

Fluctuation Analysis of Debt Based Policies for Wireless Networks with Hard Delay Constraints

Rahul Singh, I-Hong Hou, and P. R. Kumar

Department of Electrical and Computer Engineering, Texas A&M University

Email: {rsingh1,ihou,prk}@tamu.edu

Abstract—Hou et al. have analyzed wireless networks where clients served by an access point require a timely-throughput of packets to be delivered by hard per-packet deadlines and also proved the timely-throughput optimality of certain debt-based policies.

However, this is a weak notion of optimality; there might be long time intervals in which a client does not receive any packets, undesirable for real-time applications.

Motivated by this, the authors, in an earlier work, introduced a pathwise cost function based on the law of the iterated logarithm, studied in fluctuation theory, which captures the deviation from a steady stream of packet deliveries and showed that a debt-based policy is optimal if the frame length is one.

This work extends the analysis of debt-based policies to general frame lengths greater than one, as is important for general applications.

I. INTRODUCTION

We address the problem of access points (APs) serving clients that have hard per-packet deadline constraints on the packets delivered (see Figure 1). This problem is useful in applications such as video conferencing, voice over IP, networked control, and other cyberphysical systems where delay is critical.

Time is discretized by dividing it into slots of finite time duration. Slots numbered $k\tau, \dots, (k+1)\tau - 1, k = 0, 1, \dots$ form the k -th frame. A packet arrives for each client at the AP at the beginning of each frame, and has a relative deadline of τ time-slots from the time it is generated. That is, the packet generated for client i at the beginning of time-frame k becomes useless if it is not delivered to it by the time slot $(k+1)\tau - 1$, which is the end of time-frame k . Such a packet is said to have “expired”, and is never attempted by the AP again in future. We assume complete feedback from the clients to AP, i.e., the AP knows correctly whether the packet transmitted by it to client i in time-slot t was delivered successfully or not before the beginning of time-slot $t+1$. (The details on mechanisms whereby this is achieved can be found in [1]; it is done via an ACK from the client to the AP,

with the complete two-phase DATA-ACK treated as a single transaction. Hence the duration of a time-slot is large enough to accommodate a transmission from AP to a client and an acknowledgement feedback from client to AP.)

In each time slot the AP can transmit only one packet, and thus serve only one client. To model the unreliable wireless channel between the AP and clients, we assume that each time the AP transmits a packet to client i , the transmission is successful with a probability p_i . In case the transmission of a packet is a failure, the AP can try to retransmit that packet at a later time-slot if it has not expired yet.

A. Timely Throughput

It is required that each client i receives a minimum “fraction” of the packets it generates. We will term this required fraction as the *timely-throughput requirement of client i* and denote it by q_i . More precisely, the timely-throughput requirement of client i is met if

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{\text{client } i\text{'s packet is delivered in frame } t\} \geq q_i. \quad (1)$$

If it is possible for the AP to meet the timely-throughput requirement of each client, then we say that the set of clients $\{1, \dots, N\}$ is *feasible*.

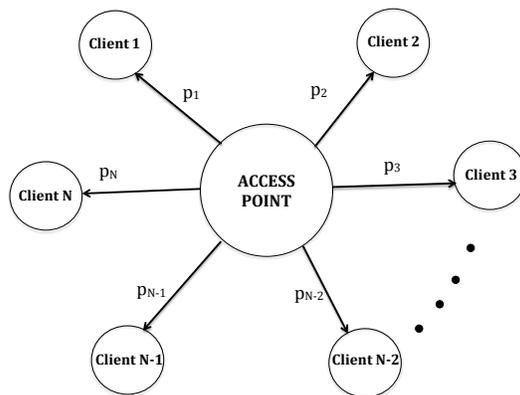


Fig. 1: An access point serving N real-time flows.

This material is based upon work partially supported by the AFOSR under Contract No. FA9550-13-1-0008, NSF under Contract Nos. CNS-1302182, CPS-1239116, CPS-1232601, CPS-1232602, Science & Technology Center Grant CCF-0939370, and USARO under Contract W911NF-08-1-0238.

