Energy-Aware Wireless Scheduling with Near Optimal Backlog and Convergence Time Tradeoffs

Michael J. Neely
University of Southern California

***Abstract*— This paper considers a wireless link with randomly arriving data that is queued and served over a time-varying channel. It is known that any algorithm that comes within ϵ of the minimum average power required for queue stability must incur average queue size at least $\Omega(\log(1/\epsilon))$. However, the optimal convergence time is unknown, and prior algorithms give convergence time bounds of $O(1/\epsilon^2)$. This paper shows that it is possible to achieve the optimal $O(\log(1/\epsilon))$ average queue size tradeoff with an improved convergence time of $O(\log(1/\epsilon)/\epsilon)$. Further, this is shown to be within a logarithmic factor of the best possible convergence time. The method uses the simple drift-plus-penalty technique with an improved convergence time analysis.**

I. Introduction

This paper considers power-aware scheduling in a wireless link with a time-varying channel and randomly arriving data. Arriving data is queued for eventual transmission. The transmission rate out of the queue is determined by the current channel state and the current power allocation decision. Specifically, the controller can make an *opportunistic scheduling* decision by observing the channel before allocating power. For a given $\epsilon > 0$, the goal is to push average power to within ϵ of the minimum possible average power required for queue stability while ensuring optimal queue size and convergence time tradeoffs.

A major difficulty is that the data arrival rate and the channel probabilities are unknown. Hence, the convergence time of an algorithm includes the *learning time* associated with estimating probability distributions or "sufficient statistics" of these distributions. The optimal learning time required to achieve the average power and backlog objectives, as well as the appropriate sufficient statistics to learn, are unknown. This open question is important because it determines how fast an algorithm can adapt to its environment. A contribution of the current paper is the development of an algorithm that, under suitable assumptions, provides an optimal power-backlog tradeoff while provably coming within a logarithmic factor of the optimal convergence time. This is done via the existing drift-plus-penalty algorithm but with an improved convergence time analysis.

Work on opportunistic scheduling was pioneered by Tassiulas and Ephremides in [1], where the Lyapunov method and the *max-weight* algorithms were introduced for queue stability. Related opportunistic scheduling work that focuses on utility optimization is given in [2][3][4][5][6][7][8][9] using dual, primal-dual, and stochastic gradient methods, and in [10] using index policies. The basic drift-plus-penalty algorithm of Lyapunov optimization can be viewed as a *dual method*, and is known to provide, for any $\epsilon > 0$, an ϵ-approximation to minimum average power with a corresponding $O(1/\epsilon)$ tradeoff in average queue size [9][11]. This tradeoff is not optimal. Work by Berry and Gallager in [12] shows that, for queues with *strictly concave* rate-power curves, any algorithm that achieves an ϵ-approximation must incur average backlog of $\Omega(\sqrt{1/\epsilon})$, even if that algorithm knows all system probabilities. Work in [13] shows this tradeoff is achievable (to within a logarithmic factor) using an algorithm that does *not* know the system probabilities. The work [13] further considers the exceptional case when rate-power curves are *piecewise linear*. In that case, an improved tradeoff of $O(\log(1/\epsilon))$ is both achievable and optimal. This is done using an exponential Lyapunov function together with a drift-steering argument. Work in [14][15] shows that similar logarithmic tradeoffs are possible via the basic drift-plus-penalty algorithm with Last-in-First-Out scheduling.

Now consider the question of *convergence time*, being the time required for the average queue size and power guarantees to kick in. This convergence time question is unique to problems of *stochastic scheduling when system probabilities are unknown*. If probabilities were known, the optimal fractions of time for making certain decisions could be computed offline (possibly via a very complex optimization), so that system averages would "kick in" immediately at time 0. Thus, convergence time in the context of this paper should not be confused with *algorithmic complexity* for non-stochastic optimization problems.

Unfortunately, prior work that treats stochastic scheduling with unknown probabilities, including the basic drift-plus-penalty algorithm as well as extensions that achieve square root and logarithmic tradeoffs, give only $O(1/\epsilon^2)$ convergence time guarantees. Recent work in [16] treats convergence time for a related problem of flow rate allocation and concludes that constraint violations decay as $c(\epsilon)/t$, where $c(\epsilon)$ is a constant that depends on ϵ and t is the total time the algorithm has been in operation. While the work [16] does not specify the size of the $c(\epsilon)$ constant, it can be shown that $c(\epsilon) = O(1/\epsilon)$. Intuitively, this is because the $c(\epsilon)$ value is related to an average queue size, which is $O(1/\epsilon)$. The time t needed to ensure constraint violations are at most ϵ is found by solving $c(\epsilon)/t = \epsilon$. The simple answer is $t = O(1/\epsilon^2)$, again exhibiting $O(1/\epsilon^2)$ convergence time! This leads one

The author is with the Electrical Engineering department at the University of Southern California, Los Angeles, CA.
This work is supported by the NSF Career grant CCF-0747525.

to suspect that $O(1/\epsilon^2)$ is *optimal*.

This paper shows, for the first time, that $O(1/\epsilon^2)$ convergence time is *not optimal*. Specifically, under the same piecewise linear assumption in [13], and for the special case of a system with just one queue, it is shown that the existing drift-plus-penalty algorithm yields an ϵ-approximation with both $O(\log(1/\epsilon))$ average queue size and $O(\log(1/\epsilon)/\epsilon)$ convergence time. This is an encouraging result that shows learning times for power-aware scheduling can be pushed much smaller than expected.

The next section specifies the problem formulation. Section III shows a lower bound on convergence time of $\Omega(1/\epsilon)$. Section IV develops an algorithm that achieves this bound to within a logarithmic factor.

II. System model

Consider a wireless link with randomly arriving traffic. The system operates in slotted time with slots $t \in \{0, 1, 2, \ldots\}$. Data arrives every slot and is queued for transmission. Define:

$$\begin{aligned} Q(t) &= \text{queue backlog on slot } t \\ a(t) &= \text{new arrivals on slot } t \\ \mu(t) &= \text{service offered on slot } t \end{aligned}$$

The values of $Q(t), a(t), \mu(t)$ are nonnegative and their units depend on the system of interest. For example, they can take integer units of *packets* (assuming packets have fixed size), or real units of *bits*. Assume the queue is initially empty, so that $Q(0) = 0$. The queue dynamics are:

$$Q(t + 1) = \max[Q(t) + a(t) - \mu(t), 0] \quad (1)$$

Assume that $\{a(t)\}_{t=0}^{\infty}$ is an independent and identically distributed (i.i.d.) sequence with mean $\lambda = \mathbb{E}[a(t)]$. For simplicity, assume the amount of arrivals in one slot is bounded by a constant a_{max}, so that $0 \leq a(t) \leq a_{max}$ for all slots t.

If the controller decides to transmit data on slot t, it uses one unit of power. Let $p(t) \in \{0, 1\}$ be the power used on slot t. The amount of data that can be transmitted depends on the current channel state. Let $\omega(t)$ be the amount of data that can be transmitted on slot t if power is allocated, so that:

$$\mu(t) = p(t)\omega(t)$$

Assume that $\omega(t)$ is i.i.d. over slots and takes values in a finite set $\Omega = \{\omega_0, \omega_1, \omega_2, \ldots, \omega_M\}$, where $\omega_0 = 0$ and ω_i is a positive real number for all $i \in \{1, \ldots, M\}$. Assume these values are ordered so that:

$$0 = \omega_0 < \omega_1 < \omega_2 < \cdots < \omega_M$$

For each $\omega_k \in \Omega$, define $\pi(\omega_k) = Pr[\omega(t) = \omega_k]$.

Every slot t the system controller observes $\omega(t)$ and then chooses $p(t) \in \{0, 1\}$. The choice $p(t) = 1$ activates the link for transmission of $\omega(t)$ units of data. Fewer than $\omega(t)$ units are transmitted if $Q(t) < \mu(t)$ (see the queue equation (1)). The largest possible average transmission rate is $\mathbb{E}[\omega(t)]$, which is achieved by using $p(t) = 1$ for all t. It is assumed throughout that $0 \leq \lambda \leq \mathbb{E}[\omega(t)]$.

