MaxWeight Scheduling: "Smoothness" of the Service Process

Rahul Singh* and Alexander Stolyar†
*Department of Electrical and Computer Engineering, Texas A&M University. Email: rsingh1@tamu.edu
†Department of Industrial and Systems Engineering, Lehigh University. Email: stolyar@lehigh.edu

***Abstract*—The model is a "generalized switch", serving multiple traffic flows in discrete time. The switch uses MaxWeight algorithm to make a service decision (scheduling choice) at each time step, depending on the current queue lengths. In some applications, it is *not* important to keep the queue lengths/delays small (e.g., when queues are virtual, rather than physical), but is important that the service processes provided to each flow remains "smooth" (i.e., without large gaps in service) even when the switch is heavily loaded. Addressing this question reduces to the analysis of the asymptotic behavior of the unscaled queue-differential process in heavy traffic. We prove that the stationary regime of this process converges to that of a positive recurrent Markov chain, whose structure we explicitly describe. This in turn implies "smoothness" of the service processes.**

I. Introduction

Suppose we have a system in which several data traffic flows share a common transmission medium (or channel). Sharing means that in each time slot a scheduler chooses a transmission mode – the subset the flows to serve and corresponding transmission rates; the outcome of each transmission (the number of successfully delivered packets) is random. Scheduler has two key objectives: (a) the time-average (successful) transmission rate of each flow i has to be at least some $\lambda_i > 0$; (b) the successful transmissions for each flow need to be spread out "smoothly" in time – without large time-gaps between succesful transmissions. Such models arise, for example, when the goal is *timely delivery* of information over a shared wireless channel [6].

A very natural way to approach this problem is to treat the model as a queueing system, where services (transmissions) are controlled by a so called MaxWeight scheduler (see [4], [10], [11]), which serves a set of *virtual queues* (one for each traffic flow), each receiving new work at the rate λ_i. (See e.g. [1].) This automatically achieves objective (a), if this is feasible at all; MaxWeight is known to be *throughput optimal* – stabilize the queues if this is feasible at all. The MaxWeight stability results, however, do not tell whether or not the objective (b) is achieved. Specifically, when the system is heavily loaded, i.e. the vector $\boldsymbol{\lambda} = (\lambda_i)$ is within the system *rate region* $\boldsymbol{V}$, but close to its boundary, the steady-state queue lengths under MaxWeight are necessarily large, and it is conceivable that this may result in large time-gaps in service for individual flows. (Note that, if (a) and (b) are the objectives and the queues are virtual, the large queue lengths in themselves are not an issue. As long as (a) and (b) are achieved, minimizing the queue lengths is not important.) Our main results show that this is *not* the case. Namely, in the heavy traffic regime, when $\boldsymbol{\lambda} \to \boldsymbol{\lambda}^\star$, where $\boldsymbol{\lambda}^\star$ is a point on the outer boundary of rate region $\boldsymbol{V}$, the service process remains "smooth", in the sense that its stationary regime converges to that of a positive recurrent Markov chain, whose structure is given explicitly.

To obtain "clean" convergence results, we assume that the amount of new work arriving in the queues in each time slot is random and has *continuous* distribution. (The amounts of service are random, but discrete.) Under this assumption, the state spaces of the processes that we consider are continuous. On one hand, this makes the analysis more involved (because the notion of positive recurrence is more involved for a continuous state space, as opposed to a countable one). But on the other hand, this makes all stationary distributions absolutely continuous w.r.t. the corresponding Lebesgue measure, making it easier to prove convergence. We emphasize that the assumption of continuous distribution of the arriving work is non-restrictive; if we create virtual queues, artificially, for the purpose of applying MaxWeight algorithm, the structure of the virtual arrival process is within our control.

The problem essentially reduces to analysis of stationary versions of the *queue-differential* process $\boldsymbol{Y}$, which is the projection of the (weighted) queue length process on the subspace $\boldsymbol{\nu}_\perp$, orthogonal to the outer *normal cone* $\boldsymbol{\nu}$ to the rate region $\boldsymbol{V}$ at the point $\boldsymbol{\lambda}^\star$. As we show, in the heavy-traffic limit, in steady-state, the values of the queue-differential process $\boldsymbol{Y}$ uniquely determine the decisions chosen by MaxWeight scheduler. Note that the process $\boldsymbol{Y}$ is obtained by projection only, without any scaling depending on the system load.

The model that we consider is essentially a "generalized switch" of [10]. Some features of our model, namely random service outcome and continuous amounts of arriving work, as well as the objective (b), are motivated by applications such as timely delivery of packets of multiple flows over a shared wireless channel [6]. The model of [6] is a special case of ours; paper [6] introduces a *debt scheme* and proves that it achieves the throughput objective (a); the objective (b) is not considered in [6].

The analysis of MaxWeight stability has a long history, starting from the seminal paper [11], which introduced MaxWeight; heavy traffic analysis of the algorithm originated in [10]. (See, e.g., [4] for an extensive recent review of MaxWeight literature.)

The line of work most closely related to this paper, is that in [4], [7], [8]. Paper [4] studies MaxWeight under heavy traffic regime and under the additional assumption that the normal cone $\boldsymbol{\nu}$ is one-dimensional, i.e. it is a ray. (The latter assumption is usually referred to in the literature as *complete resource pooling* (CRP).) Paper [4] shows, in particular, the stationary distribution tightness of what we call the queue-differential process $\boldsymbol{Y}$ in heavy traffic. Part of our analysis is also showing the stationary distribution tightness of $\boldsymbol{Y}$ – it is analogous to that in [4] (and we also borrow a lot of notation from [4]). Besides the difference in models, our proof of tightness is more general in that it applies to non-CRP case – this more general argument is close to that used in [2]. From the tightness of stationary distributions, using the structure of the corresponding continuous state space, we obtain the convergence of the stationary version of (non-Markov) process $\boldsymbol{Y}$ to that of a positive recurrent Markov chain, whose structure we explicitly describe.

Papers [7], [8] consider objective (b) in the heavy traffic regime. They introduce a modification of MaxWeight, called *regular service guarantee* (RSG) scheme, which explicitly tracks the service time-gaps for each flow to dynamically increase the scheduling priority of flows with large current time-gaps. The papers prove that RSG, under certain parameter settings, preserves heavy-traffic queue-length minimization properties of MaxWeight, under CRP condition; at the same time, the papers demonstrate via simulations that RSG improves smoothness (regularity) of the service process. Recall that in this paper we focus on the "pure" MaxWeight, without CRP, and formally show the service process smoothness in the heavy traffic limit.

The rest of the paper is organized as follows. The formal model is presented in Section II. Section III describes the MaxWeight algorithm and the heavy traffic asymptotic regime. Our main results, Theorems 2 and 21 are described in Section IV. (Formal statement of Theorem 21 is in Section IX.) An important contribution of this work is that our main results do *not* require the CRP condition. However, due to space constraints, here we give a detailed proof for the CRP case only – the generalization to the non-CRP case can be found in [12]. The CRP condition is defined in Section V. In Section VI we provide some necessary background and results for general state-space Markov chains. Sections VII – IX we prove our results for the special case when CRP holds. Finally, in Section X we present simulation results for a heavily loaded system, which show smoothness of the service process under MaxWeight; to demonstrate that the smoothness property is not guaranteed under any throughput optimal scheduling, we also simulate an algorithm which can be viewed as a crude version of 802.11 Carrier-Sensing-Multiple-Access (CSMA) MAC algorithm – the service process under this algorithm has large gaps.