B. Necessary Conditions for Feasibility

The set of timely throughput vectors (q_1, \dots, q_N) that are feasible is characterized in [1]. We briefly describe this. From the discussion of the system model above, it follows that the time taken to deliver a packet of client i generated in frame s is geometrically distributed with parameter p_i . Denote this random time by γ_i . Clearly if $\sum_{i=1}^N \gamma_i < \tau$, then the AP will have delivered all the packets generated in the frame s before the completion of frame s , and, since it has no more packets to attempt, it is forced to be idle. Thus $\tau - \sum_{i=1}^N \gamma_i$ number of time-slots would have effectively been “wasted”. If the AP only has to serve a subset S of clients $\{1, \dots, N\}$, then

$$I_S \triangleq \frac{1}{\tau} E \left(\left[\tau - \sum_{i=1}^N \gamma_i \right]^+ \right), \text{ where } x^+ = \max(x, 0), \quad (2)$$

is the expected fraction of time in each frame which is wasted by the AP as idle time. We note that the above quantity is independent of the scheduling policy implemented at the AP, as long as it is non-idling. In [1] it is shown that a set of clients is feasible if and only if

$$\sum_{i \in S} \frac{q_i}{\tau p_i} \leq 1 - I_S, \forall S \subseteq \{1, \dots, N\}. \quad (3)$$

Henceforth we will call this set (3) the “rate-region”, and a scheduling policy which satisfies any set of clients having the vector of timely throughput requirements lying in the rate-region is said to be a “feasibility optimal” policy.

C. Maximum Debt First Policy

In [1], a class of policies called debt-based policies is introduced and is shown to satisfy any set of clients whose timely-throughput requirement vector lies in the rate-region. The weighted maximum debt first (MDF) policy keeps track of the “delivery-based debt” of each client, where the debt of client i at the beginning of time-frame t is defined as

$$d_i(t) := q_i t - \text{number of packets delivered for client } i \text{ in the frames } \{0, \dots, t-1\}. \quad (4)$$

At the beginning of time-frame t , the MDF policy arranges the clients in the decreasing order of weighted-debts $\frac{d_i(t)}{p_i}$ and then serves the clients in that order during the frame t . Subsequent works such as [2] have generalized the results to the scenario where the clients have different deadlines and the case where the channel reliabilities are time-varying.

D. Problem Motivation

Note that the timely throughput requirements only captures the long term average throughput, i.e.,

$$\liminf_{T \rightarrow \infty} \frac{\text{Number of packets of user delivered till frame } T}{T}.$$

This is however a loose requirement in the following sense. It might be the case that there are long time intervals during which a client i does not receive any packets, even though the relation (1) is satisfied. Thus there is a need to pay attention to the short term behaviour of the debts, which is not captured in the definition of timely throughput.

We address a stronger sense of performance in this paper. It can be regarded as the analog of the difference between the law of large numbers and the central limit theorem. Or, more precisely, the difference between throughput optimality and low delay, since several policies can be throughput optimal while still having very poor delay performance.

E. Fluctuation Analysis Approach

We consider a precise pathwise notion of performance. We study the debt process by $\phi(t) := \sqrt{2t \log \log t}$, and study the almost sure limit of the scaled process $\frac{d_i(t)}{\sqrt{2t \log \log t}}$. This is similar to the studies done in fluctuation theory [3]. A survey of results of fluctuation theory can be found in [3].

The above is similar to deriving characterizations for the “workload process” associated with a queueing system in the heavy traffic regime via fluid limits or diffusive scaling. The difference is that in our case the system is taken to be exactly in heavy traffic, and thus we do not need to consider a sequence of systems whose mean arrivals approach the maximum possible mean departure. Also, our characterization by the law of the iterated logarithm (LIL) is a precise sample path characterization of the performance in heavy traffic.

The max weight scheduler has been shown to be “optimal” in the “heavy traffic regime” in [4], [5]. In [4], results are derived for the diffusive scaled “workload processes”. Reference [5] shows that Max Weight scheduler is optimal with respect to the cost defined as the sum of expected (average) queue lengths in steady state multiplied by $(1 - \rho)$, where ρ is the traffic intensity of the queueing system.

We note that the optimality with respect to average queue lengths doesn’t guarantee the absence of fluctuations in the queue lengths, or, as in our case, the debt vector. For example the infinite sequences $(10, 0, 10, 0, \dots)$ and $(5, 5, \dots)$ have the same average, but the first sequence has a peak value twice that of the second. This raises the important question as to whether there exists a well-defined envelope containing the queue lengths, i.e., a non-random function upper-bounding the queue lengths, with moreover the queue lengths “touching” this envelope infinitely often. Our analysis in this paper shows that such a well-defined envelope does indeed exist for the queue lengths associated with a system operating under the MDF policy, and describes methods to calculate it. We also note that on the other hand optimality with respect to this envelope doesn’t guarantee optimality with respect to average queue lengths.