A. Optimization goal

For a real-valued random process $b(\tau)$ that evolves over slots $\tau \in \{0, 1, 2, \ldots\}$, define its time average expectation over $t > 0$ slots as:

$$\overline{b}(t) \triangleq \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E}[b(\tau)] \quad (2)$$

where "$\triangleq$" represents "defined to be equal to." With this notation, $\overline{\mu}(t)$, $\overline{p}(t)$, $\overline{Q}(t)$ respectively denote the time average expected transmission rate, power, and queue size over the first t slots.

The basic stochastic optimization problem of interest is:

$$\begin{aligned} \text{Minimize:} \quad & \limsup_{t\to\infty} \overline{p}(t) & (3) \\ \text{Subject to:} \quad & \liminf_{t\to\infty} \overline{\mu}(t) \geq \lambda & (4) \\ & p(t) \in \{0, 1\} \quad \forall t \in \{0, 1, 2, \ldots\} & (5) \end{aligned}$$

The assumption $\lambda \leq \mathbb{E}[\omega(t)]$ ensures the above problem is always *feasible*, so that it is possible to satisfy constraints (4)-(5) using $p(t) = 1$ for all t. Define p^* as the infimum average power for the above problem. An algorithm is said to produce an *ϵ-approximation* at time t if, for a given $\epsilon \geq 0$:

$$\begin{aligned} \overline{p}(t) &\leq p^* + \epsilon \\ \lambda - \overline{\mu}(t) &\leq \epsilon \end{aligned}$$

An algorithm is said to produce an *$O(\epsilon)$-approximation* if the ϵ symbols on the right-hand-side of the above two inequalities are replaced by some constant multiples of ϵ.

Fix $\epsilon > 0$. This paper shows that a simple drift-plus-penalty algorithm that takes ϵ as an input parameter (and that has no knowledge of the arrival rate or channel probabilities) can be used to ensure there is a time T_ϵ, called the *convergence time*, for which:

- The algorithm produces an $O(\epsilon)$-approximation for all $t \geq T_\epsilon$.
- The algorithm ensures the following for all $t \in \{0, 1, 2, \ldots\}$:
$$\overline{Q}(t) \leq O(\log(1/\epsilon)) \quad (6)$$
- $T_\epsilon = O(\log(1/\epsilon)/\epsilon)$.

The average queue size bound (6) is known to be optimal, in the sense that no algorithm can provide a sub-logarithmic guarantee [13]. The next section shows that the convergence time $O(\log(1/\epsilon)/\epsilon)$ is within a logarithmic factor of the optimal convergence time.

III. A lower bound on convergence time

A. Intuition

One type of power allocation policy is an *ω-only policy* that, every slot t, observes $\omega(t)$ and independently chooses $p(t) \in \{0, 1\}$ according to some stationary conditional probabilities $Pr[p(t) = 1|\omega(t) = \omega]$ that are specified for all $\omega \in \Omega$. The resulting average power and transmission rate is:

$$\begin{aligned} \mathbb{E}[p(t)] &= \textstyle\sum_{k=0}^{M} \pi(\omega_k) Pr[p(t) = 1|\omega(t) = \omega_k] \\ \mathbb{E}[\mu(t)] &= \textstyle\sum_{k=0}^{M} \pi(\omega_k)\omega_k Pr[p(t) = 1|\omega(t) = \omega_k] \end{aligned}$$

It is known that the problem (3)-(5) is solvable over the class of ω-only policies [9]. Specifically, if the arrival rate λ and the channel probabilities $\pi(\omega_k)$ were known in advance, one could offline compute an ω-only policy to satisfy:

$$\mathbb{E}[p(t)] = p^* \quad (7)$$
$$\mathbb{E}[\mu(t)] = \lambda \quad (8)$$

This is a 0-approximation for all $t \geq 0$. However, such an algorithm would typically incur infinite average queue size (since the service rate equals the arrival rate). Further, it is not possible to implement this algorithm without perfect knowledge of λ and $\pi(\omega_k)$ for all $\omega_k \in \Omega$.

Suppose one temporarily allows for infinite average queue size. Consider the following thought experiment (similar to that considered for utility optimal flow allocation in [16]). Consider an algorithm that does not know the system probabilities and hence makes a *single mistake* at time 0, so that:

$$\mathbb{E}[p(0)] = p^* + c$$

where $c > 0$ is some constant gap away from the optimal average power p^*. However, suppose a genie gives the controller perfect knowledge of the system probabilities at time 1, and then for slots $t \geq 1$ the network makes decisions to achieve the ideal averages (7)-(8). The resulting time average expected power over the first $t > 1$ slots is:

$$\overline{p}(t) = \frac{p^* + c}{t} + \frac{(t-1)p^*}{t} = p^* + \frac{c}{t}$$

Thus, to reach an ϵ-approximation, this genie-aided algorithm requires a convergence time $t = c/\epsilon = \Theta(1/\epsilon)$.

B. An example with $\Omega(1/\epsilon)$ convergence time

The above thought experiment does not prove an $\Omega(1/\epsilon)$ bound on convergence time because it assumes the algorithm makes decisions according to (7)-(8) for all slots $t \geq 1$, which may not be the optimal way to compensate for the mistake on slot 0. This section defines a simple system for which convergence time is at least $\Omega(1/\epsilon)$ under any algorithm.

Consider a system with deterministic arrivals of 1 packet every slot (so $\lambda = 1$). There are three possible channel states $\omega(t) \in \{1, 2, 3\}$, with probabilities:

$$\pi(3) = y, \pi(2) = z, \pi(1) = 1 - y - z$$

For each slot $t > 0$, define the *system history* $\mathcal{H}(t) = \{(a(0), \omega(0), p(0)), \ldots, (a(t-1), \omega(t-1), p(t-1))\}$. Define $\mathcal{H}(0) = 0$. For each slot t, a general algorithm observes the history $\mathcal{H}(t)$, the current $\omega(t)$, and makes a (possibly randomized) decision for $p(t) \in \{0, 1\}$. On a single slot, it is not difficult to show that the minimum average power $\mathbb{E}[p(t)]$ required to achieve a given average service rate $\mu = \mathbb{E}[\mu(t)]$ is characterized by the following function $h(\mu)$:

$$h(\mu) \triangleq \begin{cases} \mu/3 & \text{if } 0 \leq \mu \leq 3y \\ y + (\mu - 3y)/2 & \text{if } 3y \leq \mu \leq 3y + 2z \\ \mu - 2y - z & \text{if } 3y + 2z \leq \mu \leq 2y + z + 1 \end{cases}$$

There are two significant vertex points $(\mu, h(\mu))$ for this function. The first is $(3y, y)$, achieved by allocating power if and only if $\omega(t) = 3$. The second is $(3y + 2z, y + z)$, achieved by allocating power if and only if $\omega(t) \in \{2, 3\}$. For $\lambda \in [0, 2y + z + 1]$, it is known that $h(\lambda)$ is the minimum average power required to support an arrival rate of λ [9].