A. Basic notation

Elements of a Euclidean space $\mathbb{R}^N$ will be viewed as row-vectors, and written in bold font; $\|\boldsymbol{a}\|$ is the usual Euclidean norm of vector $\boldsymbol{a}$. For two vectors $\boldsymbol{a}$ and $\boldsymbol{b}$, $\boldsymbol{a} \cdot \boldsymbol{b}$ denotes their scalar (dot) product; vector inequalities are understood componentwise; zero vector and the vector of all ones are denoted $\boldsymbol{0}$ and $\boldsymbol{1}$, respectively; $\boldsymbol{ab}$ will denote the vector obtained by componentwise multiplication; if all components of $\boldsymbol{b}$ are non-zero, $\frac{\boldsymbol{a}}{\boldsymbol{b}}$ will denote the vector obtained by componentwise division; statement "$\boldsymbol{a}$ is a positive vector" means $\boldsymbol{a} > \boldsymbol{0}$. The closed ball of radius r centered at $\boldsymbol{x}$ is $B_r(\boldsymbol{x})$. The positive orthant of $\mathbb{R}^N$ is denoted $\mathbb{R}^N_+ = \{\boldsymbol{x} \in \mathbb{R}^N : \boldsymbol{x} \geq \boldsymbol{0}\}$. For numbers a and b, we denote $a \vee b = \max(a,b)$, $a \wedge b = \min(a,b)$, $a^+ = a \vee 0$. For vectors $\boldsymbol{a} \leq \boldsymbol{b}$, we denote by $[\boldsymbol{a}, \boldsymbol{b}]$ the rectangle $\times_{i=1}^N [a_i, b_i]$ in $\mathbb{R}^N$. We always consider Borel σ-algebra $\mathcal{B}(\mathbb{R}^N)$ (resp. $\mathcal{B}(\mathbb{R}^N_+)$) on $\mathbb{R}^N$ (resp. $\mathcal{B}(\mathbb{R}^N_+)$), when the latter is viewed as measurable space. Lebesgue measure on $\mathbb{R}^N$ is denoted by $\mathcal{L}$. When we consider a linear subspace of $\mathbb{R}^N$, we endow it with the Euclidean metric and the corresponding Borel σ-algebra and Lebesgue measure. For a random process $\boldsymbol{W}(t),\ t = 0, 1, 2, \ldots$, we often use notation $\boldsymbol{W}(\cdot)$ or simply $\boldsymbol{W}$.

II. System Model

We consider a system of N flows served by a "switch", which evolves in discrete time $t = 0, 1, \ldots$. At the beginning of each time-slot, the scheduler has to choose from a finite number K of "service-decisions". If the service decision $k \in \{1, \ldots, K\}$ is chosen, then independently of the past history the flows get an amount of service, given by a random non-negative vector. Furthermore, we assume that (if decision k is chosen), there is a finite number $\mathcal{O}_k$ of possible service-vector outcomes, i.e. with probability $p^{k,j}, j = 1, \ldots, \mathcal{O}_k$, it is given by a non-negative vector $\boldsymbol{v}^{k,j} = (v_1^{k,j}, \ldots, v_N^{k,j})$. The expected service vector for decision k is denoted $\boldsymbol{\mu}^k = (\mu_1^k, \ldots, \mu_N^k) = \sum_{j=1}^{\mathcal{O}_k} \boldsymbol{v}^{k,j} p^{k,j}$. We assume that vectors $\boldsymbol{\mu}^k$ are non-zero and different from each other; and that for each i there exists k such that $\mu_i^k > 0$. We will use notations

$$S_i^{max} = \max v_i^{k,j} \quad \text{over all } k \text{ and } j;$$
$$\boldsymbol{S}^{max} = (S_1^{max}, \ldots, S_N^{max}).$$

We denote by $\boldsymbol{S}(t) = (S_1(t), \ldots, S_N(t))$ the (random) realization of the service vector at time t, and call $\boldsymbol{S}(\cdot)$ the service process.

After the service at time t is completed, a random amount of work arrives into the queues, and it is given by a non-negative vector $\boldsymbol{A}(t) = (A_1(t), \ldots, A_N(t))$. The values of $\boldsymbol{A}(t)$ are i.i.d. across times t, and $\boldsymbol{A}(\cdot)$ is called the arrival process. The mean arrival rates of this process are given by vector

$$\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_N) = E\boldsymbol{A}(t).$$

We will now make assumptions on the distribution of $\boldsymbol{A}(t)$. The distribution is absolutely continuous w.r.t. Lebesgue measure, it is concentrated on the rectangle $[\boldsymbol{0}, \boldsymbol{A}^{\max}]$ for some constant vector $\boldsymbol{A}^{\max} > \boldsymbol{S}^{\max}$; moreover, on this rectangle the distribution density $f(\boldsymbol{x})$ is both upper and lower bounded by positive constants, i.e. $0 < \delta_* \leq f(\boldsymbol{x}) \leq \delta^*$. These "continuous" assumptions on the arrival process are used, in particular, in Lemmas 6-8, 11, 14.

If $\boldsymbol{Q}(t) \triangleq (Q_1(t), \ldots, Q_N(t))$ is the vector of queue lengths at time t, then for each $i = 1, \ldots, N$

$$\begin{aligned} Q_i(t+1) &= (Q_i(t) - S_i(t))^+ + A_i(t), \\ &= Q_i(t) + A_i(t) - S_i(t) + U_i(t), \end{aligned} \qquad (1)$$

where $U_i(t) = (S_i(t) - Q_i(t))^+$ is the amount of service "wasted" by flow i at time t.

III. MaxWeight scheduling scheme. Heavy traffic regime

A. MaxWeight definition

Let a vector $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_N) > \boldsymbol{0}$ be fixed. MaxWeight scheduling algorithm chooses, at each time t, a service decision

$$k \in \arg\max_l \left((\boldsymbol{\gamma} \boldsymbol{Q}(t)) \cdot \boldsymbol{\mu}^l \right); \qquad (2)$$

with ties broken according to any well defined rule.

Under MaxWeight, the queue length process $\boldsymbol{Q}(\cdot)$ is a discrete time Markov chain with (continuous) state space $\mathbb{R}_+^N$. System stability is understood as positive Harris recurrence of this Markov chain.

Denote the system *rate region* by

$$\begin{aligned} \boldsymbol{V} \triangleq \Big\{ & \boldsymbol{x} \in \mathbb{R}_+^N : \\ & \boldsymbol{x} \le \sum_k \psi_k \boldsymbol{\mu}^k \text{ for some } \psi_k \ge 0, \sum_k \psi_k = 1 \Big\} \end{aligned} \qquad (3)$$

It is well known (see [4], [10], [11]) that, in general, under MaxWeight the system is stable as long as the vector of mean arrival rates $\boldsymbol{\lambda}$ is such that $\boldsymbol{\lambda} < \boldsymbol{x} \in \boldsymbol{V}$. (Scheduling rules having this property are sometimes called "throughput-optimal".) This is true for our model as well as will be shown in Section VII. (Establishing this fact is not difficult, but it does not directly follow from previous work, because we have continuous state space.)

B. Heavy traffic regime

We will consider a sequence of systems, indexed by $n \to \infty$, operating under MaxWeight scheduling. (Variables pertaining to n-th system will be supplied superscript (n).) The switch parameters will remain unchanged, but the distribution of $\boldsymbol{A}^{(n)}(t)$ changes with n: namely, for each n it has density $f^{(n)}$ which satisfies all conditions specified in Section II, and $f^{(n)}$ uniformly converges to some density f^*. Note that, automatically, the limiting density $f^\star$ (as well as each $f^{(n)}$) satisfies bounds $0 < \delta_* \le f^\star(\boldsymbol{x}) \le \delta^*$ in the rectangle $[\boldsymbol{0}, \boldsymbol{A}^{max}]$, and is zero elsewhere. The arrival process $\boldsymbol{A}^\star(\cdot)$, such that the distribution of $\boldsymbol{A}^\star(t)$ has density $f^\star$, has the arrival rate vector $\boldsymbol{\lambda}^\star$. Correspondingly, $\boldsymbol{\lambda}^{(n)} \to \boldsymbol{\lambda}^\star$.

We assume that $\boldsymbol{\lambda}^\star > \boldsymbol{0}$ is a maximal element of rate region $\boldsymbol{V}$, i.e. $\boldsymbol{x} \ge \boldsymbol{\lambda}^\star$ and $\boldsymbol{x} \in \boldsymbol{V}$ only when $\boldsymbol{x} = \boldsymbol{\lambda}^\star$. Thus, $\boldsymbol{\lambda}^\star$ lies on the outer boundary of $\boldsymbol{V}$. We further assume that for each n, $\boldsymbol{\lambda}^{(n)}$ lies in the *interior* of $\boldsymbol{V}$; therefore, the system is stable for each n (under the MaxWeight algorithm).

The (limiting) system, with arrival process $\boldsymbol{A}^\star(\cdot)$ is called critically loaded.

IV. Main Results

Consider the sequence of systems described in Section III, in the heavy traffic regime. Under any throughput-optimal scheduling algorithm, for each n, the steady-state average amount of service provided to each flow i is greater or equal to its arrival rate λ_i. (It may, and typically will, be greater if the wasted service is taken into account.)

We now define the notion of *asymptotic smoothness* of the steady-state service process. Informally, it means the property that as the system load approaches critical, the steady state service processes are such that for each flow the probability of a T-long gap (without any service at all) uniformly vanishes, as $T \to \infty$.