We also note in passing that our problem differs from the general switch considered in [4], [5] in that the debt processes can become negative, unlike queue lengths.

II. ASSUMPTIONS AND RELEVANT PRIOR RESULTS.

For a system operating under the MDF policy, the debt vector at time t will be denoted by $\mathbf{d}(t) \triangleq (d_1(t), \dots, d_N(t))$. Throughout vectors will be denoted in bold. Similarly the vectors of timely-throughput requirements and channel reliabilities will be denoted by $\mathbf{q} \triangleq (q_1, \dots, q_N)$ and $\mathbf{p} \triangleq (p_1, \dots, p_N)$, respectively. We consider a timely throughput vector \mathbf{q} that lies on exactly one of the hyperplanes describing the rate-region. More precisely \mathbf{q} is such that,

$$\begin{aligned} \sum_{j \in S} \frac{q_j}{p_j} < \tau - \tau I_S, \text{ where } S \subsetneq \{1, \dots, N\}, \text{ and,} \\ \sum_{j=1}^N \frac{q_j}{p_j} = \tau - \tau I_{\{1, \dots, N\}}; \text{ this is the heavy traffic condition.} \end{aligned} \quad (5)$$

As a technical assumption, we will suppose that \mathbf{q} and \mathbf{p} are rational, i.e., all their entries are rational. Under the above assumptions the debt vector process for a system operating under the MDF policy is a Markov process with a countable state space \mathbb{Q}^N . The i -th entry of the debt vector process, i.e., the debt of client i , evolves as

$$\begin{aligned} d_i(t+1) &= d_i(t) + q_i \\ &\quad - \mathbb{1}\{\text{packet for user } i \text{ is delivered in frame } t\}, \end{aligned} \quad (6)$$

with $d_i(0) \equiv 0$ for all $i = 1, \dots, N$.

Since the MDF policy fixes an order at the beginning of each frame, and serves the clients in that order, we will index the $N!$ possible permutations and then use the term ‘‘order k ’’ to mean that the MDF policy serves the clients in the order described by the k -th permutation.

We now state a few results from [1], which will be useful in later sections. Under the MDF policy, for a vector of timely throughput that satisfies (5), we have

$$\limsup_{t \rightarrow \infty} \frac{d_j(t)}{t} = 0. \quad (7)$$

We note the following. For a fixed frame t , if we renumber the clients so that under the MDF policy they are served in order $1, 2, \dots, N$, then if π_j denotes the probability of user j receiving packet successfully in frame t , we have

$$\sum_{j=1}^n \frac{\pi_j}{p_j} = \tau - \tau I_{\{1, \dots, n\}}, \text{ where } n = 1, \dots, N. \quad (8)$$

III. NOTATIONS

Throughout, we denote vectors in boldfont, for example $\mathbf{a} \triangleq (a_1, \dots, a_N)$. Let $W_i(t) := \frac{d_i(t)}{p_i}$, which we will refer to as the ‘‘workload’’ for user i , with $\sum_{i=1}^N W_i(t)$ the total workload at time t . The vector obtained by arranging the workloads $W_i(t), i = 1, \dots, N$ in decreasing order will be denoted by $\mathbf{W}^0(t)$ with entries $W_1^0(t), \dots, W_N^0(t)$.

Also $\frac{W_j(t)}{\phi(t)}$ will be denoted as $\tilde{W}_j(t)$ and called the scaled workload for user j . The scalar $\frac{\sum_{j=1}^N W_j(t)}{\phi(t)}$ will be

denoted as $\tilde{W}(t)$ and called the scaled total workload, while $W(t) := \sum_{j=1}^N W_j(t)$ will be called total workload.

Stochastic processes will be mentioned without time-index, e.g., W refers to $W(t)$.

For two random variables Y and Z , $Y \leq_{st} Z$ or equivalently Z is stochastically greater than Y would mean that $P(Y \geq x) \leq P(Z \geq x)$, [6].