Define $\mathcal{R}$ as the set of points (μ, p) that lie on or above the curve $h(\mu)$:

$$\mathcal{R} = \{(\mu, p) \in \mathbb{R}^2 | 0 \leq \mu \leq 2y + z + 1, h(\mu) \leq p \leq 1\}$$

The set $\mathcal{R}$ is convex. Under any algorithm one has:

$$(\mathbb{E}[\mu(\tau)], \mathbb{E}[p(\tau)]) \in \mathcal{R} \quad \forall \tau \in \{0, 1, 2, \ldots\}$$

For a given $t > 1$, the following two vectors must be in $\mathcal{R}$:

$$(\mu_0, p_0) \triangleq (\mathbb{E}[\mu(0)], \mathbb{E}[p(0)])$$
$$(\mu_1, p_1) \triangleq \frac{1}{t-1} \sum_{\tau=1}^{t-1} (\mathbb{E}[\mu(\tau)], \mathbb{E}[p(\tau)])$$

That (μ_1, p_1) is in $\mathcal{R}$ follows because it is the average of points in $\mathcal{R}$, and $\mathcal{R}$ is convex. By definition of $(\overline{\mu}(t), \overline{p}(t))$:

$$(\overline{\mu}(t), \overline{p}(t)) = \frac{1}{t}(\mu_0, p_0) + \frac{1-t}{t}(\mu_1, p_1) \quad (9)$$

Fix ϵ such that $0 < \epsilon < 1/64$. An ϵ-approximation to the target point $(1, h(1))$ requires $\overline{\mu}(t)$ and $\overline{p}(t)$ to satisfy:

$$\overline{\mu}(t) \geq 1 - \epsilon \ , \ \overline{p}(t) \leq h(1) + \epsilon \quad (10)$$

Fix a particular algorithm and define $\theta_i(0) = Pr[p(t) = 1 | \omega(0) = i]$ for $i \in \{1, 2, 3\}$. The controller has no information about the probabilities y and z at time 0, and so the algorithm must specify the $\theta_i(0)$ values without knowledge of y and z. Suppose a genie reveals y and z on slot 1, and the network makes decisions on slots $\{1, \ldots, t-1\}$ that result in a (μ_1, p_1) vector that *optimally compensates for any mistake on slot* 0. Thus, (μ_1, p_1) is assumed to be the vector in $\mathcal{R}$ that ensures (9) produces an ϵ-approximation in the smallest time t.

The following proof considers the cases $\theta_2(0) \leq 1/2$ and $\theta_2(0) > 1/2$. In both cases, the nonlinear structure of the $h(\mu)$ curve prevents a fast recovery from the initial mistake.

- Case 1: Suppose $\theta_2(0) \leq 1/2$. Consider $y = 0, z = 1/4$. Then $\pi(1) = 3/4, \pi(2) = 1/4$, and $\omega(t) = 2$ is the most efficient state. The $h(\mu)$ curve is shown in Fig. 1. The minimum average power to support $\lambda = 1$ is $h(1) = 3/4$, and so the target point is $X = (1, 3/4)$. The point $(\mu_0, p_0) = (\mathbb{E}[\mu(0)], \mathbb{E}[p(0)])$ is:

$$(\mu_0, p_0) = \left(\frac{\theta_2(0)}{2} + \frac{3\theta_1(0)}{4}, \frac{\theta_2(0)}{4} + \frac{3\theta_1(0)}{4}\right)$$

The set of possible (μ_0, p_0) is formed by considering all $\theta_2(0) \in [0, 1/2]$, $\theta_1(0) \in [0, 1]$. This set lies inside the left (orange) shaded region of Fig. 1. To see this, note that if $\theta_2(0)$ is fixed at a certain value, the resulting (μ_0, p_0) point lies on a line segment of slope 1 that is formed by sweeping $\theta_1(0)$ through the interval $[0, 1]$. If $\theta_2(0) = 1/2$, that line segment is between points $(1/4, 1/8)$ and $(1, 7/8)$ in Fig. 1. If $\theta_2(0) < 1/2$ then the line segment is shifted to the left.

The small triangular (green) shaded region in Fig. 1, with one vertex at point A, is the *target region*. The

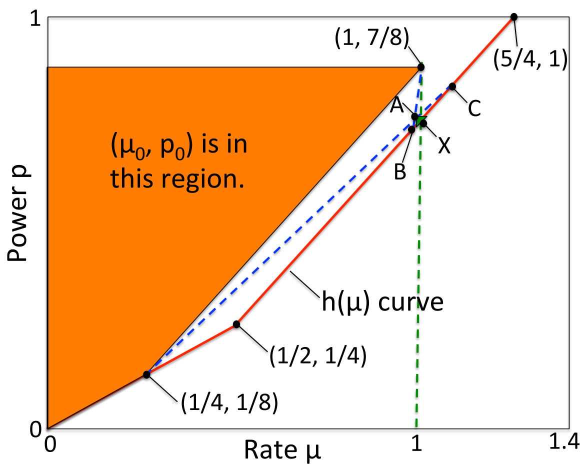

Fig. 1. The performance region for case 1. The line segments between $(1/4, 1/8)$ and C and between $(1, 7/8)$ and B intersect at point A.

vector $(\overline{\mu}(t), \overline{p}(t))$ must lie in this region to be an ϵ-approximation. The point A is defined:

$$A = X + (-\epsilon, \epsilon) = (1-\epsilon, 3/4+\epsilon)$$

It suffices to search for an optimal compensation vector (μ_1, p_1) on the curve $(\mu, h(\mu))$. This is because the average power from a point (μ_1, p_1) *above* the curve $(\mu, h(\mu))$ can be reduced, without affecting μ_1, by choosing a point *on* the curve. By geometry, (μ_1, p_1) must be on the line segment between points B and C in Fig. 1, where:

$$\begin{aligned} B &= X - \left(\frac{\epsilon}{1-16\epsilon}, \frac{\epsilon}{1-16\epsilon}\right) \\ C &= X + \left(\frac{11\epsilon}{1-16\epsilon}, \frac{11\epsilon}{1-16\epsilon}\right) \end{aligned}$$

Indeed, if (μ_1, p_1) is on the curve $(\mu, h(\mu))$ but not in between B and C, then a convex combination of (μ_0, p_0) and (μ_1, p_1) cannot lie in the target region (as required by (9)-(10)). Observe that:

$$\begin{aligned} ||(\overline{\mu}(t), \overline{p}(t)) - X|| &\leq \epsilon\sqrt{2} && (11) \\ ||(\mu_1, p_1) - X|| &\leq O(\epsilon) && (12) \\ ||(\mu_0, p_0) - (\mu_1, p_1)|| &\geq \sqrt{2}/16 && (13) \end{aligned}$$

where (11) follows by considering the maximum distance between X and any point in the (green) triangular region that defines an ϵ-approximation, (12) holds because any vector on the line segment between B and C is $O(\epsilon)$ distance away from X, and (13) holds because the distance between any point on the line segment between B and C and a point in the left (orange) shaded region is at least $\sqrt{2}/16$ (being the distance between the two parallel lines of slope 1). Starting from (11) one has:

$$\begin{aligned} \epsilon\sqrt{2} &\geq ||(\overline{\mu}(t), \overline{p}(t)) - X|| \\ &= ||(1/t)(\mu_0, p_0) + (1-1/t)(\mu_1, p_1) - X|| \\ &= ||(1/t)[(\mu_0, p_0) - (\mu_1, p_1)] - [X - (\mu_1, p_1)]|| \\ &\geq (1/t)||(\mu_0, p_0) - (\mu_1, p_1)|| - ||X - (\mu_1, p_1)|| \\ &\geq \sqrt{2}/(16t) - O(\epsilon) \end{aligned}$$

where the first equality holds by (9), the second-to-last inequality uses the triangle inequality $||W - Z|| \geq ||W|| - ||Z||$ for any vectors W, Z, and the final inequality uses (12) and (13). So $\sqrt{2}/(16t) \leq O(\epsilon)$. It follows that $t \geq \Omega(1/\epsilon)$.

- Case 2: Suppose $\theta_2(0) > 1/2$. However, suppose $y = z = 1/2$. Again it can be shown that $t \geq \Omega(1/\epsilon)$ (see [17]) for details).

IV. The Dynamic Algorithm

This section shows that a simple drift-plus-penalty algorithm achieves $O(\log(1/\epsilon)/\epsilon)$ convergence time and $O(\log(1/\epsilon))$ average queue size.

A. *Problem structure*

Without loss of generality, assume $\pi(\omega_k) > 0$ for all $k \in \{1, \ldots, M\}$ (else, remove ω_k from the set Ω). The value of $\pi(\omega_0)$ is possibly zero. For each $\mu \in [0, \mathbb{E}[\omega(t)]]$, define $h(\mu)$ as the minimum average power required to achieve an average transmission rate of μ. It is known that $p^* = h(\lambda)$. Further, it is not difficult to show that $h(\mu)$ is non-decreasing, convex, and *piecewise linear* with $h(0) = 0$ and $h(\mathbb{E}[\omega(t)]) = 1 - \pi(\omega_0)$. The point $(0, 0)$ is a vertex point of the piecewise linear curve $h(\mu)$. There are M other vertex points, achieved by the ω-only policies of the form:

$$p(t) = \begin{cases} 1 & \text{if } \omega(t) \geq \omega_k \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

for $k \in \{1, \ldots, M\}$. This means that a vertex point is achieved by only using channel states $\omega(t)$ that are on or above a certain threshold ω_k. Lowering the threshold value by selecting a smaller ω_k allows for a larger $\mathbb{E}[\mu(t)]$ at the expense of sometimes using less efficient channel states. The proof that this class of policies achieves the vertex points follows by a simple interchange argument that is omitted for brevity.