For each n, consider the cumulative service process $\boldsymbol{G}^{(n)}(\cdot)$ in steady state. Namely,

$$\boldsymbol{G}^{(n)}(t) \triangleq \sum_{\tau=1}^{t} \boldsymbol{S}^{(n)}(\tau), \quad t = 1, 2, \ldots$$

Definition 1: We call the service process asymptotically smooth, if

$$\max_i \lim_{T\to\infty} \left(\limsup_{n\to\infty} P\left(G_i^{(n)}(T) = 0 \right) \right) = 0. \qquad (4)$$

Our key result (Theorem 21 in Section IX) shows that a "queue-differential" process, which determines scheduling decisions in the system under MaxWeight in heavy traffic, is such that its stationary version converges to that of stationary positive Harris recurrent Markov chain, whose structure we describe explicitly. This result, in particular, will imply the following

Theorem 2: Consider the sequence of systems described in Section III, in the heavy traffic regime. Under MaxWeight scheduling, the service process is asymptotically smooth.

The proof is given in Section IX.

V. Complete Resource Pooling condition

COM As stated earlier, to improve exposition, we give detailed proofs of our main results for the special case, when the following *complete resource pooling* (CRP) condition holds. (The generalization to the non-CRP case can be found in [12].) Assume that vector $\boldsymbol{\lambda}^\star$ is such that there is a unique (up to scaling) outer normal vector $\boldsymbol{\nu} > \boldsymbol{0}$ to $\boldsymbol{V}$ at point $\boldsymbol{\lambda}^\star$; we choose $\boldsymbol{\nu}$ so that $\|\boldsymbol{\nu}\| = 1$. Denote by

$$\boldsymbol{V}^\star \triangleq \arg\max_{\boldsymbol{x} \in \boldsymbol{V}} \boldsymbol{\nu} \cdot \boldsymbol{x} \qquad (5)$$

the outer face of $\boldsymbol{V}$ where $\boldsymbol{\lambda}^\star$ lies. Given our assumptions on $\boldsymbol{\lambda}^\star$, it lies in the relative interior of $\boldsymbol{V}^\star$.

By $\boldsymbol{\nu}_\perp$ we denote the subspace of $\mathbb{R}^N$ orthogonal to $\boldsymbol{\nu}$. For any vector $\boldsymbol{a}$, we denote by $\boldsymbol{a}_\star \triangleq (\boldsymbol{a} \cdot \boldsymbol{\nu}) \boldsymbol{\nu}$ its orthogonal projection on the (one-dimensional) subspace spanned by $\boldsymbol{\nu}$, and by $\boldsymbol{a}_\perp \triangleq \boldsymbol{a} - \boldsymbol{a}_\star$ its orthogonal projection on the $(N-1)$-dimensional subspace $\boldsymbol{\nu}_\perp$.

The following observations and notations will be useful. There is a $\delta > 0$ such that the entire set

$$B^{\delta}_{\boldsymbol{\lambda}^\star} \triangleq \{\boldsymbol{y} \in \boldsymbol{V}^\star : \|\boldsymbol{y} - \boldsymbol{\lambda}^\star\| \leq \delta\}, \quad (6)$$

also lies in the relative interior of $\boldsymbol{V}^\star$.

VI. Background on general-state-space discrete-time Markov Chains

We will briefly discuss some notions and results from [9] and [5] on the stability of discrete time Markov Chains (MC), which will be used in later sections. Throughout this section we will assume that the Markov Chain $\boldsymbol{\Phi} = \{\Phi(0), \Phi(1), \ldots\}$ is evolving on a locally compact separable metric space $\boldsymbol{X}$ whose Borel σ-algebra will be denoted by $\mathcal{B}$. P_η and E_η are used to denote the probabilities and expectations conditional on Φ_0 having distribution η, while $P_{\boldsymbol{x}}$ and $E_{\boldsymbol{x}}$ are used when η is concentrated at $\boldsymbol{x}$. The transition function of $\boldsymbol{\Phi}$ is denoted by $P(\boldsymbol{x}, A), \boldsymbol{x} \in \boldsymbol{X}, A \in \mathcal{B}$. The iterates $P^t,\ t = 0, 1, 2, \ldots$, are defined inductively by

$$P^0 \triangleq I, P^t \triangleq PP^{t-1}, t \geq 1,$$

where I is the identity transition function.

Definition 3: (i) ϕ-irreducibility: A Markov Chain $\boldsymbol{\Phi} = \{\Phi(0), \Phi(1), \ldots\}$ is called ϕ irreducible if there exists a finite measure ϕ such that $\sum_{k=1}^{\infty} P^k(\boldsymbol{x}, A) > 0$ for all $\boldsymbol{x} \in \boldsymbol{X}$ whenever $\phi(A) > 0$. Measure ϕ is called an irreducibility measure.

(ii) Harris Recurrence: If $\boldsymbol{\Phi}$ is ϕ-irreducible and $P_x(\Phi(t) \in A \quad \text{i.o.}) \equiv 1$ whenever $\phi(A) > 0$, then $\boldsymbol{\Phi}$ is called Harris recurrent. [Abbreviation 'i.o.' means 'infinitely often'.]

(iii) Invariant Measure: A σ-finite measure π on $\mathcal{B}$ with the property

$$\pi\{A\} = \pi P\{A\} \triangleq \int \pi(d\boldsymbol{x}) P(\boldsymbol{x}, A), \forall A \in \mathcal{B},$$

is called an invariant measure.

(iv) Positive Harris Recurrence: If $\boldsymbol{\Phi}$ is Harris Recurrent with a finite invariant measure π, then it is called positive Harris Recurrent.

(v) Boundedness in Probability: If for any $\epsilon > 0$ and any $\boldsymbol{x} \in \boldsymbol{X}$, there exists a compact set D such that

$$\liminf_{t \to \infty} P_{\boldsymbol{x}}(\Phi(t) \in D) \geq 1 - \epsilon, \quad (7)$$

then the Markov process $\boldsymbol{\Phi}$ is called bounded in probability.

(vi) Small Sets: A set C is called small if for all $\boldsymbol{x} \in C$ and some integer $l \geq 1$, we have

$$P^l(\boldsymbol{x}, \cdot) \geq \nu(\cdot), \quad (8)$$

where $\nu(\cdot)$ is a sub-probability measure, i.e. $\nu(\boldsymbol{X}) \leq 1$.

(vii) For a probability distribution $\boldsymbol{a} = (a_1, a_2 \ldots)$ on $\{1, 2, \ldots\}$, the Markov transition function $K_{\boldsymbol{a}}$ is defined as

$$K_{\boldsymbol{a}} \triangleq \sum_{i=1}^{\infty} a_i P^i.$$

(viii) Petite Sets: A set $A \in \mathcal{B}$ and a sub-probability measure ψ on $\mathcal{B}(\boldsymbol{X})$ are called petite if for some probability distribution $\boldsymbol{a}$ on $\{1, 2, \ldots\}$ we have

$$K_{\boldsymbol{a}}(\boldsymbol{x}, \cdot) \geq \psi(\cdot), \forall \boldsymbol{x} \in A.$$

(ix) Non-evanescence: A Markov chain $\boldsymbol{\Phi}$ is called non-evanescent if $P_{\boldsymbol{x}}\{\boldsymbol{\Phi} \to \infty\} = 0$ for each $\boldsymbol{x} \in \boldsymbol{X}$. [Event $\{\boldsymbol{\Phi} \to \infty\}$ consists of the outcomes such that the sequence $\Phi(t)$ visits any compact set at most a finite number of times.]

The following proposition states some results from [9].

Proposition 4: (i) If a set A is small and for some probability distribution $\boldsymbol{a}$ on $\{1, 2, \ldots\}$ and a set $B \in \mathcal{B}$, we have

$$\inf_{\boldsymbol{x} \in B} K_{\boldsymbol{a}}(\boldsymbol{x}, A) > 0, \quad (9)$$

then B is petite.

(ii) Suppose that every compact subset of $\boldsymbol{X}$ is petite. Then $\boldsymbol{\Phi}$ is positive Harris recurrent if and only if it is bounded in probability.

(iii) Suppose that every compact subset of $\boldsymbol{X}$ is petite. Then $\boldsymbol{\Phi}$ is Harris recurrent if and only if it is non-evanescent.

The following result is form from [5]. It is stated in a form convenient for its application in this paper.