IV. THE NOTION OF PATHWISE OPTIMALITY.

We now introduce an appropriate notion of optimality for the heavy traffic regime. We will restrict ourselves to policies satisfying $\limsup_{t \rightarrow \infty} \tilde{W}_i(t) < +\infty$ almost surely (a.s.) for all i . If this is the case then clearly $\limsup_{t \rightarrow \infty} \max_{i=1, \dots, N} \tilde{W}_i(t) < +\infty$ a.s.. The class of policies having this property will be denoted by \mathcal{P} . Note that for a policy $\Pi \in \mathcal{P}$, the quantity $\limsup_{t \rightarrow \infty} \max_{i=1, \dots, N} \tilde{W}_i(t)$ can be random. If a policy $\Pi^* \in \mathcal{P}$ is such that

$$\limsup_{t \rightarrow \infty} \max_{i=1, \dots, N} \tilde{W}_i^{\Pi^*}(t) \leq_{st} \limsup_{t \rightarrow \infty} \max_{i=1, \dots, N} \tilde{W}_i^{\Pi}(t), \forall \Pi \in \mathcal{P}, \quad (9)$$

where the superscript denotes the dependence on the policy, then we say that the policy Π^* is optimal. Hence to any policy $\Pi \in \mathcal{P}$ we have attached a cost

$$\limsup_{t \rightarrow \infty} \max_{i=1, \dots, N} \tilde{W}_i^{\Pi}(t). \quad (10)$$

Note that since the relation \leq_{st} is a partial order, it is not clear from the definition (9) that there exists an optimal policy. Next we define another notion of optimality.

If a policy $\Pi^* \in \mathcal{P}$ is such that

$$\limsup_{t \rightarrow \infty} \tilde{W}^{\Pi^*}(t) \leq_{st} \limsup_{t \rightarrow \infty} \tilde{W}^{\Pi}(t), \quad \forall \Pi \in \mathcal{P}, \quad (11)$$

then we say that the policy Π^* is optimal with respect to the limiting total scaled workload. Once again there might not exist a policy which is optimal as in (11). However, we show in later sections that indeed there exists a policy which is optimal as in (11).

We will consider the problem of minimizing the above cost (9). We show in the sequel how minimization as in (9) yields a non-random constant. Minimizing the performance cost defined in (10) will also entail guaranteeing a sort of ‘‘fairness’’ amongst the different clients, i.e., the fluctuations in the total workload process will be spread out evenly across all the clients.

V. AN OVERVIEW OF THE APPROACH

A cost similar to the one defined in equation (10) was introduced in [7], and the MDF policy was shown to be optimal with respect to it. However the analysis was stringently restricted to the case of $\tau = 1$, i.e., it was assumed that the system defined in Section I is operating with a frame length of one time slot. In this work we analyze the important general case of $\tau > 1$. We now briefly describe the approach taken in this work. One may note that the cost in (10) can be random. However it will be shown in later sections that for a considered class of

policies, the limiting scaled total workload is a constant a.s., and that the MDF policy lies in this class. Assuming that we have shown this fact, one may observe that for the MDF policy to optimize the cost described in equation (10) it suffices to

- 1) Minimize $\limsup_{t \rightarrow \infty} \sum_{i=1}^N \tilde{W}_i(t)$, and, meanwhile
- 2) Keep the limiting scaled workloads equal, i.e. $\lim_{t \rightarrow \infty} |\tilde{W}_i(t) - \tilde{W}_j(t)| = 0$ for all i, j .

This is the approach we take in analyzing the MDF policy with respect to the defined cost function. It turns out that it is relatively easy to show the second of the two conditions mentioned above. We define the notions of “weak” and “strong” state space collapse (SSC) and show that if under a policy strong SSC occurs then the condition 2) is true. We show that for the system operating under an MDF policy, strong SSC does indeed occur. From this we will conclude that condition 2) is true for the MDF policy.

We then proceed to check the conjecture that 1) is true for the MDF policy. We provide an expression for $\limsup_{t \rightarrow \infty} \sum_{i=1}^N \tilde{W}_i(t)$ for a system operating under the MDF policy. We briefly describe the class of policies mentioned that has a constant nonrandom limit for the total scaled workload. We find that if weak SSC occurs under a policy $\Pi \in \mathcal{P}$, then $\limsup_{t \rightarrow \infty} \sum_{i=1}^N \tilde{W}_i(t)$ is a constant a.s. for a system operating under this policy. Since strong SSC would imply weak SSC, for the MDF policy, the quantities $\limsup_{t \rightarrow \infty} \tilde{W}_i(t)$ are non-random almost surely and independent of i . Section VI derives the state space collapse property of a system operating under the MDF policy. Section VII discusses some key properties of martingales, which are used in Section VIII to derive the expression for the limiting total scaled workload. Finally Section X derives some easily computable bounds on the scaled limiting workloads.