For ease of notation, define $\omega_{M+1} \triangleq \infty$ and $\mu_{M+1} \triangleq 0$. Let $\{\mu_1, \mu_2, \ldots, \mu_M, \mu_{M+1}\}$ be the set of transmission rates at which there are vertex points. Specifically, for $k \in \{1, \ldots, M\}$, μ_k corresponds to the threshold ω_k in the policy (14). That is:

$$\mu_k \triangleq \mathbb{E}[\omega(t) | \omega(t) \geq \omega_k] = \sum_{i=k}^{M} \omega_i \pi(\omega_i) \quad (15)$$

Note that:

$$0 = \mu_{M+1} < \mu_M < \mu_{M-1} < \cdots < \mu_1 = \mathbb{E}[\omega(t)]$$

It follows that $h(\mu_k)$ is the corresponding average power for vertex k, so that $(\mu_k, h(\mu_k))$ is a vertex point of the curve $h(\mu)$:

$$h(\mu_k) = Pr[\omega(t) \geq \omega_k] = \sum_{i=k}^{M} \pi(\omega_i) \quad (16)$$

The numbers $\{\mu_1, \mu_2, \ldots, \mu_M, \mu_{M+1}\}$ represent a set of measure 0 in the interval $[0, \mathbb{E}[\omega(t)]]$. It is assumed that the arrival rate λ is a number in $[0, \mathbb{E}[\omega(t)]]$ that lies strictly between two points μ_{b+1} and μ_b for some index $b \in \{1, \ldots, M\}$. That is:

$$\mu_{b+1} < \lambda < \mu_b$$

Thus, the point $(\lambda, h(\lambda))$ can be achieved by timesharing between the vertex points $(\mu_{b+1}, h(\mu_{b+1}))$ and $(\mu_b, h(\mu_b))$:

$$\begin{aligned} \lambda &= \theta\mu_{b+1} + (1-\theta)\mu_b && (17) \\ p^* = h(\lambda) &= \theta h(\mu_{b+1}) + (1-\theta)h(\mu_b) && (18) \end{aligned}$$

for some probability θ that satisfies $0 < \theta < 1$.

B. The drift-plus-penalty algorithm

For each slot $t \in \{0, 1, 2, \ldots\}$, define $L(t) = \frac{1}{2}Q(t)^2$ and $\Delta(t) = L(t+1) - L(t)$. Let V be a nonnegative real number. The *drift-plus-penalty* algorithm from [9][11] makes a power allocation decision that, every slot t, minimizes a bound on $\Delta(t) + Vp(t)$. The value V can be chosen as desired and affects a performance tradeoff. This technique is known to yield average queue size of $O(V)$ with deviation from optimal average power no more than $O(1/V)$ [9][11]. This holds for general multi-queue networks. By defining $\epsilon = 1/V$, this produces an $O(\epsilon)$ approximation with average queue size $O(1/\epsilon)$. Further, it can be shown that convergence time is $O(1/\epsilon^2)$ (see Appendix D in [18] and/or [19]).

In the context of the simple one-queue system of the current paper, the drift-plus-penalty algorithm reduces to the following: Every slot t, observe $Q(t)$ and $\omega(t)$ and choose $p(t) \in \{0, 1\}$ to minimize:

$$Vp(t) - Q(t)\omega(t)p(t)$$

That is, choose $p(t)$ according to the following rule:

$$p(t) = \begin{cases} 1 & \text{if } Q(t)\omega(t) \geq V \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

The current paper shows that, for this special case of a system with only one queue, the above algorithm leads to an improved queue size and convergence time tradeoff.

C. The induced Markov chain

The drift-plus-penalty algorithm induces a Markov structure on the system. The system state is $Q(t)$ and the state space is the set of nonnegative real numbers. Observe from (19) that the drift-plus-penalty algorithm has the following behavior:

- $Q(t) \in [V/\omega_{b+1}, V/\omega_b) \implies p(t) = 1$ if and only if $\omega(t) \geq \omega_{b+1}$. In this case one has (from (15) and (16)):

$$\begin{aligned} \mathbb{E}[\mu(t)|Q(t) \in [V/\omega_{b+1}, V/\omega_b)] &= \mu_{b+1} && (20) \\ \mathbb{E}[p(t)|Q(t) \in [V/\omega_{b+1}, V/\omega_b)] &= h(\mu_{b+1}) && (21) \end{aligned}$$

- $Q(t) \in [V/\omega_b, V/\omega_{b-1}) \implies p(t) = 1$ if and only if $\omega(t) \geq \omega_b$. In this case one has:

$$\begin{aligned} \mathbb{E}[\mu(t)|Q(t) \in [V/\omega_b, V/\omega_{b-1})] &= \mu_b && (22) \\ \mathbb{E}[p(t)|Q(t) \in [V/\omega_b, V/\omega_{b-1})] &= h(\mu_b) && (23) \end{aligned}$$

where $V/0$ is defined as ∞ (in the case $\omega_{b-1} = \omega_0 = 0$), and $\omega_{M+1} = \infty$ so that $V/\omega_{M+1} = 0$.

Now define intervals $\mathcal{I}^{(1)}, \mathcal{I}^{(2)}, \mathcal{I}^{(3)}, \mathcal{I}^{(4)}$ (see Fig. 2):

$$\begin{aligned} \mathcal{I}^{(1)} &\triangleq [0, V/\omega_{b+1}) \\ \mathcal{I}^{(2)} &\triangleq [V/\omega_{b+1}, V/\omega_b) \\ \mathcal{I}^{(3)} &\triangleq [V/\omega_b, V/\omega_{b-1}) \\ \mathcal{I}^{(4)} &\triangleq [V/\omega_{b-1}, \infty) \end{aligned}$$

If $V/\omega_{b+1} = 0$ then $\mathcal{I}^{(1)}$ is defined as the empty set, and if $V/\omega_{b-1} = \infty$ then $\mathcal{I}^{(4)}$ is defined as the empty set. The equalities (20)-(23) can be rewritten as:

$$\begin{aligned} \mathbb{E}\left[\mu(t)|Q(t) \in \mathcal{I}^{(2)}\right] &= \mu_{b+1} && (24) \\ \mathbb{E}\left[p(t)|Q(t) \in \mathcal{I}^{(2)}\right] &= h(\mu_{b+1}) && (25) \\ \mathbb{E}\left[\mu(t)|Q(t) \in \mathcal{I}^{(3)}\right] &= \mu_b && (26) \\ \mathbb{E}\left[p(t)|Q(t) \in \mathcal{I}^{(3)}\right] &= h(\mu_b) && (27) \end{aligned}$$