Proposition 5: Let $L(\boldsymbol{x})$ be a non-negative (Lyapunov) function such that the Markov process $\boldsymbol{\Phi}$ satisfies the following two conditions, for some positive constants κ, δ, D:
(a) $E\left[L(\Phi(t+1)) - L(\Phi(t)) | \Phi(t) = \boldsymbol{x}\right] < -\delta$, for any state $\boldsymbol{x}$ such that $L(\boldsymbol{x}) \geq \kappa > 0$.
(b) $|L(\Phi(t+1)) - L(\Phi(t))| < D$.
Then there exist constants $\eta > 0$ and $0 < \rho < 1$ such that

$$P(L(\Phi(t)) \geq u \mid L(\Phi(0)) = b) \leq$$

$$\rho^t \exp(\eta(b - u)) + \frac{1 - \rho^t}{1 - \rho} D \exp(\eta(\kappa - u)),\ u \geq 0. \quad (10)$$

VII. Queue length process

Recall that $\boldsymbol{Q}^{(n)}(\cdot)$ is the queue length process for the n-th system under MaxWeight. In this section we prove that for all n, the process $\boldsymbol{Q}^{(n)}(\cdot)$ is positive Harris recurrent. The proof uses a Lyapunov drift argument which is fairly standard (in fact, there is more than one way to prove stability of $\boldsymbol{Q}^{(n)}(\cdot)$), except, since our state space is continuous, as a first step we will show that all compact sets are petite.

The following Lemmas 6 – 8 are rather simple technical statements. Their proof can found in [12].

Lemma 6: (i) The points $\boldsymbol{x} \in \mathbb{R}^N_+$, such that $k \in \arg\max_\ell (\boldsymbol{\gamma x}) \cdot \boldsymbol{\mu}^l$ is non-unique, form a set of zero Lebesgue measure. Moreover, if $\boldsymbol{x} > \boldsymbol{0}$ is such that $k \in \arg\max_\ell (\boldsymbol{\gamma x}) \cdot \boldsymbol{\mu}^l$ is unique, then for a sufficiently small $\epsilon > 0$ the decision k is also the unique element of $\arg\max_\ell (\boldsymbol{\gamma y}) \cdot \boldsymbol{\mu}^l$ for all $\boldsymbol{y} \in B_\epsilon(\boldsymbol{x})$.
(ii) The one-step transition function $P^{(n)}(\boldsymbol{x}, \cdot)$ of the process $\boldsymbol{Q}^{(n)}(\cdot)$ is such that, uniformly in n and $\boldsymbol{x} \in \mathbb{R}^N_+$, the distribution $P^{(n)}(\boldsymbol{x}, \cdot)$ is absolutely continuous with the density upper bounded by δ^* and, in the rectangle $[\boldsymbol{0}, \boldsymbol{A}^{max} - \boldsymbol{S}^{max}]$, lower bounded by δ_*.

Lemma 7: For any $\boldsymbol{x} > \boldsymbol{0}$, there exists $\epsilon > 0$ such that the set $B_\epsilon(\boldsymbol{x})$ is small for the process $\boldsymbol{Q}^{(n)}(\cdot)$.

Lemma 8: For the Markov process $\boldsymbol{Q}^{(n)}(\cdot)$, any compact set is petite.

To prove stability, we will apply Proposition 4 which requires the following

Lemma 9: Consider the scalar projection $\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(\cdot)\|$, $t = 0, 1, \ldots$ of the the Markov process $\boldsymbol{Q}^{(n)}$ starting with a fixed initial state $\boldsymbol{Q}^{(n)}(0)$, such that $\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(0)\| = b$. Then, uniformly on all large n we have,

$$P(\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\| \geq u) \leq \rho^t \exp(\eta(b-u)) + \frac{1-\rho^t}{1-\rho} D \exp(\eta(\kappa - u)),\ u \geq 0, \tag{11}$$

for some constants $\eta, \kappa, D > 0$ and $1 > \rho > 0$ which depend on n. Consequently, the process $\boldsymbol{Q}^{(n)}(\cdot)$ is bounded in probability.

Proof: We will use notation $L(\boldsymbol{x}) = \|\sqrt{\boldsymbol{\gamma}}\boldsymbol{x}\|$. Then $L(\boldsymbol{Q}^{(n)}(0)) = b$. Clearly, $|L(\boldsymbol{Q}^{(n)}(t+1)) - L(\boldsymbol{Q}^{(n)}(t))|$ is uniformly bounded by a constant, given our assumptions on the arrival and service processes. We will show that the drift (average increment) of $L(\boldsymbol{Q}^{(n)}(t+1))) - L(\boldsymbol{Q}^{(n)}(t)))$ is upper bounded by some $-\tilde{\delta} < 0$ when $\|L(\boldsymbol{Q}^{(n)}(t))\| \geq \kappa$ for some $\kappa > 0$.

Consider a fixed $\boldsymbol{Q}^{(n)}(t)$ and denote $\Delta L = E[L(\boldsymbol{Q}^{(n)}(t+1)) - L(\boldsymbol{Q}^{(n)}(t))]$. Clearly,

$$\begin{aligned}\Delta L &= E\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t+1)\| - \|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\| \\ &\leq \frac{1}{2\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\|}\left(E\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t+1)\|^2 \right. \\ &\qquad \left. - \|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\|^2\right),\end{aligned} \tag{12}$$

where the inequality follows from the concavity of the function $\sqrt{x}$. Substitute the value of $\boldsymbol{Q}^{(n)}(t+1)$ from equation (1), concentrate on the numerator of the above expression to obtain,

$$\begin{aligned}&E\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t+1)\|^2 - \|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\|^2 \\ &= E\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t) + \sqrt{\boldsymbol{\gamma}}\left(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t)\right)\|^2 \\ &\qquad - \|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\|^2 \\ &= E[\|\sqrt{\boldsymbol{\gamma}}(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t))\|^2 \\ &\qquad + 2\left(\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\right) \cdot \left(\sqrt{\boldsymbol{\gamma}}\left(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t)\right)\right)] \\ &= E[\|\sqrt{\boldsymbol{\gamma}}(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t))\|^2 \\ &\qquad + 2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot (\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t))] \\ &= E[\|\sqrt{\boldsymbol{\gamma}}(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t))\|^2 \\ &\qquad + 2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot \boldsymbol{U}^{(n)}(t) \\ &\qquad + 2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot \left(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t)\right)] \\ &\leq b_1 + b_2 \\ &\qquad + 2E\left[\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot \left(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t)\right) | \boldsymbol{Q}^{(n)}(t)\right],\end{aligned} \tag{13}$$

where b_1 is a uniform bound on $\|\sqrt{\boldsymbol{\gamma}}(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t) + \boldsymbol{U}^{(n)}(t))\|^2$, and b_2 is a uniform bound on $\|2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot \boldsymbol{U}^{(n)}(t)\|$ which follows from the property that $U_i(t) > 0$ only when $Q_i(t)$ is sufficiently small.

To simplify exposition and avoid introducing additional notation, let us assume that $\boldsymbol{\lambda}^{(n)} - \boldsymbol{\lambda}^\star = -\epsilon\boldsymbol{\nu}$ for some $\epsilon > 0$. (If not, then instead of $\boldsymbol{\lambda}^\star$ in this proof we can use $\boldsymbol{\lambda}^{\star\star}$, which the orthogonal projection of $\boldsymbol{\lambda}^{(n)}$ on $\boldsymbol{V}^\star$.) Combining (12) and (13), we obtain

$$\begin{aligned}&2\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\|\Delta L \\ &\leq b_1 + b_2 + 2E\left[\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot \left(\boldsymbol{A}^{(n)}(t) - \boldsymbol{S}^{(n)}(t)\right)\right] \\ &= b_1 + b_2 + 2E\left[\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)\right. \\ &\qquad \left. \cdot \left(\boldsymbol{A}^{(n)}(t) - \boldsymbol{\lambda}^\star + \boldsymbol{\lambda}^\star - \boldsymbol{S}^{(n)}(t)\right)\right] \\ &= b_1 + b_2 - 2\epsilon\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\star\| \\ &\qquad + 2E\left[\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right) \cdot \left(\boldsymbol{\lambda}^\star - \boldsymbol{S}^{(n)}(t)\right)\right] \\ &\leq b_1 + b_2 - 2\epsilon\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\star\| - \delta\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\perp\|,\end{aligned} \tag{14, 15}$$

where the last inequality follows from the definition of MaxWeight (see (2)) and the set $B^\delta_{\boldsymbol{\lambda}^\star}$ (see (6)). If $\|\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\| \geq x$, then at least one of $\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\star\|$ or $\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\perp\|$ is greater than or equal to $x/\sqrt{2}$. After some algebraic manipulations we obtain ($\gamma_{min} = \min_i \gamma_i$),

$$\begin{aligned}&\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\| > x \\ &\implies \|\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\| > \sqrt{\gamma_{\min}}x \\ &\implies \|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\star\| \vee \|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\perp\| \geq \frac{\sqrt{\gamma_{\min}}x}{\sqrt{2}} \\ &\implies \delta\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\star\| + \epsilon\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t)\right)_\perp\| \\ &\qquad \geq (\epsilon \wedge \delta)\frac{\sqrt{\gamma_{\min}}x}{\sqrt{2}}.\end{aligned}$$

