarantees on QoS requirements. Specifically, the queuing policy is designed to shape the distribution of packets in the queue to a target distribution which captures the QoS requirements on the throughput, latency, delay jitter, and packet loss rate. We present a convex optimization framework for obtaining the optimal queueing policy that drives any initial distribution to a target distribution in the fastest time. We then show how to extend the proposed technique to obtain a queueing policy that produces ϵ-approximation to the given distribution with even faster convergence time. The former is useful in settings whose network conditions change slowly, while the later is appropriate for fast-changing network conditions. Both simulation and theoretical results verify the benefits of the proposed approach.

References

[1] S. Blake, D. Black, M. Carson, E. Davis, Z. Wang, and W. Weiss, "An architecture for differentiated services," in *RFC2475*, December 1998.

[2] I. S. 802.11e/D6.0, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Quality of Service (QoS) Enhancements*, 2003.

[3] J. Nagle, "On packet switches with infinite storage," *Communications, IEEE Transactions on*, vol. 35, no. 4, pp. 435 – 438, apr 1987.

[4] D. Stiliadis and A. Varma, "Latency-rate servers: a general model for analysis of traffic scheduling algorithms," *Networking, IEEE/ACM Transactions on*, vol. 6, no. 5, pp. 611 –624, oct 1998.

[5] D. A.Levin, Y. Peres, and E. L.Wilmer, *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

[6] S. Boyd, P. Diaconis, and L. Xiao, "Fastest mixing markov chain on a graph," *SIAM Review*, vol. 46(4), pp. 667–689, December 2004.

[7] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21," ../../cvx, Apr. 2011.

Appendix

Proposition 6

Proof: (sketch)

We assume P has n eigenvalues $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$ and n left eigenvectors $\{v_1, v_2, \dots, v_n\}$ such that: $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq -1$.

Let $\langle f,g\rangle_{\frac{1}{\pi}} := \sum_{x\in\Omega} \frac{f(i)g(i)}{\pi(i)}$ denote the inner product with respect to $\pi(x)$. Due to the reversibility of P, it can be shown that the set of eigenvectors $\{v_i\}$ forms an orthonormal basis with $\langle .,.\rangle_{\frac{1}{\pi}}$. The eigenvector corresponds to the largest eigenvalue $\lambda_1 = 1$ is equal to the stationary distribution: $v_1 = \pi$. We have:

$$\pi^{*T} - \pi^T = \sum_{i=1}^{n} \langle \pi^* - \pi, v_i\rangle_{\frac{1}{\pi}} v_i^T$$

Since $v_i^T P = \lambda_i v_i^T$,

$$(\pi^{*T} - \pi^T)(P - I) = \sum_{i=1}^{n} (\lambda_i - 1)\langle \pi^* - \pi, v_i\rangle_{\frac{1}{\pi}} v_i^T$$

Also,

$$\begin{aligned}\langle \pi^* - \pi, v_1\rangle_{\frac{1}{\pi}} &= \langle \pi^* - \pi, \pi\rangle_{\frac{1}{\pi}} \\ &= \sum_{i=1}^{n} (\pi^*(i) - \pi(i)) \\ &= 0\end{aligned}$$

Then

$$||\pi^{*T} - \pi^T||_{\frac{1}{\pi}} = ||\pi^* - \pi||_{\frac{1}{\pi}} = \sqrt{\sum_{i=2} \langle \pi^* - \pi, v_i\rangle^2_{\frac{1}{\pi}}}$$

and

$$||(\pi^{*T} - \pi^T)(P - I)||_{\frac{1}{\pi}} = \sqrt{\sum_{i=2} (\lambda_i - 1)^2 \langle \pi^* - \pi, v_i\rangle^2_{\frac{1}{\pi}}}$$

Therefore:

$$||(\pi^{*T} - \pi^T)(P - I)||_{\frac{1}{\pi}} \geq \min_{i=2,\dots,n} |1 - \lambda_i| ||(\pi^* - \pi)||_{\frac{1}{\pi}}$$

$$\rightarrow ||(\pi^{*T} - \pi^T)(P - I)||_{\frac{1}{\pi}} \geq (1 - \lambda_2)||(\pi^* - \pi)||_{\frac{1}{\pi}}$$

$$\rightarrow ||(\pi^{*T} * P - \pi^{*T})||_{\frac{1}{\pi}} \geq (1 - \lambda_2)||(\pi^* - \pi)||_{\frac{1}{\pi}}$$

Since for any vector x:

$$\frac{||x||_2}{\sqrt{\pi_{min}}} \geq ||x||_{\frac{1}{\pi}} \geq \frac{||x||_2}{\sqrt{\pi_{max}}}$$

Then we conclude

$$||\pi^* - \pi||_2 \leq \frac{\pi_{max}^{1/2}}{\pi_{min}^{1/2}} \frac{||\pi^* P - \pi^*||_2}{1 - \lambda_2}$$