Recall that under the drift-plus-penalty algorithm (19), if $Q(t) \in \mathcal{I}^{(2)}$ then the set of all $\omega(t)$ that lead to a transmission is equal to $\{\omega \in \Omega | \omega \geq \omega_{b+1}\}$. If $Q(t) \in \mathcal{I}^{(1)}$, then the set of all $\omega(t)$ that lead to a transmission depends on the particular value of $Q(t)$. However, since interval $\mathcal{I}^{(1)}$ is to the left of interval $\mathcal{I}^{(2)}$, the set of all $\omega(t)$ that lead to a transmission when $Q(t) \in \mathcal{I}^{(1)}$ is always a subset of $\{\omega \in \Omega | \omega \geq \omega_{b+1}\}$. Similarly, since $\mathcal{I}^{(4)}$ is to the right of $\mathcal{I}^{(3)}$, the set of all $\omega(t)$ that lead to a transmission when $Q(t) \in \mathcal{I}^{(4)}$ is a superset of the set of all $\omega(t)$ that lead to a transmission when $Q(t) \in \mathcal{I}^{(3)}$. Therefore, under the drift-plus-penalty algorithm one has:

$$\begin{aligned} \mathbb{E}\left[\mu(t)|Q(t) \in \mathcal{I}^{(1)}\right] &\leq \mu_{b+1} && (28) \\ \mathbb{E}\left[p(t)|Q(t) \in \mathcal{I}^{(1)}\right] &\leq h(\mu_{b+1}) && (29) \\ \mathbb{E}\left[\mu(t)|Q(t) \in \mathcal{I}^{(4)}\right] &\geq \mu_b && (30) \\ \mathbb{E}\left[p(t)|Q(t) \in \mathcal{I}^{(4)}\right] &\geq h(\mu_b) && (31) \end{aligned}$$

For each $i \in \{1, 2, 3, 4\}$ define the indicator function:

$$1\{Q(t) \in \mathcal{I}^{(i)}\} = \begin{cases} 1 & \text{if } Q(t) \in \mathcal{I}^{(i)} \\ 0 & \text{otherwise} \end{cases}$$

For each slot $t > 0$ and each $i \in \{1, 2, 3, 4\}$, define $\overline{1}^{(i)}(t)$ as the expected fraction of time that $Q(t) \in \mathcal{I}^{(i)}$:

$$\overline{1}^{(i)}(t) \triangleq \frac{1}{t}\sum_{\tau=0}^{t-1} \mathbb{E}\left[1\{Q(t) \in \mathcal{I}^{(i)}\}\right]$$

It follows that (using (25), (27), (29)):

$$\begin{aligned} \overline{p}(t) \leq\ & \overline{1}^{(2)}(t)h(\mu_{b+1}) + \overline{1}^{(3)}(t)h(\mu_b) \\ & + \overline{1}^{(1)}(t)h(\mu_{b+1}) + \overline{1}^{(4)}(t) && (32) \end{aligned}$$

where the final term follows because $p(t) \leq 1$ for all slots t. Similarly (using (24), (26), (28)):

$$\begin{aligned} \overline{\mu}(t) \leq\ & \overline{1}^{(2)}(t)\mu_{b+1} + \overline{1}^{(3)}(t)\mu_b \\ & + \overline{1}^{(1)}(t)\mu_{b+1} + \overline{1}^{(4)}(t)\mathbb{E}[\omega(t)] && (33) \end{aligned}$$

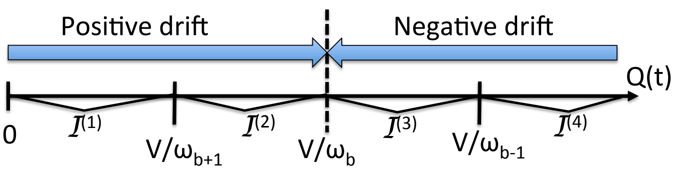

Fig. 2. An illustration of the four intervals $\mathcal{I}_i$ for $i \in \{1, 2, 3, 4\}$.

where the final term follows because $\mathbb{E}\left[\mu(t)|Q(t) \in \mathcal{I}^{(4)}\right] \leq \mathbb{E}\left[\omega(t)\right]$. Likewise (using (24), (26), (30)):

$$\overline{\mu}(t) \quad \geq \quad \overline{1}^{(2)}(t)\mu_{b+1} + \overline{1}^{(3)}(t)\mu_b + \overline{1}^{(4)}(t)\mu_b \tag{34}$$

which holds because $\mathbb{E}\left[\mu(t)|Q(t) \in \mathcal{I}^{(1)}\right] \geq 0$.

In the next section it is shown that:

- $\overline{\mu}(t)$ is close to λ when t is sufficiently large.
- $\overline{1}^{(2)}(t)$ and $\overline{1}^{(3)}(t)$ are close to θ and $1-\theta$, respectively, when t and V are sufficiently large.

V. Analysis

A. The distance between $\overline{\mu}(t)$ and λ

Recall that ω_M is the largest possible value of $\omega(t)$. Assume that $V \geq \omega_M^2$.

Lemma 1: If $V \geq \omega_M^2$, then under the drift-plus-penalty algorithm:

a) One has $p(t) = \mu(t) = 0$ whenever $Q(t) < \omega_M$.

b) The queueing equation (1) can be replaced by the following for all slots $t \in \{0, 1, 2, \ldots\}$:

$$Q(t+1) = Q(t) + a(t) - \mu(t)$$

Proof: Suppose $V \geq \omega_M^2$. To prove (a), suppose that $Q(t) < \omega_M$. Since $\omega(t) \leq \omega_M$ for all t, one has:

$$Q(t)\omega(t) \leq Q(t)\omega_M < \omega_M^2 \leq V$$

and so the algorithm (19) chooses $p(t) = 0$, so that $\mu(t)$ is also 0. This proves part (a). Part (a) implies that $Q(t) \geq \mu(t)$ for all slots t, which immediately implies part (b). □

Lemma 2: If $V \geq \omega_M^2$ and $Q(0) = q_0$ with probability 1 (for some constant $q_0 \geq 0$), then for every slot $t > 0$:

$$\overline{\mu}(t) = \lambda - \mathbb{E}\left[Q(t) - q_0\right]/t$$

Proof: By Lemma 1 one has for all slots $\tau \in \{0, 1, 2, \ldots\}$:

$$Q(\tau+1) - Q(\tau) = a(\tau) - \mu(\tau)$$

Taking expectations gives:

$$\mathbb{E}\left[Q(\tau+1)\right] - \mathbb{E}\left[Q(\tau)\right] = \lambda - \mathbb{E}\left[\mu(\tau)\right]$$

Summing the above over $\tau \in \{0, 1, 2, \ldots, t-1\}$ gives:

$$\mathbb{E}\left[Q(t)\right] - \mathbb{E}\left[Q(0)\right] = \lambda t - \sum_{\tau=0}^{t-1} \mathbb{E}\left[\mu(\tau)\right]$$

Dividing by t proves the result. □

B. The distance between $\overline{1}^{(2)}(t)$ and θ

The following lemma shows that if $\overline{1}^{(1)}(t)$, $\overline{1}^{(4)}(t)$, and $\mathbb{E}\left[Q(t) - q_0\right]/t$ are close to 0, then $\overline{1}^{(2)}(t)$ is close to θ.

Lemma 3: If $V \geq \omega_M^2$ and $Q(0) = q_0$ with probability 1 (for some constant $q_0 \geq 0$), then for all slots $t > 0$:

$$\begin{aligned} \theta - \frac{[\mu_b \overline{1}^{(1)}(t) - \psi(t)]}{\mu_b - \mu_{b+1}} \quad &\leq \quad \overline{1}^{(2)}(t) \\ &\leq \quad \theta + \frac{\overline{1}^{(4)}(t)\mathbb{E}\left[\omega(t)\right] + \psi(t)}{\mu_b - \mu_{b+1}} \end{aligned}$$

where $\psi(t) \triangleq \mathbb{E}\left[Q(t) - q_0\right]/t$.