Substituting the above in inequality (14) we see that the drift is upper bounded by

$$-(\epsilon \wedge \delta)\frac{\sqrt{\gamma_{\min}}x}{2\sqrt{2}} + \frac{b_1 + b_2}{\|\sqrt{\boldsymbol{\gamma}}\boldsymbol{Q}^{(n)}(t)\|}.$$

This quantity is uniformly bounded by a negative constant for sufficiently large x. Application of Proposition 5 completes the proof. ■

Now the positive recurrence of $\boldsymbol{Q}^{(n)}(\cdot)$ follows from Proposition 4. In fact, we will prove the following stronger statement.

Theorem 10: For each $n = 1, 2, \ldots$, the Markov process $\boldsymbol{Q}^{(n)}(\cdot)$ is positive Harris recurrent and hence has a unique invariant probability distribution, which will be denoted $\chi^{(n)}$.

Moreover, if $\boldsymbol{Q}^{(n)}(\infty)$ is the (random) process state in stationary regime (i.e. it has distribution $\chi^{(n)}$),

$$E[\|\boldsymbol{Q}^{(n)}(\infty)\|^r] < \infty, \forall r > 0.$$

Proof: By Lemma 8 any compact set is petite. Since $\boldsymbol{Q}^{(n)}(\cdot)$ is also bounded in probability (Lemma 9), by Proposition 4 $\boldsymbol{Q}^{(n)}(\cdot)$ is positive Harris recurrent.

For a function $f(\cdot)$ and fixed $b > 0$, denote $T_b f(\cdot) = f(\cdot) \wedge b$. Consider the process starting from an arbitrary fixed initial state $\boldsymbol{Q}^{(n)}(0)$. Since the process is positive Harris recurrent, we can apply the ergodic theorem to obtain (note that $T_b\|\cdot\|$ is a bounded continuous function):

$$E\left(T_b\|\boldsymbol{Q}^{(n)}(\infty)\|^r\right) = \lim_{m\to\infty} \frac{1}{m}\sum_{t=0}^{m} E\left[T_b\|\boldsymbol{Q}^{(n)}(t)\|^r\right]. \tag{16}$$

On the other hand,

$$\begin{aligned}
\lim_{m\to\infty} \frac{1}{m}\sum_{t=0}^{m} & E\left[T_b\|\boldsymbol{Q}^{(n)}(t)\|^r\right] \\
&\leq \lim_{m\to\infty} \frac{1}{m}\sum_{t=0}^{m} E\left[\|\boldsymbol{Q}^{(n)}(t)\|^r\right] \\
&< C,
\end{aligned} \tag{17}$$

for some constant $C > 0$, where the second inequality follows from (11). Combining (16) and (17), we have

$$E\left(T_b\|\boldsymbol{Q}^{(n)}(\infty)\|^r\right) \leq C, \quad \forall b > 0, \tag{18}$$

and therefore, by monotone convergence theorem,

$$E\left(\|\boldsymbol{Q}^{(n)}(\infty)\|^r\right) = \lim_{b\to\infty} E\left(T_b\|\boldsymbol{Q}^{(n)}(\infty)\|^r\right) \leq C.$$

■

Lemma 11: Uniformly on all (large) n and the distributions of $\boldsymbol{Q}^{(n)}(0)$, the distribution of $\boldsymbol{Q}^{(n)}(1)$ is absolutely continuous w.r.t. Lebesgue measure, with the density upper bounded by δ^*.

We omit the proof, which is straightforward, given our assumptions on the distribution of $\boldsymbol{A}^{(n)}(t)$.

Lemma 12: As $n \to \infty$, $\|\boldsymbol{Q}^{(n)}(\infty)\| \to \infty$ in probability.

Proof: The proof is by contradiction. Suppose, for some fixed $C > 0$ the compact set $D = \{\boldsymbol{x} \in \mathbb{R}^N : \|\boldsymbol{x}\| \leq C\}$ is such that

$$\limsup_{n\to\infty} \chi^{(n)}(D) = \beta > 0. \tag{19}$$

Suppose $\boldsymbol{Q}^{(n)}(t) \in D$. Then, using the same argument as in the proof of Lemma 8, it is easy to see that for any $\epsilon > 0$ there exists time $\tau \geq 1$, such that $P\{\|\boldsymbol{Q}^{(n)}(t+\tau)\| \leq \epsilon\} \geq \beta_1 > 0$. This in turn implies that, with probability at least some $\beta_2 > 0$, for at least one flow i the amount of wasted service $U_i^{(n)}(t+\tau) \geq \epsilon_2 > 0$. This implies that, for at least one i,

$$\limsup_{n\to\infty} E[U_i^{(n)}(\infty)] \geq \beta_1\beta_2\epsilon_2 > 0.$$

This, however, contradicts the fact that the process is stable for all large n. ■

VIII. Steady-state queue lengths deviations from $\boldsymbol{\nu}$

Let us consider the process $\boldsymbol{Y}^{(n)}(\cdot)$, defined as

$$\boldsymbol{Y}^{(n)}(t) := (\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(t))_\perp.$$

Lemma 13: The steady-state expected norm $E\|\boldsymbol{Y}^{(n)}(\infty)\|$ is uniformly bounded in n.

Proof: As we did in the proof of Lemma 9, to simplify exposition, assume that $\boldsymbol{\lambda}^{(n)} - \boldsymbol{\lambda}^\star = -\epsilon\boldsymbol{\nu}$. (If not, in this proof we would consider the projection $\boldsymbol{\lambda}^{\star\star}$ of $\boldsymbol{\lambda}^{(n)}$ on $\boldsymbol{V}^\star$, instead of $\boldsymbol{\lambda}^\star$. Consider Lyapunov function $L(\boldsymbol{Q}) = \sum_{i=1}^N \gamma_i Q_i^2$. By Theorem 10, $EL(\boldsymbol{Q}^{(n)}(\infty)) < \infty$. The conditional drift of $L(\boldsymbol{Q})$ in one time step is given by (let $\boldsymbol{Q}^{(n)}(t) = \boldsymbol{Q}^{(n)}$, $\boldsymbol{A}^{(n)}(t) = \boldsymbol{A}^{(n)}$, and so on, to simplify notation)

$$\begin{aligned}
&E\left[L(\boldsymbol{Q}^{(n)}(t+1)) - L(\boldsymbol{Q}^{(n)}(t))|\boldsymbol{Q}^{(n)}\right] \\
&= E\left[\sum_{i=1}^N \gamma_i \left(Q_i^{(n)} + A_i^{(n)} - S_i^{(n)} + U_i^{(n)}\right)^2 |\boldsymbol{Q}^{(n)}\right] \\
&\quad - \sum_{i=1}^N \gamma_i \left(Q_i^{(n)}\right)^2 \\
&= E\left[\sum_{i=1}^N \gamma_i \left(A_i^{(n)} - S_i^{(n)} + U_i^{(n)}\right)\left(2Q_i^{(n)}\right.\right. \\
&\qquad\qquad \left.\left. + A_i^{(n)} - S_i^{(n)} + U_i^{(n)}\right) |\boldsymbol{Q}^{(n)}\right] \\
&= E\left[\sum_{i=1}^N \gamma_i \left(A_i^{(n)} - S_i^{(n)} + U_i^{(n)}\right)^2\right. \\
&\qquad \left. + 2\gamma_i Q_i^{(n)}\left(A_i^{(n)} - S_i^{(n)} + U_i^{(n)}\right) |\boldsymbol{Q}^{(n)}\right] \\
&\leq b_1 + 2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right) \cdot \left(\boldsymbol{\lambda}^{(n)} - E\left(\boldsymbol{S}^{(n)}|\boldsymbol{Q}^{(n)}\right)\right) \\
&= b_1 \\
&+ 2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right) \cdot \left(\boldsymbol{\lambda}^{(n)} - \boldsymbol{\lambda}^\star + \boldsymbol{\lambda}^\star - E\left(\boldsymbol{S}^{(n)}|\boldsymbol{Q}^{(n)}\right)\right) \\
&= b_1 - 2\epsilon\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right)_\star\| \\
&\qquad + 2\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right) \cdot \left(\boldsymbol{\lambda}^\star - E\left(\boldsymbol{S}^{(n)}|\boldsymbol{Q}^{(n)}\right)\right) \\
&\leq b_1 - 2\epsilon\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right)_\star\| \\
&\quad + 2\min_{\boldsymbol{y}\in B_\nu^\delta} \left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right) \cdot (\boldsymbol{\lambda}^\star - \boldsymbol{y}) \\
&\leq b_1 - 2\epsilon\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}\right)_\star\| - 2\delta\|(\boldsymbol{\gamma}\boldsymbol{Q})_\perp\|,
\end{aligned} \tag{20}$$