■

Proposition 8

Proof: (sketch)

Denote a vector $s = \sqrt{\pi^*} - \sqrt{\pi}$ then $|s_i| \leq \Delta \ \forall i \in \Omega$ where $\Delta = \frac{\epsilon_\pi}{\sqrt{\pi^*_{min}}} > 0$

Using Taylor series for function $f(x) = \frac{1}{c+x}$ at point $x = 0$ in the interval $x \in (-\Delta, \Delta)$, we have:

$$\frac{1}{\sqrt{\pi_i^*}} = \frac{1}{\sqrt{\pi_i + s_i}} = \frac{1}{\sqrt{\pi_i}} - \frac{1}{\pi_i} s_i + R_i$$

where R_i is the Taylor Remainder then $|R_i| \leq \frac{1}{{\pi^*_{min}}^{3/2}}\Delta^2$. Denote R is a vector whose entries are R_i then

$$\begin{cases} D_{\pi^*}^{1/2} = D_\pi^{1/2} + D_s \\ D_{\pi^*}^{-1/2} = D_\pi^{-1/2} - D_{s/\pi} + D_R \end{cases}$$

We also denote:

$$\begin{cases} A = D_{\pi^*}^{1/2} P D_{\pi^*}^{-1/2} - \sqrt{\pi^*}(\sqrt{\pi^*})^T \rightarrow \mu_3 = ||A||_2 \\ B = D_\pi^{1/2} P D_\pi^{-1/2} - \sqrt{\pi}(\sqrt{\pi})^T \rightarrow \mu_2 = ||B||_2 \end{cases} \tag{15}$$

Then we have:

$$\begin{aligned} A &= (D_\pi^{1/2} + D_s)P(D_\pi^{-1/2} - D_{s/\pi} + D_R) \\ &\quad - (\sqrt{\pi} + s)(\sqrt{\pi} + s)^T \\ &= B + D_s P D_\pi^{-1/2} - D_\pi^{1/2} P D_{s/\pi} \\ &\quad - D_s P D_{s/\pi} + D_\pi^{1/2} P D_R + D_s P D_R \\ &\quad - s(\sqrt{\pi})^T - \sqrt{\pi} s^T - s s^T \end{aligned} \tag{16}$$

Since $||P|| = 1$, using sub-multiplicative property of matrix norm each element in the right side of (16) (except B) can be bound as following:

$$\begin{cases} ||D_s P D_\pi^{-1/2}|| \leq max_i |\frac{s_i}{\sqrt{\pi_i}}| = \frac{\Delta}{\sqrt{\pi^*_{min}} - \Delta} \\ ||D_\pi^{1/2} P D_{s/\pi}|| \leq max_i |\frac{s_i}{\sqrt{\pi_i}}| = \frac{\Delta}{\sqrt{\pi^*_{min}} - \Delta} \\ ||D_s P D_{s/\pi}|| \leq max_i |\frac{s_i^2}{\pi_i}| = \frac{\Delta^2}{(\sqrt{\pi^*_{min}} - \Delta)^2} \\ ||D_\pi^{1/2} P D_R|| \leq max_i |\sqrt{\pi_i}||R_i| = (\sqrt{\pi^*_{max}} + \Delta)\frac{\Delta^2}{{\pi^*_{min}}^{3/2}} \\ ||D_s P D_R|| \leq \max_i |s_i||R_i| = \frac{\Delta^3}{{\pi^*_{min}}^{3/2}} \\ ||s(\sqrt{\pi})^T|| \leq |\Omega| \max_i |s_i(\sqrt{\pi_i^*} - s_i)| = |\Omega|\delta(\sqrt{\pi^*_{max}} + \Delta) \\ ||(\sqrt{\pi})s^T|| \leq |\Omega| \max_i |s_i(\sqrt{\pi_i^*} - s_i)| = |\Omega|\delta(\sqrt{\pi^*_{max}} + \Delta) \\ ||s s^T|| \leq |\Omega| \max_i |s_i^2| = |\Omega|\Delta^2 \end{cases}$$

Sum up all these elements, we now have:

$$\begin{aligned} ||A - B|| \leq C &= \frac{\Delta(2\sqrt{\pi^*_{min}} - \Delta)}{(\sqrt{\pi^*_{min}} - \Delta)^2} + (\sqrt{\pi^*_{max}} + 2\Delta)\frac{\Delta^2}{{\pi^*_{min}}^{3/2}} \\ &\quad + |\Omega|\Delta(2\sqrt{\pi^*_{max}} + 3\Delta) \end{aligned} \tag{17}$$

Also

$$\min ||A|| \geq \min ||B|| - \max ||A - B||$$

From (15), we have:

$$\mu_3 \geq \mu_2 - D \tag{18}$$

Clearly, solution set of (9) includes that of (8) then

$$\mu_3 \leq \mu_2 \tag{19}$$

(17), (18) and (19) complete the proof.

■