Proof: Fix $t > 0$. Lemma 2 implies:

$$\begin{aligned} \lambda \quad &= \quad \overline{\mu}(t) + \psi(t) \\ &\geq \quad \overline{1}^{(2)}(t)\mu_{b+1} + \overline{1}^{(3)}(t)\mu_b + \overline{1}^{(4)}(t)\mu_b + \psi(t) \\ &= \quad \overline{1}^{(2)}(t)\mu_{b+1} + (1 - \overline{1}^{(2)}(t))\mu_b - \overline{1}^{(1)}(t)\mu_b + \psi(t) \end{aligned}$$

where the first inequality holds by (34). Substituting the identity for λ given in (17) into the above inequality gives:

$$\begin{aligned} &\theta\mu_{b+1} + (1-\theta)\mu_b \\ &\geq \overline{1}^{(2)}(t)\mu_{b+1} + (1 - \overline{1}^{(2)}(t))\mu_b - \overline{1}^{(1)}(t)\mu_b + \psi(t) \end{aligned}$$

Rearranging terms proves that:

$$\theta - \frac{[\mu_b \overline{1}^{(1)}(t) - \psi(t)]}{\mu_b - \mu_{b+1}} \leq \overline{1}^{(2)}(t)$$

The second inequality is proven similarly (see [17]). □

C. Positive and negative drift

Define $\mathbb{E}\left[Q(t+1) - Q(t)|Q(t)\right]$ as the *conditional drift*. Assume that $V \geq \omega_M^2$, so that Lemma 1 implies $Q(t+1) - Q(t) = a(t) - \mu(t)$ for all slots t. Thus:

$$\begin{aligned} \mathbb{E}\left[Q(t+1) - Q(t)|Q(t)\right] \quad &= \quad \mathbb{E}\left[a(t) - \mu(t)|Q(t)\right] \\ &= \quad \lambda - \mathbb{E}\left[\mu(t)|Q(t)\right] \end{aligned}$$

where the final equality follows because $a(t)$ is independent of $Q(t)$. From (24) and (28) one has for all slots t:

$$\mathbb{E}\left[\mu(t)|Q(t)\right] \leq \mu_{b+1} \quad \text{if } Q(t) < V/\omega_b$$

Likewise, from (26) and (30) one has:

$$\mathbb{E}\left[\mu(t)|Q(t)\right] \geq \mu_b \quad \text{if } Q(t) \geq V/\omega_b$$

Define positive constants β_L and β_R (associated with drift when $Q(t)$ is to the *Left* and *Right* of the threshold V/ω_b) by:

$$\beta_L \triangleq \lambda - \mu_{b+1} \quad , \quad \beta_R \triangleq \mu_b - \lambda$$

It follows that:

$$\mathbb{E}\left[Q(t+1) - Q(t)|Q(t)\right] \geq \beta_L \quad \text{if } Q(t) < V/\omega_b \tag{35}$$

$$\mathbb{E}\left[Q(t+1) - Q(t)|Q(t)\right] \leq -\beta_R \quad \text{if } Q(t) \geq V/\omega_b \tag{36}$$

In particular, the system has *positive drift* if $Q(t) < V/\omega_b$, and *negative drift* otherwise (see Fig. 2). Remarkably, the threshold drift structure naturally achieved by the drift-plus-penalty algorithm (which does not have knowledge of system probabilities) is qualitatively similar to the structure intentionally designed in [12] using full probability information.

D. A basic drift lemma

Consider a real-valued random process $Z(t)$ over slots $t \in \{0, 1, 2, \ldots\}$. The following drift lemma is similar in spirit to results in [20][15], but focuses on a finite time horizon with an arbitrary initial condition $Z(0) = z_0$ (rather than on steady state), and on expectations at a given time (rather than time averages). The lemma will be applied using $Z(t) = Q(t)$ for bounds on $\overline{Q}(t)$ and $\overline{1}^{(4)}(t)$. The lemma will be applied using $Z(t) = V/\omega_b - Q(t)$ to bound $\overline{1}^{(1)}(t)$. Assume there is a constant $\delta_{max} > 0$ such that with probability 1:

$$|Z(t+1) - Z(t)| \leq \delta_{max} \ \forall t \in \{0, 1, 2, \ldots\} \tag{37}$$

Suppose there are constants $\theta \in \mathbb{R}$ and $\beta > 0$ such that:

$$\mathbb{E}\left[Z(t+1) - Z(t)|Z(t)\right] \leq \begin{cases} \delta_{max} & \text{if } Z(t) < \theta \\ -\beta & \text{if } Z(t) \geq \theta \end{cases} \tag{38}$$

Note that if (37) holds then (38) automatically holds for the special case $Z(t) < \theta$. Thus, the negative drift case $Z(t) \geq \theta$ is the important case for condition (38). Further, if (37)-(38) both hold, then the constant β necessarily satisfies:

$$0 < \beta \leq \delta_{max}$$

Lemma 4: Suppose $Z(t)$ is a random process that satisfies (37)-(38) for given constants θ, δ_{max}, β (with $\theta \in \mathbb{R}$ and $0 < \beta \leq \delta_{max}$). Suppose $Z(0) = z_0$ (with probability 1) for some $z_0 \in \mathbb{R}$. Then for every slot $t \geq 0$ the following holds:

$$\mathbb{E}\left[e^{rZ(t)}\right] \leq D + \left(e^{rz_0} - D\right)\rho^t \tag{39}$$

where constants r, ρ, D are defined:

$$\begin{aligned} r &\triangleq \frac{\beta}{\delta_{max}^2 + \delta_{max}\beta/3} && (40) \\ \rho &\triangleq 1 - r\beta/2 && (41) \\ D &\triangleq \frac{(e^{r\delta_{max}} - \rho)e^{r\theta}}{1-\rho} && (42) \end{aligned}$$

Note that the property $0 < \beta \leq \delta_{max}$ can be used to show that $0 < \rho < 1$.

Proof: (Lemma 4) The proof is by induction. The inequality (39) trivially holds for $t = 0$. Suppose (39) holds at some slot $t \geq 0$. The goal is to show that it also holds on slot $t+1$. Let r be a positive number that satisfies $0 < r\delta_{max} < 3$. It is known from results in [20] that for any real number x that satisfies $|x| \leq \delta_{max}$:

$$e^{rx} \leq 1 + rx + \frac{(r\delta_{max})^2}{2(1 - r\delta_{max}/3)} \tag{43}$$

Define $\delta(t) = Z(t+1) - Z(t)$ and note that $|\delta(t)| \leq \delta_{max}$ for all t. Then:

$$\begin{aligned} e^{rZ(t+1)} &= e^{rZ(t)}e^{r\delta(t)} \\ &\leq e^{rZ(t)}\left[1 + r\delta(t) + \frac{(r\delta_{max})^2}{2(1 - r\delta_{max}/3)}\right] && (44) \end{aligned}$$

where the final inequality holds by (43). Choose r such that:

$$\frac{(r\delta_{max})^2}{2(1 - r\delta_{max}/3)} \leq \frac{r\beta}{2} \tag{45}$$

It is not difficult to show that the value of r given in (40) simultaneously satisfies (45) and $0 < r\delta_{max} < 3$. For this value of r, substituting (45) into (44) gives:

$$e^{rZ(t+1)} \leq e^{rZ(t)}\left[1 + r\delta(t) + r\beta/2\right] \tag{46}$$

Now consider the following two cases:

- Case 1: Suppose $Z(t) \geq \theta$. Taking conditional expectations of (46) gives:

$$\begin{aligned} \mathbb{E}\left[e^{rZ(t+1)}|Z(t)\right] &\leq \mathbb{E}\left[e^{rZ(t)}[1 + r\delta(t) + r\beta/2]|Z(t)\right] \\ &\leq e^{rZ(t)}[1 - r\beta + r\beta/2] && (47) \\ &= e^{rZ(t)}\rho \end{aligned}$$

  where (47) follows by (38), and the final equality holds by definition of ρ in (41).
- Case 2: Suppose $Z(t) < \theta$. Then:

$$\begin{aligned} \mathbb{E}\left[e^{rZ(t+1)}|Z(t)\right] &= \mathbb{E}\left[e^{rZ(t)}e^{r\delta(t)}|Z(t)\right] \\ &\leq e^{rZ(t)}e^{r\delta_{max}} \end{aligned}$$

Putting these two cases together gives:

$$\begin{aligned} &\mathbb{E}\left[e^{rZ(t+1)}\right] \\ &\leq \rho\mathbb{E}\left[e^{rZ(t)}|Z(t) \geq \theta\right] Pr[Z(t) \geq \theta] \\ &\quad + e^{r\delta_{max}}\mathbb{E}\left[e^{rZ(t)}|Z(t) < \theta\right] Pr[Z(t) < \theta] \\ &= \rho\mathbb{E}\left[e^{rZ(t)}\right] \\ &\quad + (e^{r\delta_{max}} - \rho)\mathbb{E}\left[e^{rZ(t)}|Z(t) < \theta\right] Pr[Z(t) < \theta] \\ &\leq \rho\mathbb{E}\left[e^{rZ(t)}\right] + (e^{r\delta_{max}} - \rho)e^{r\theta} \end{aligned}$$

where the final inequality uses the fact that $e^{r\delta_{max}} > 1 > \rho$. By the induction assumption it is known that (39) holds on slot t. Substituting (39) into the right-hand-side of the above inequality gives:

$$\begin{aligned} \mathbb{E}\left[e^{rZ(t+1)}\right] &\leq \rho\left[D + \left(e^{rz_0} - D\right)\rho^t\right] \\ &\quad + (e^{r\delta_{max}} - \rho)e^{r\theta} \\ &= D + \left(e^{rz_0} - D\right)\rho^{t+1} \end{aligned}$$

where the final equality holds by the definition of D in (42). This completes the induction step. □

Let $1\{Z(\tau) \geq \theta + c\}$ be an indicator function that is 1 if $Z(\tau) \geq \theta + c$, and 0 else.