where b_1 depends only on $\boldsymbol{\gamma}, \boldsymbol{A}^{max}, \boldsymbol{S}^{max}$, and the last inequality follows from the definition of MaxWeight and $B_{\boldsymbol{\lambda}^\star}^\delta$. Now consider the process $\boldsymbol{Q}^{(n)}(\cdot)$ in stationary regime, and take the expectation of both parts of (20). We obtain,

$$\begin{aligned}
2\delta E&\left[\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(\infty)\right)_\perp\|\right] + 2\epsilon E\left[\|\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(\infty)\right)_\star\|\right] \\
&\leq b_1.
\end{aligned} \tag{21}$$

Recalling that $\left(\boldsymbol{\gamma}\boldsymbol{Q}^{(n)}(\infty)\right)_\perp = \boldsymbol{Y}^{(n)}(\infty)$, we see that $E\|\boldsymbol{Y}^{(n)}(\infty)\|$ is uniformly bounded. ■

IX. Limit of the queue-differential process

We now define a Markov chain $\boldsymbol{Y}^\star(\cdot)$, which, in the sense that will be made precise later, is a limit of the (non-Markov) process $\boldsymbol{Y}^{(n)}(\cdot)$ as $n \to \infty$.

Define $\boldsymbol{Y}^{(n)}(t)$ as the orthogonal projection of $\boldsymbol{\gamma Q}^{(n)}(t)$ on the subspace $\boldsymbol{\nu}_\perp$. We call $\boldsymbol{Y}^{(n)}(\cdot)$ a *queue-differential* process. (Obviously, *under the CRP condition*, the queue-differential process is equal to the "queue deviation" process $\boldsymbol{Y}^{(n)}(\cdot) = (\boldsymbol{\gamma Q}^{(n)}(t))_\perp$ in Section VIII. When CRP does not hold, the "deviation" and "differential" processes are defined differently. See [12].) Denote by $\boldsymbol{Y}^{(n)}(\infty)$ the corresponding projection of the steady-state $\boldsymbol{Q}^{(n)}(\infty)$, and by $\Gamma^{(n)}$ its distribution.

Markov chain $\boldsymbol{Y}^\star(\cdot)$ is defined formally as follows. (We will show below that, in fact, the distribution $\Gamma^{(n)}$ converges to the stationary distribution $\Gamma^\star$ of $\boldsymbol{Y}^\star(\cdot)$.) The state space of $\boldsymbol{Y}^\star(\cdot)$ is $\boldsymbol{\nu}_\perp$. Assume that at time t the "scheduler" chooses decision

$$k \in \arg\max_{l:\boldsymbol{\mu}^l \in \boldsymbol{V}^\star} (\boldsymbol{Y}^\star(t)) \cdot \boldsymbol{\mu}^l, \tag{22}$$

which determines the corresponding random amount of service $\boldsymbol{S}(t)$, provided to the "queues" given by vector $\boldsymbol{Q}^\star(t) = \boldsymbol{Y}^\star(t)/\boldsymbol{\gamma}$. After that the (random) amount $\boldsymbol{A}^\star(t)$ of new "work" arrives and is added to the "queues." Finally, the new queue lengths vector $\boldsymbol{Q}^\star(t) - \boldsymbol{S}(t) + \boldsymbol{A}^\star(t)$ is transformed into $\boldsymbol{Y}^\star(t+1)$ via componentwise multiplication by $\boldsymbol{\gamma}$ and orthogonal projection on $\boldsymbol{\nu}_\perp$. (Note that both $\boldsymbol{Q}^\star(t)$ and $\boldsymbol{Y}^\star(t)$ may have components of any sign. Also, there is no "wasted service" here.) In summary, the one step evolution is described by

$$\boldsymbol{Y}^\star(t+1) = \boldsymbol{Y}^\star(t) + (\boldsymbol{\gamma A}^\star(t) - \boldsymbol{\gamma S}(t))_\perp . \tag{23}$$

Informally, one can interpret the process $\boldsymbol{Y}^\star(\cdot)$ as the queue-differential process $\boldsymbol{Y}^{(n)}(\cdot)$, when n is very large and the queue length vector $\boldsymbol{Q}^{(n)}$ is both large and has a small angle with $\boldsymbol{\nu}$. Under these conditions, *the only service decisions k that can be chosen are such that $\boldsymbol{\mu}^k \in \boldsymbol{V}^\star$, and the choice is uniquely determined by $\boldsymbol{Y}^{(n)}(\cdot)$.*

Let $\tilde{P}(\boldsymbol{x}, \cdot)$ denote the one-step transition function for the Markov process $\boldsymbol{Y}^\star(\cdot)$. If $\boldsymbol{x} \in \boldsymbol{\nu}_\perp$, then let $\tilde{B}_\epsilon(\boldsymbol{x}) := \{\boldsymbol{y} \in \boldsymbol{\nu}_\perp : ||\boldsymbol{y} - \boldsymbol{x}|| \le \epsilon\}$. The following fact is analogous to Lemma 6.

Lemma 14: (i) The points $\boldsymbol{y} \in \boldsymbol{\nu}_\perp$, such that

$$k \in \arg\max_{l:\boldsymbol{\mu}^l \in \boldsymbol{V}^\star} \boldsymbol{y} \cdot \boldsymbol{\mu}^l \tag{24}$$

is non-unique, form a set of zero Lebesgue measure. Moreover, if $\boldsymbol{y}$ is such that the corresponding decision k is unique, then for a sufficiently small $\epsilon > 0$ the decision k is also the unique element of

$$\arg\max_{l:\boldsymbol{\mu}^l \in \boldsymbol{V}^\star} \boldsymbol{z} \cdot \boldsymbol{\mu}^l$$

for all $\boldsymbol{z} \in \tilde{B}_\epsilon(\boldsymbol{y})$.
(ii) There exist small $\epsilon > 0$ and constant $c_* > 0$, $c^* > 0$ such that $\tilde{P}(\boldsymbol{x}, \cdot)$ is absolutely continuous and, moreover, uniformly in $\boldsymbol{x} \in \boldsymbol{\nu}_\perp$, the density of $\tilde{P}(\boldsymbol{x}, \cdot)$ is lower bounded by c_* on set $\tilde{B}_\epsilon(\boldsymbol{x})$ and is upper bounded by c^* everywhere.

Proof: Statement (i) is obvious. Statement (ii) follows from our assumptions on the distribution of $\boldsymbol{A}^\star(t)$, the fact that $\boldsymbol{A}^{max} > \boldsymbol{S}^{max}$, and the one-step evolution rule (23). We omit details. ■

Lemma 15: For the Markov chain $\boldsymbol{Y}^\star(\cdot)$, every compact set is petite.

The proof easily follows from Lemma 14, by using the argument analogous to that in the proof of Lemma 8. We omit details.

Next, we establish some properties of a stationary distribution $\Gamma^\star$ of the Markov process $\boldsymbol{Y}^\star(\cdot)$, *assuming a stationary distribution exists*. This will help us later prove that the stationary distribution in fact exists and is unique.

Lemma 16: If $\Gamma^\star$ is a stationary distribution of $\boldsymbol{Y}^\star(\cdot)$, then $\Gamma^\star$ is equivalent to the Lebesgue measure $\tilde{\mathcal{L}}$, i.e. $\Gamma^\star \ll \tilde{\mathcal{L}}$ and $\tilde{\mathcal{L}} \ll \Gamma^\star$.