VI. STATE SPACE COLLAPSE

State space collapse for a general switch operating under the Max Weight scheduler has been shown to occur in [4], [5]. However, we note that the debt vector process in our case is different from the vector-valued “queue length process” encountered while studying the general switch, which has queue lengths as its components. The components of debt vector process can become negative, which is not the case with the queue length process for a general switch. We use techniques similar to [5] to show state space collapse in our case. The formal definition of state space collapse is given in subsection VI-B. The main result of this section can be summarized as saying that for the MDF policy, the limiting scaled workloads for various clients are the same, i.e., $\tilde{W}_i(t)$ is independent of i modulo $o(\phi(t))$.

The first subsection (VI-A) introduces two lemmas which are crucial in showing the main result, which is introduced in the second subsection (VI-B).

A. Preliminaries

The following Lemma is taken from [5], [8].

Lemma 1: For an irreducible and aperiodic Markov Chain $\{X(t)\}_{t \geq 0}$ over a countable state space \mathcal{X} , suppose $Z : \mathcal{X} \rightarrow \mathbb{R}_+$ is a non-negative valued function. Define the drift of Z at X as

$$\Delta Z(X) \triangleq [Z(X(t+1)) - Z(X(t))] \mathbb{1}\{X(t) = X\}. \quad (12)$$

Let us suppose that the drift satisfies the following conditions:

- Condition \mathcal{C}_1 : There exists an $\eta > 0$, and a $\kappa < \infty$ such that

$$E[\Delta Z(X) | X(t) = X] \leq -\eta, \forall X \in \mathcal{X} \text{ with } Z(X) \geq \kappa.$$

- Condition \mathcal{C}_2 : There exists a $D < \infty$ such that

$$P(|\Delta Z(X)| \leq D) = 1, \quad \forall X \in \mathcal{X}.$$

Then there exists a $\theta^* > 0$ and a $C^* < \infty$ such that

$$\limsup_{t \rightarrow \infty} E[\exp^{\theta^* Z(X[t])}] \leq C^*. \quad (13)$$

Also, the Markov Chain $X(t)$ is positive recurrent and the process $Z(t) \triangleq Z(X(t))$ converges in distribution to a random variable Z_∞ for which

$$E[\exp^{\theta^* Z_\infty}] \leq C^*,$$

which implies that all moments of Z_∞ exist and are finite. Moreover the process $Z(t)$ satisfies

$$P(Z(t) \geq b | \mathcal{F}_0) \leq \rho^t \exp(\zeta(Z(0) - b)) + \frac{1 - \rho^t}{1 - \rho} D \exp(\zeta(\kappa - b)), \quad (14)$$

where $0 < \rho < 1$ and $\zeta < 0$. Hence if X_∞ denotes the version of the process $X(t)$ distributed according to the stationary distribution, then

$$P(Z_\infty \geq b) \leq \frac{D}{1 - \rho} \exp(\zeta(b - a)). \quad (15)$$

If a process $Y(t)$ satisfies the condition (14), then we will say that $Y(t)$ is **stochastically bounded** (by an exponential random variable, modulo transient behaviour).

Lemma 2: If a random process $X(t)$ is stochastically bounded and has bounded increments, then $\lim_{t \rightarrow \infty} \frac{|X(t)|}{\phi(t)} = 0$.

Proof: We notice that if inequality (14) is satisfied, then it implies inequality (13). We will show that $\sum_{t=1}^{\infty} P\left(\frac{|X(t)|}{\phi(t)} > \delta\right) < \infty$, from which the result follows via Borel Cantelli Lemma.