Corollary 1: If the assumptions of Lemma 4 hold, then for any $c > 0$ and any slots T and t that satisfy $0 \leq T < t$:

$$\begin{aligned} &\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left[1\{Z(\tau) \geq \theta + c\}\right] \\ &\leq \frac{(e^{r\delta_{max}} - \rho)e^{-rc}}{1-\rho} + \left[\frac{T}{t} + \frac{e^{r(z_0 - c - \theta)}\rho^T}{t(1-\rho)}\right] && (48) \end{aligned}$$

where r and ρ are defined in (40)-(41). Further, if $z_0 \leq \theta$ then for any $t > 0$:

$$\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left[1\{Z(\tau) \geq \theta + c\}\right] \leq \frac{e^{-rc}(e^{r\delta_{max}} - \rho + 1/t)}{(1-\rho)} \tag{49}$$

Proof: Inequality (49) follows from (48) by using $T = 0$. See [17] for a proof of (48). $\square$

The intuition behind the right-hand-side of (48) is that the first term represents a "steady state" bound as $t \to \infty$, which decays like e^{-rc}. The last two terms (in brackets) are due to the transient effect of the initial condition z_0. This transient can be significant when $z_0 > \theta$. In that case, $e^{r(z_0-c-\theta)}$ might be large, and a time T is required to shrink this term by multiplication with the factor ρ^T.

E. Bounding $\mathbb{E}[Q(t)]$ and $\overline{1}^{(4)}(t)$

Let $Q(t)$ be the backlog process under the drift-plus-penalty algorithm. Assume that $V \geq \omega_M^2$ and the initial condition is $Q(0) = q_0$ for some constant q_0. Define $\delta_{max} \triangleq \max[\omega_M, a_{max}]$ as the largest possible change in $Q(t)$ over one slot, so that:

$$|Q(t+1) - Q(t)| \leq \delta_{max} \;\; \forall t \in \{0, 1, 2, \ldots\}$$

From (36) it holds that:

$$\mathbb{E}[Q(t+1) - Q(t)|Q(t)] \leq \begin{cases} \delta_{max} & \text{if } Q(t) < V/\omega_b \\ -\beta_R & \text{if } Q(t) \geq V/\omega_b \end{cases}$$

It follows that the process $Q(t)$ satisfies the conditions (37)-(38) required for Lemma 4. Specifically, define $Z(t) = Q(t)$, $z_0 = q_0$, $\theta = V/\omega_b$, $\beta = \beta_R$.

Lemma 5: If $0 \leq q_0 \leq V/\omega_b$ and $V \geq \omega_M^2$, then for all slots $t \geq 0$ one has:

$$\mathbb{E}[Q(t)] \leq \frac{V}{\omega_b} + \frac{1}{r_R}\log\left(1 + \frac{e^{r_R\delta_{max}} - \rho_R}{1-\rho_R}\right) = O(V)$$

where constants r_R and ρ_R are defined:

$$r_R \triangleq \frac{\beta_R}{\delta_{max}^2 + \delta_{max}\beta_R/3} \tag{50}$$

$$\rho_R \triangleq 1 - r_R\beta_R/2 \tag{51}$$

Proof: For ease of notation, let "r" and "ρ" respectively denote "r_R" and "ρ_R" given in (50) and (51). Define $\theta = V/\omega_b$ and $\beta = \beta_R$. By (39) one has for all $t \geq 0$ (using $Z(0) = Q(0) = q_0$):

$$\begin{aligned} \mathbb{E}\left[e^{rQ(t)}\right] &\leq D + (e^{rq_0} - D)\rho^t \\ &\leq D + e^{rV/\omega_b} \end{aligned}$$

where D is given in (42), and where the final inequality uses $D\rho^t \geq 0$ and $q_0 \leq V/\omega_b$. Using Jensen's inequality gives:

$$e^{r\mathbb{E}[Q(t)]} \leq D + e^{rV/\omega_b}$$

Taking a log of both sides and dividing by r leads to the result. $\square$

Lemma 6: If $0 \leq q_0 \leq V/\omega_b$ and $V \geq \omega_M^2$, then for all slots $t > 0$:

$$\overline{1}^{(4)}(t) \leq O(e^{-r_R V(\frac{1}{\omega_{b-1}} - \frac{1}{\omega_b})})$$

where r_R is given by (50).

Proof: For ease of notation, this proof uses "r" to denote "r_R." If the interval $\mathcal{I}^{(4)}$ does not exist then $\overline{1}^{(4)}(t) = 0$ and the result is trivial. Now suppose interval $\mathcal{I}^{(4)}$ exists (so that the interval $\mathcal{I}^{(3)}$ is not the final interval in Fig. 2). Define $\theta = V/\omega_b$, $c = V(1/\omega_{b-1} - 1/\omega_b)$, $\beta = \beta_R$, $\rho = 1 - r\beta_R/2$. Then $1\{Q(\tau) \geq \theta + c\} = 1$ if and only if $Q(\tau) \geq V/\omega_{b-1}$, which holds if and only if $Q(\tau) \in \mathcal{I}_4$. Thus, for all slots $t > 0$:

$$\begin{aligned} \overline{1}^{(4)}(t) &= \frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}[1\{Q(\tau) \geq \theta + c\}] \\ &\leq \frac{e^{-rc}(e^{r\delta_{max}} - \rho + 1/t)}{1-\rho} \\ &= \frac{e^{-rV(\frac{1}{\omega_{b-1}} - \frac{1}{\omega_b})}(e^{r\delta_{max}} - \rho + 1/t)}{1-\rho} \end{aligned} \tag{52}$$

where (52) holds by (49) (which applies since $z_0 = q_0 \leq \theta$). The right-hand-side of the above inequality is indeed of the form $O(e^{-rV(\frac{1}{\omega_{b-1}} - \frac{1}{\omega_b})})$. $\square$

F. Bounding $\overline{1}^{(1)}(t)$

One can similarly prove a bound on $\overline{1}^{(1)}(t)$. The intuition is that the positive drift in region $\mathcal{I}^{(2)}$ of Fig. 2, together with the fact that the size of interval $\mathcal{I}^{(2)}$ is $\Theta(V)$, makes the fraction of time the queue is to the left of V/ω_b decay exponentially as we move further left. The result is given below. Recall that $Q(0) = q_0$ for some constant $q_0 \geq 0$.

Lemma 7: If $q_0 \geq 0$ and $V \geq \omega_M^2$, then for all slots $t > 0$ one has:

$$\overline{1}^{(1)}(t) \leq O(V)/t + O(e^{-r_L V(\frac{1}{\omega_b} - \frac{1}{\omega_{b+1}})})$$

where r_L is defined:

$$r_L \triangleq \frac{\beta_L}{\delta_{max}^2 + \delta_{max}\beta_L/3}$$

Intuitively, the first term in the above lemma (that is, the $O(V)/t$ term) bounds the contribution from the *transient time* starting from the initial state $Q(0) = q_0$ and ending when the threshold V/ω_b is crossed. The second term represents a "steady state" probability assuming an initial condition V/ω_b. The proof defines a new process $Z(t) = V/\omega_b - Q(t)$. It then applies inequality (48) of Corollary 1, with a suitably large time $T > 0$, to handle the initial condition $z_0 = V/\omega_b - q_0$.