Proof: $\Gamma^\star \ll \tilde{\mathcal{L}}$: This follows from Lemma 14.
$\tilde{\mathcal{L}} \ll \Gamma^\star$: It suffices to show that $\Gamma^\star(\tilde{B}_r(\boldsymbol{z})) > 0$ for any $\boldsymbol{z} \in \boldsymbol{\nu}_\perp$ and $r > 0$. Consider the process $\boldsymbol{Y}^\star(\cdot)$ with the distribution of $\boldsymbol{Y}^\star(0)$ equal to $\Gamma^\star$. (Then the process is of course stationary.) Fix any $0 < \beta < 1$ and choose a compact set $D \subset \boldsymbol{\nu}_\perp$ such that $\Gamma^\star(D) \ge \beta$. Using Lemma 14 we can easily show that there exists time $\tau > 0$ and a constant $\Delta > 0$, such that, uniformly in $\boldsymbol{Y}^\star(0) = \boldsymbol{x} \in D$,

$$P\{\boldsymbol{Y}^\star(\tau) \in \tilde{B}_r(\boldsymbol{z}) \mid \boldsymbol{Y}^\star(0) = \boldsymbol{x}\} \ge \Delta,$$

and therefore

$$\Gamma^\star(\tilde{B}_r(\boldsymbol{z})) \ge \beta\Delta > 0.$$

■

Lemma 17: Suppose $\Gamma^\star$ is a stationary distribution of $\boldsymbol{Y}^\star(\cdot)$. Then $\tilde{P}_{\boldsymbol{x}}(\boldsymbol{Y}^\star \to \infty) = 0,\ \Gamma^\star -$ a.s., and hence $\tilde{P}_{\boldsymbol{x}}(\boldsymbol{Y}^\star(t) \to \infty) = 0,\ \tilde{\mathcal{L}} -$ a.s..

Proof: The proof is by contradiction. Let $\boldsymbol{Y}^\star(0)$ have the stationary distribution $\Gamma^\star$, and assume that $\exists \quad \epsilon > 0, \epsilon_1 > 0$ such that

$$\Gamma^\star(\{\boldsymbol{x} : \tilde{P}_{\boldsymbol{x}}(\boldsymbol{Y}^\star \to \infty) \ge \epsilon_1\}) \ge \epsilon.$$

This would imply that $\limsup_{t\to\infty} P(\boldsymbol{Y}^\star(t) \in D) \le 1 - \epsilon\epsilon_1$ for every compact set $D \subset \boldsymbol{\nu}_\perp$. This is impossible, because the distribution of $\boldsymbol{Y}^\star(t)$ is equal to $\Gamma^\star$ for all t. ■

Lemma 18: If process $\boldsymbol{Y}^\star(\cdot)$ has a stationary distribution, it is non-evanescent.

Proof: Consider process $\boldsymbol{Y}^\star(\cdot)$ with fixed initial state $\boldsymbol{Y}^\star(0) = \boldsymbol{x}$. Consider one-step transition. The distribution of $\boldsymbol{Y}^\star(1)$ is absolutely continuous with respect to $\tilde{\mathcal{L}}$. Thus, by Lemma 17, with probability 1, $\boldsymbol{z} = \boldsymbol{Y}^\star(1)$ is such that $\tilde{P}_{\boldsymbol{z}}(\boldsymbol{Y}^\star \to \infty) = 0$. Then, $\tilde{P}_{\boldsymbol{x}}(\boldsymbol{Y}^\star \to \infty) = 0$. ■

Lemma 19: Suppose $\Gamma^\star$ is a stationary distribution of $\boldsymbol{Y}^\star(\cdot)$. Then, the Markov chain is positive Harris recurrent, and therefore $\Gamma^\star$ is its unique stationary distribution.

Proof: Since every compact set is petite (Lemma 15) and the process is non-evanescent (Lemma 18), it is Harris recurrent by Proposition 4. But since it has a finite invariant measure $\Gamma^\star$, $\boldsymbol{Y}^\star(\cdot)$ is positive Harris recurrent. ■

We now show the existence of a stationary distribution of $\boldsymbol{Y}^{(\star)}(\cdot)$.

Lemma 20: Every weak limit point $\Gamma^{(\star)}$ of the sequence of distributions $\Gamma^{(n)}$ is a stationary distribution of the process $\boldsymbol{Y}^{(\star)}(\cdot)$.

Proof: Let $\Gamma^\star$ be a weak limit of $\Gamma^{(n)}$ along a subsequence on n. We can make the following observations.

(a) Observe that uniformly on all (large) n and the distributions of $\boldsymbol{Q}^{(n)}(0)$, the distribution of $\boldsymbol{Y}^{(n)}(1)$ is absolutely continuous w.r.t. Lebesgue measure, with the upper bounded density. (This easily follows from Lemma 11 and the fact that $\|\boldsymbol{Q}^{(n)}(1) - \boldsymbol{Q}^{(n)}(0)\|$ is uniformly bounded.) Then, we see that $\Gamma^\star$ is absolutely continuous with bounded density.

(b) Consider any point $\boldsymbol{y} \in \boldsymbol{\nu}_\perp$ such that the decision k in (24) is unique and a small $\epsilon > 0$ such that this decision k is also unique for all $\boldsymbol{z} \in \tilde{B}_\epsilon(\boldsymbol{y})$. (See Lemma 14(i).) Then, there exists a sufficiently large $C > 0$ such that, uniformly in n, conditions $\|\boldsymbol{Q}^{(n)}(t)\| \geq C$ and $\boldsymbol{Y}^{(n)}(t) \in \tilde{B}_\epsilon(\boldsymbol{y})$ imply that the same decision k will be unique at time t for the process $\boldsymbol{Q}^{(n)}(\cdot)$.

Using these two observations, Lemma 12, and the fact that the distribution of $\boldsymbol{A}^{(n)}(t)$ converges to that of $\boldsymbol{A}^\star(t)$, we can choose a further subsequence of n along which the following property holds. The stationary versions of processes $\boldsymbol{Q}^{(n)}(\cdot)$ and the process $\boldsymbol{Y}^\star(\cdot)$ with distribution of $\boldsymbol{Y}^\star(0)$ equal to $\Gamma^\star$, can be constructed on one common probability space, so that with probability 1:

(c) for all large n, the same decision k is chosen at time 0 in the processes $\boldsymbol{Q}^{(n)}(\cdot)$ and $\boldsymbol{Y}^\star(\cdot)$;

(d) $\boldsymbol{Y}^{(n)}(0) \to \boldsymbol{Y}^\star(0)$ and $\boldsymbol{Y}^{(n)}(1) \to \boldsymbol{Y}^\star(1)$.

This, in turn, implies that for any bounded continuous function g we have,

$$E\left[g(\boldsymbol{Y}^\star(0))\right] = \lim_{n\to\infty} E\left[g(\boldsymbol{Y}^{(n)}(0))\right],$$
$$E\left[g(\boldsymbol{Y}^\star(1))\right] = \lim_{n\to\infty} E\left[g(\boldsymbol{Y}^{(n)}(1))\right].$$

But, $E\left[g(\boldsymbol{Y}^{(n)}(0))\right] = E\left[g(\boldsymbol{Y}^{(n)}(1))\right]$ for all n. Therefore, $E\left[g(\boldsymbol{Y}^\star(0))\right] = E\left[g(\boldsymbol{Y}^\star(1))\right]$. This proves stationarity of $\Gamma^\star$. ■

Theorem 21: The Markov process $\boldsymbol{Y}^\star(\cdot)$ is positive Harris recurrent. The sequence $\Gamma^{(n)}$ [i.e., the distributions of $\boldsymbol{Y}^{(n)}(\infty)$] weakly converges to the unique stationary distribution $\Gamma^\star$ of $\boldsymbol{Y}^\star(\cdot)$.

Proof: This follows from Lemma 20 and Lemma 19. ■

We are finally in position to give a

Proof of Theorem 2: By Theorem 21, the process $\boldsymbol{Y}^\star(\cdot)$ is positive Harris recurrent. Moreover, we know that it is such that every compact set is petite. We can pick any compact set D such that $\Gamma^\star(D) > 0$, and using Nummelin splitting view the process $\boldsymbol{Y}^\star(\cdot)$ as having an atom state, with finite average return time to this atom. We see that the cumulative "service process" $\boldsymbol{G}^\star(\cdot)$ corresponding to $\boldsymbol{Y}^\star(\cdot)$ in steady-state is such that

$$\max_i \lim_{T\to\infty} P(G_i^\star(T) = 0) = 0.$$

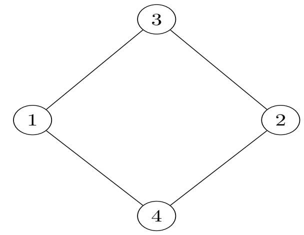

Fig. 1: Interference graph of a network comprising 4 wireless links.