Since the process $X(t)$ is stochastically bounded, for every $\epsilon > 0$, there exists an $N(\epsilon)$ such that $E(|X(t)|) < C^* + \epsilon, \forall t > N(\epsilon)$. Moreover since $X(t)$ has bounded increments, $|X(t)| < B, \forall t \leq N(\epsilon)$ for some $0 < B < \infty$. Combining these two, we have

$$\sum_{t=1}^{\infty} P\left(\frac{|X(t)|}{\phi(t)} > \delta\right) \leq \sum_{t=1}^{N(\epsilon)} \frac{B\delta}{\phi(t)}$$

$$\begin{aligned}
& + \sum_{t=N(\epsilon)+1}^{\infty} P(\exp(\theta^*|X(t)|) > \exp(\phi(t)\delta)) \\
& \leq \sum_{t=1}^{N(\epsilon)} \frac{B\delta}{\phi(t)} + \sum_{t=N(\epsilon)+1}^{\infty} \frac{C^* + \epsilon}{\exp(\delta\sqrt{t})}, \\
& < \infty,
\end{aligned}$$

where the inequalities result from Markov's inequality, and the fact that $\sqrt{t} < \phi(t)$ for large t . ■

B. Drift Analysis

We begin with some definitions. Since we want to show that the limiting scaled workloads for various clients are the same for all clients, we define a vector-valued process

$$\mathcal{W} := (W_2 - W_1, \dots, W_N - W_1). \quad (16)$$

Notice that since the MDF policy decides the schedule based on the order of workloads of users, W_1, \dots, W_N , the process \mathcal{W} is Markov with components evolving as

$$\begin{aligned}
\mathcal{W}_i(t+1) &= \mathcal{W}_i(t) + \frac{q_i}{p_i} - \frac{q_1}{p_1} \\
&\quad - \frac{\mathbb{1}\{\text{packet of user } i \text{ is delivered in frame } t\}}{p_i} \\
&\quad + \frac{\mathbb{1}\{\text{packet of user 1 is delivered in frame } t\}}{p_1}.
\end{aligned}$$

Definition 1: 1) Weak State Space Collapse (Weak SSC): If the Markov process \mathcal{W} associated with a system described in Section I operating under a policy Π is positive recurrent, then we say that the scheduling policy Π induces weak SSC.

2) Strong State Space Collapse (Strong SSC). : If the Markov process \mathcal{W} associated with the system described in Section I operating under a policy Π is such that each component of \mathcal{W} is stochastically bounded, i.e., the processes $W_i - W_1$ are stochastically bounded for all $i = 2, \dots, N$, then we say that the scheduling policy Π induces Strong SSC.

Note that strong SSC implies weak SSC. This is true since the Markov chains discussed here evolve on a countable state space.

For inferring that the MDF policy induces SSC, instead of studying the process \mathcal{W} , we will analyze the process $\mathbf{W}^0(t)$. The difference $W_i^0 - W_{i+1}^0$ being difficult to analyze, we introduce the notion of ‘‘centre of mass’’ of a set,

$$CM_S(\mathbf{d}) = \frac{1}{|S|} \sum_{l \in S} W_l^0, \text{ where } S \subseteq \{1, \dots, N\}. \quad (17)$$

Note that

$$CM_{\{1, \dots, n\}}(\mathbf{d}) - CM_{\{n+1, \dots, N\}}(\mathbf{d}) \geq W_n^0 - W_{n+1}^0. \quad (18)$$

It is easy to show that these distances between CM's are stochastically bounded and hence we will use these distances as Lyapunov functions. The next Theorem states that the MDF policy induces strong SSC. The proof is given in Appendix, Section XII-A.

Theorem 1: Consider the system described in Section I operating under the MDF policy. Then $W_i - W_j$ is stochastically bounded for all i, j and hence $\lim_{t \rightarrow \infty} \bar{W}_i$ is independent of i . The MDF policy induces strong SSC.

Remark 1: Following the approach above, it can also be shown that under the time-based debt policy, we have SSC for the time-based debt vector, i.e., $d_i - d_j$ is stochastically bounded, where d_i is the time-based debt of user i .

After showing that the workloads of all users are the same modulo $o(\phi(t))$, we would like to analyze the total workload of the system. Since the workloads of all users are the same, dividing it by N would yield the workload of an individual user. The remaining sections discuss various approaches for studying the total workload.

VII. SOME USEFUL RESULTS ON MARTINGALES

Let $(Y_n, \mathcal{F}_n)_{n \geq 0}$ be a Martingale with $X_n = Y_n - Y_{n-1}$ and $s_n^2 = \sum_{i=1}^n E[X_i^2 | \mathcal{F}_{i-1}]$. Also let $u_n = \sqrt{2 \log \log s_n^2}$. The following is the Law of Large Numbers for martingales, [9]:

Theorem 2: (LLN) Let $Y_n = \sum_{i=1}^n X_i$ be a Martingale such that $\sum_{n=1}^{\infty} \frac{E(X_n^2)}{n^2} < \infty$. Then $\frac{Y_n}{n} \rightarrow 0$ a.s.