Proof: (Lemma 7) Omitted for brevity (see [17]). $\square$

G. Optimal backlog and near-optimal convergence time

Define:

$$\gamma \triangleq \min\left[r_R\left(\frac{1}{\omega_{b-1}} - \frac{1}{\omega_b}\right), r_L\left(\frac{1}{\omega_b} - \frac{1}{\omega_{b+1}}\right)\right]$$

Results of Lemmas 5-7 imply that if the drift-plus-penalty algorithm (19) is used with $V \geq \omega_M^2$, and if the initial queue state satisfies $0 \leq q_0 \leq V/\omega_b$, then for all $t > 0$:

$$\overline{Q}(t) \leq O(V) \tag{53}$$

$$\mathbb{E}[Q(t)]/t \leq O(V)/t \tag{54}$$

$$\overline{1}^{(4)}(t) \leq O(e^{-\gamma V}) \tag{55}$$

$$\overline{1}^{(1)}(t) \leq O(e^{-\gamma V}) + O(V)/t \tag{56}$$

Indeed, (53)-(54) follow from Lemma 5, while (55) and (56) follow from Lemmas 6 and 7, respectively.

Fix $\epsilon > 0$ and define:

$$\begin{aligned} V &= \max[(1/\gamma)\log(1/\epsilon), \omega_M^2] \\ T_\epsilon &= \log(1/\epsilon)/\epsilon \end{aligned}$$

Then $V = O(\log(1/\epsilon))$ and $e^{-\gamma V} = O(\epsilon)$. Inequalities (53)-(56) immediately imply the following facts:

- Fact 1: For all slots $t > 0$ one has $\overline{Q}(t) \leq O(\log(1/\epsilon))$.
- Fact 2: For all slots $t > T_\epsilon$ one has $\mathbb{E}[Q(t)]/t \leq O(\epsilon)$.
- Fact 3: For all slots $t > 0$ one has $\overline{1}^{(4)}(t) \leq O(\epsilon)$.
- Fact 4: For all slots $t > T_\epsilon$ one has $\overline{1}^{(1)}(t) \leq O(\epsilon)$.

For example, Fact 2 follows from (54) since:

$$\mathbb{E}[Q(t)]/t \leq O(V)/t \leq O(V)/T_\epsilon = O(\epsilon)$$

Fact 2 and Lemma 2 ensure that for $t > T_\epsilon$:

$$\overline{\mu}(t) \geq \lambda - O(\epsilon) \tag{57}$$

Facts 2, 3, 4 and Lemma 3 ensure that for $t > T_\epsilon$:

$$|\overline{1}^{(2)}(t) - \theta| \leq O(\epsilon) \quad , \quad |\overline{1}^{(3)}(t) - (1-\theta)| \leq O(\epsilon)$$

Substituting the above into (32) proves that for $t > T_\epsilon$:

$$\begin{aligned} \overline{p}(t) &\leq \theta h(\mu_{b+1}) + (1-\theta)h(\mu_b) + O(\epsilon) \\ &= p^* + O(\epsilon) \end{aligned} \tag{58}$$

where the final equality holds by (18). The guarantees (57) and (58) show that the drift-plus-penalty algorithm gives an $O(\epsilon)$-approximation with convergence time $T_\epsilon = O(\log(1/\epsilon)/\epsilon)$. This is within a factor $\log(1/\epsilon)$ of the convergence time lower bound given in Section III. Hence, the algorithm has near-optimal convergence time.

Further, it is known that if the rate-power curve $h(\mu)$ has at least two piecewise linear segments and if the point $(\lambda, h(\lambda))$ does not lie on the segment closest to the origin, then any algorithm that yields an $O(\epsilon)$-approximation must have average queue size that satisfies $\overline{Q}(t) \geq \Omega(\log(1/\epsilon))$ [13]. Fact 1 shows the drift-plus-penalty algorithm meets this bound with equality. Hence, it provides an *optimal* average queue size tradeoff, and *near optimal* convergence time.

VI. Conclusions

This paper considers convergence time for minimizing average power in a wireless transmission link with time varying channels and random traffic. Prior algorithms produce an ϵ-approximation with convergence time $O(1/\epsilon^2)$. This paper shows, for a simple example, that no algorithm can get convergence time better than $O(1/\epsilon)$. It then shows that this ideal convergence time tradeoff can be approached to within a logarithmic factor. Furthermore, the resulting average queue size is at most $O(\log(1/\epsilon))$, which is known to be an optimal tradeoff. This establishes fundamental convergence time, queue size, and power characteristics of wireless links. It shows that learning times in an unknown environment can be pushed much faster than expected.

Acknowledgement

Work supported in part by NSF Career grant CCF-0747525.

References

[1] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466-478, March 1993.

[2] A. Eryilmaz and R. Srikant. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. *IEEE/ACM Transactions on Networking*, vol. 15, no. 6, pp. 1333-1344, Dec. 2007.

[3] A. Eryilmaz and R. Srikant. Joint congestion control, routing, and MAC for stability and fairness in wireless networks. *IEEE Journal on Selected Areas in Communications, Special Issue on Nonlinear Optimization of Communication Systems*, vol. 14, pp. 1514-1524, Aug. 2006.

[4] J. W. Lee, R. R. Mazumdar, and N. B. Shroff. Opportunistic power scheduling for dynamic multiserver wireless systems. *IEEE Transactions on Wireless Communications*, vol. 5, no.6, pp. 1506-1515, June 2006.

[5] H. Kushner and P. Whiting. Asymptotic properties of proportional-fair sharing algorithms. *Proc. 40th Annual Allerton Conf. on Communication, Control, and Computing, Monticello, IL*, Oct. 2002.

[6] R. Agrawal and V. Subramanian. Optimality of certain channel aware scheduling policies. *Proc. 40th Annual Allerton Conf. on Communication, Control, and Computing, Monticello, IL*, Oct. 2002.

[7] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, vol. 50, no. 4, pp. 401-457, 2005.

[8] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, vol. 16, no. 2, pp. 396-409, April 2008.

[9] M. J. Neely. *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.

[10] X. Liu, E. K. P. Chong, and N. B. Shroff. A framework for opportunistic scheduling in wireless networks. *Computer Networks*, vol. 41, no. 4, pp. 451-474, March 2003.

[11] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 2915-2934, July 2006.

[12] R. Berry and R. Gallager. Communication over fading channels with delay constraints. *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1135-1149, May 2002.

[13] M. J. Neely. Optimal energy and delay tradeoffs for multi-user wireless downlinks. *IEEE Transactions on Information Theory*, vol. 53, no. 9, pp. 3095-3113, Sept. 2007.

[14] L. Huang, S. Moeller, M. J. Neely, and B. Krishnamachari. LIFO-backpressure achieves near optimal utility-delay tradeoff. *Proc. WiOpt*, May 2011.

[15] L. Huang and M. J. Neely. Delay reduction via Lagrange multipliers in stochastic network optimization. *IEEE Transactions on Automatic Control*, vol. 56, no. 4, pp. 842-857, April 2011.

[16] B. Li and A. Eryilmaz. Wireless scheduling for network utility maximization with optimal convergence speed. In *Proc. IEEE INFOCOM*, Turin, Italy, April 2013.

[17] M. J. Neely. Energy-aware wireless scheduling with near optimal backlog and convergence time tradeoffs. *ArXiv technical report, arXiv:1411.4740*, Nov. 2014.

[18] M. J. Neely. Distributed stochastic optimization via correlated scheduling. *ArXiv technical report, arXiv:1304.7727v2*, May 2013.

[19] M. J. Neely. A simple convergence time analysis of drift-plus-penalty for stochastic optimization and convex programs. *ArXiv technical report, arXiv:1412.0791v1*, Dec. 2014.

[20] F. Chung and L. Lu. Concentration inequalities and martingale inequalities–a survey. *Internet Mathematics*, vol. 3, pp. 79-127, 2006.