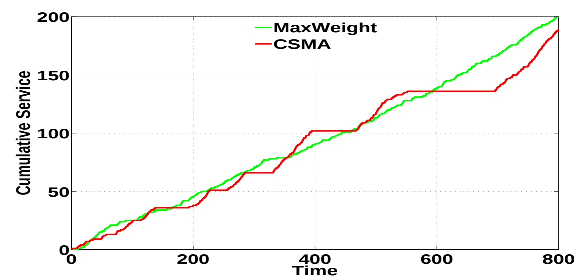

Fig. 2: $\lambda = .245$, $\mu = .5$, load $= .98$, $\gamma_i \equiv 1$.

Finally, the argument used in the proof of Lemma 20 shows that the stationary versions of processes $\boldsymbol{Y}^\star(\cdot)$ and $\boldsymbol{Q}^{(n)}(\cdot)$ for all (large) n can be constructed on a common probability space in a way such that, w.p.1, for any $T > 0$

$$\boldsymbol{G}^{(n)}(T) \to \boldsymbol{G}^\star(T).$$

This implies (4). ■

X. Simulations

In this section we illustrate via simulation that, as our theoretical results show, the service process under MaxWeight is smooth, while it can be non-smooth under other throughput-optimal policies. Specifically, we compare MaxWeight to a "CSMA" algorithm, which can be viewed as a crude approximation of 802.11 Medium Access Control (MAC) [3]. Also, in the simulations we use arrival processes consisting of discrete customers, as opposed to continuously distributed "amounts of work", which we assumed in our analysis for technical convenience.

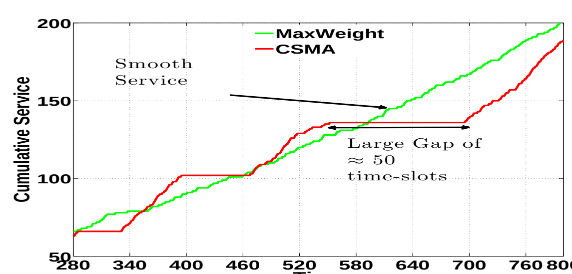

Fig. 3: Figure 2 is enlarged to show large gaps in the service process under CSMA.

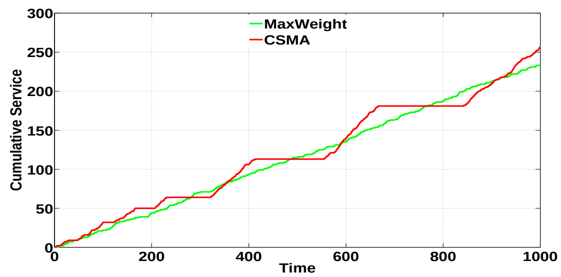

Fig. 4: $\lambda = .249$, $\mu = .5$, load $= .996$, $\gamma_i \equiv 1$.

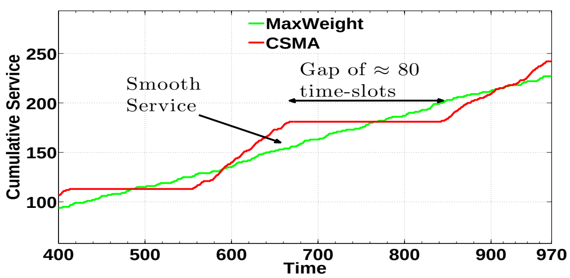

Fig. 5: Figure 4 is enlarged to show large gaps in the service process under CSMA.

Consider a network of 4 wireless links, shown as vertices $1, 2, 3, 4$ in Figure 1; an edge between two vertices indicates interference between the corresponding two links – they cannot be used for transmission simultaneously. Therefore, a scheduling policy can either simultaneously activate the links 1 and 2 (we label this scheduling choice as $(1, 2)$), or the links 3 and 4. Each link is associated with exactly one flow. The arrival process for each link is an independent Bernoulli process with mean $\lambda_l \equiv \lambda$; i.e. in each time slot one customer arrives with probability λ. The services are also Bernoulli: when a link l is scheduled, one customer is served with probability $\mu_l \equiv \mu$.

Let $\gamma_l = 1, l = 1, 2, 3, 4$ be the weights associated with the links $1, 2, 3, 4$. Then the MaxWeight policy chooses schedule $(1, 2)$ in time slot t if $\gamma_1 Q_1(t) + \gamma_2 Q_2(t) > \gamma_3 Q_3(t) + \gamma_4 Q_4(t)$, and schedule $(3, 4)$ otherwise.

The "CSMA" policy that we consider emulates the behavior of 802.11 Medium Access Control (MAC), operating in discrete time (equal size packets), and with a large maximum-backoff parameter. Under these conditions, and given the interference structure of the network, if say links 1 and 2 transmit packets simultaneously, then (under the 802.11 MAC) as long as both queues 1 and 2 are non-empty, the schedule $(1, 2)$ will persist for a long time. Our goal here is just to illustrate the fact that the service process can be "non-smooth" under some not unreasonable scheduling policies. That is why we pick a simple policy emulating the behavior of 802.11 MAC., as opposed to simulating the latter in detail. Specifically, we simulate the following "CSMA" policy. If the schedule $(1, 2)$ is used in time-slot t, then it will persist till time s when either $Q_1(s) = 0$ or $Q_2(s) = 0$. When this occurs, if both $Q_1(s) + Q_2(s) > 0$ and $Q_3(s) + Q_4(s) > 0$, one of the two schedules is chosen with equal probability $\frac{1}{2}$. If $Q_1(s) + Q_2(s) > 0$, but $Q_3(s) + Q_4(s) = 0$, schedule $(1, 2)$ is chosen. If, $Q_3(s) + Q_4(s) > 0$, but $Q_1(s) + Q_2(s) = 0$, schedule $(3, 4)$ is chosen. The scheduling rule is symmetric w.r.t. schedules $(1, 2)$ and $(3, 4)$. It is not hard to verify that this algorithm is throughput-optimal for our system.

Figures 2, 3, 4 and 5 summarize the results. We plot the cumulative amount of service obtained by flow 1 as a function of time under the MaxWeight and CSMA policies. As expected, the service process under MaxWeight remains smooth even as traffic intensity is close to 1, while the service process under CSMA algorithm has large gaps.

References

[1] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, P. Whiting. Providing Quality of Service over a Shared Wireless Link. *IEEE Communications Magazine*, 2001, Vol.39, No.2, pp.150-154.

[2] M. Andrews, K. Jung, A. L. Stolyar. Stability of the Max-Weight Routing and Scheduling Protocol in Dynamic Networks and at Critical Loads. STOC'07, San Diego, CA, June 11-13, 2007.

[3] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 2000, Vol. 18, No. 3, pp. 535-547.

[4] A. Eryilmaz and R. Srikant. Asymptotically Tight Steady-State Queue Length Bounds Implied by Drift Conditions. *Queueing Systems*, 2012, Vol.72, No.3-4, pp.311-359.

[5] B. Hajek. Hitting-Time and Occupation-Time Bounds Implied by Drift Analysis with Applications. *Advances in Applied Probability*, 1982, Vol.14, No.3, pp. 502-525.

[6] I-H. Hou, V. Borkar and P.R. Kumar. A Theory of QoS for Wireless. *INFOCOM 2009*.

[7] R. Li, A. Eryilmaz, B. Li. Throughput-Optimal Wireless Scheduling with Regulated Inter-Service Times. *INFOCOM 2013*.

[8] B. Li, R. Li, A. Eryilmaz. Heavy-Traffic-Optimal Scheduling with Regular Service Guarantees in Wireless Networks. *MobiHoc 2013*.

[9] S. P. Meyn and R.L. Tweedie. Stability of Markovian Processes I: Criteria for Discrete-Time chains. *Advances in Applied Probability*, 1992, Vol.24, No. 3, pp. 542-574.

[10] A.L. Stolyar. Max Weight Scheduling in a Generalized Switch: State Space Collapse and Workload Minimization in Heavy Traffic. *Annals of Applied Probability*, 2004, Vol.14, No. 1, pp. 1-53.

[11] L. Tassiulas, E. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Control*, 1992, Vol. 37, No. 1, pp. 1936 - 1948.

[12] R. Singh, A. Stolyar. MaxWeight Scheduling: Asymptotic Behavior of Unscaled Queue-Differentials in Heavy Traffic. http://arxiv.org/abs/1502.03793