The following is the law of iterated logarithm for martingales [10]:

Theorem 3: (LIL) If $s_n^2 \rightarrow \infty$ and $|X_n| \leq \frac{K_n s_n}{u_n}$, where K_n are \mathcal{F}_{n-1} measurable with $K_n \rightarrow 0$, then

$$\limsup_{n \rightarrow \infty} \frac{Y_n}{s_n u_n} = 1. \quad (19)$$

VIII. LIMIT THEOREM FOR TOTAL WORKLOAD

This section firstly introduces a martingale, then relates it to the total workload of the system. Finally the LIL for martingales is used to study the total workload scaling. Consider the martingale difference sequence

$$\begin{aligned}
m(s) &= \sum_{j=1}^N [\text{number of slots given to user } j \text{ in frame } s \\
&\quad - \frac{\mathbb{1}\{\text{a packet of user } j \text{ is delivered in frame } s\}}{p_j}] \\
&\quad + i(s) - \tau I_{1, \dots, N},
\end{aligned} \quad (20)$$

where $i(s)$ denotes the idle time in the frame s . Note that the total time-slots given to all users and the idle time add to τ . Hence combining the first and third terms in the expression on the RHS of equation (20), we can rewrite (20) as

$$\begin{aligned}
m(s) &= \tau - \tau I_{\{1, \dots, N\}} \\
&\quad - \sum_{j=1}^N \frac{\mathbb{1}\{\text{a packet of user } j \text{ is delivered in frame } s\}}{p_j}
\end{aligned}$$

$$= \sum_{j=1}^N \left(\frac{q_j}{p_j} - \frac{\mathbb{1}\{\text{a packet of user } j \text{ is delivered in frame } s\}}{p_j} \right), \quad (21)$$

where the last equality follows from (5).

Summing (21) over s varying from 1 to t yields,

$$M(t) = \sum_{j=1}^N W_j(t), \quad (22)$$

where $M(t)$ is a martingale with the difference sequence $m(s)$. Let

$$\alpha_k(t) = \frac{\sum_{s=1}^t \mathbb{1}\{\text{order } k \text{ is followed in frame } s\}}{t}, \quad (23)$$

which denotes the fraction of time order k is followed till time t . Application of LIL to the martingale $M(t)$ gives the following:

Theorem 4:

$$\limsup_{t \rightarrow \infty} \tilde{W}(t) = \lim_{t \rightarrow \infty} \left(\sum_{k=1}^{N!} \alpha_k(t) V_k \right)^{\frac{1}{2}}, \quad (24)$$

where $V_k = E[m^2(t) | \text{order } k]$. Hence the RHS has a limit a.s.

Proof: Consider the martingale $M(t)$ with the difference sequence $m(s)$. Note that $m(s)$ defined in equation (20) is bounded. Thus we apply Theorem 3 by setting $K_t = \frac{u_t}{s_t}$, where $s_t^2 = \sum_{k=1}^{N!} \alpha_k(t) V_k$ and $u_t = \sqrt{2 \log \log s_t^2}$. Clearly $K_t \rightarrow 0$, and since $\frac{M(t)}{\phi(t)} = \tilde{W}(t)$, this gives us

$$\limsup_{t \rightarrow \infty} \tilde{W}(t) = \lim_{t \rightarrow \infty} s_t. \quad \blacksquare$$

Remark 2: Though the above theorem yields the expression for the limiting scaled total workload, it is not clear whether the expression in the RHS of (24) is a constant a.s.. Next we show that if the scheduling policy is such that it induces weak SSC, then the limit in the expression (24) is indeed a constant.

IX. WORKLOAD UNDER SSC

We will consider the ‘‘scheduling process,’’ by which we mean a process which at time t denotes the order being followed by the MDF scheduler in the frame t . Notice that because the MDF policy decides the schedule based on the order of workloads of users W_1, \dots, W_N , one can equivalently describe the scheduling process by the vector process \mathcal{W} . Since the limiting scaled workload in Theorem 4 depends only on the scheduling process, it suffices to study the process \mathcal{W} in order to study the limit in the expression (24).

Now assume the process \mathcal{W} is positive recurrent and let μ be the corresponding stationary measure. Since the

V_k 's as defined in Theorem 4 are bounded, using the ergodic theorem for Markov processes,

$$\lim_{t \rightarrow \infty} \left(\sum_{k=1}^{N!} \alpha_k(t) V_k \right)^{\frac{1}{2}} = \left(\sum_{k=1}^{N!} \mu_k V_k \right)^{\frac{1}{2}}, \quad \text{where} \\ \mu_k \triangleq \mu(\{\mathbf{d} : \text{order } k \text{ is followed} \\ \text{when the debt vector is } \mathbf{d}\}). \quad (25)$$

This gives the following.

Theorem 5: Consider the Markov process \mathcal{W} associated with the system described in Section I evolving under a policy Π such that weak SSC occurs. Then the process \mathcal{W} is ergodic and hence $\limsup_{t \rightarrow \infty} \tilde{W}(t)$ is constant and given by

$$\left(\sum_{k=1}^{N!} \mu_k V_k \right)^{\frac{1}{2}}. \quad (26)$$

Corollary 1: Since the MDF policy induces strong SSC on the system described in Section I, $\limsup_{t \rightarrow \infty} \tilde{W}(t)$ is constant.

In the next few sections, we derive bounds on the limiting workloads via several techniques.

X. BOUNDS ON WORKLOADS.

A. Bounding the Workload via a Linear Program.

Consider the martingale difference sequence,

$$m_j^k(s) := \frac{\mathbb{1}\{\text{order } k \text{ is followed in frame } s, \text{ packet of } j \text{ is delivered}\}}{\pi_j^k}, \quad (27)$$

with π_j^k denoting the probability of delivering packet of user j under the condition that order k is followed. The following Lemma is useful in deriving bounds on the workload and its proof is given in Appendix, Section XII-B.

Lemma 3: For any non-idling feasibility optimal policy,

$$\lim_{t \rightarrow \infty} \sum_k \pi_j^k \alpha_k(t) = q_j, \quad (28)$$

where $\alpha_k(t)$ are as defined in equation (23).

Lemma 3 when applied with Theorem 5 yields the following result regarding the total workload scaling in the system:

Theorem 6: Consider the linear program,

$$\min \sum_{k=1}^{N!} x_k V_k \quad \text{such that} \\ \mathbf{x}A = \mathbf{q} \quad \text{and} \quad \sum_{i=1}^{N!} x_i = 1. \quad (29)$$

where $\mathbf{x} = (x_1, \dots, x_{N!})$ and $A_{k,j} = \pi_j^k$. Let $\mathbf{x}^* \triangleq (x_1^*, \dots, x_{N!}^*)$ and c^* denote the solution and the optimal

value of the linear program, respectively. For any throughput optimal policy, the total limiting scaled workload, i.e., $\limsup_{t \rightarrow \infty} \tilde{W}(t)$ is lower bounded by $\sqrt{c^*}$. Similarly the cost (10) for any policy in \mathcal{P} is lower bounded by $\frac{\sqrt{c^*}}{N}$.

The scheduling policy which schedules clients by time-sharing amongst the $N!$ orders, with the k -th order being employed for a fraction x_k^* of time, minimizes the limiting scaled total workload. Moreover replacing \min by \max in the above program yields an upperbound on the total workload under any throughput-optimal policy.

Proof: Applying Lemma 3 in conjunction with Theorem 4 and Theorem 5 yields the result. ■

Since we have already proved the SSC property for the MDF policy, the above result provides bounds on the workloads of each user.

Corollary 2: The limiting scaled workloads $\tilde{W}_i, i = 1, \dots, N$ for the system described in Section I and operating under the MDF policy are bounded by $\frac{\sqrt{c^*}}{N}$.

Though the LP gives upper and lower bounds, and, indirectly, the policy which is optimal with respect to total workload, it is computationally expensive. In the remainder of the paper, we derive some easily computable bounds on the workload.

B. Comparing Workloads via Coupling.

Instead of writing the difference sequence in (20) in terms of events in a frame, we can write it in terms of events in a time-slot. This gives,

$$M(t) = \sum_{j=1}^N \sum_{s=1}^{t \times \tau} [\mathbb{1}\{\text{user } j \text{ is attempted in slot } s\} - \frac{\mathbb{1}\{\text{packet of user } j \text{ is delivered in slot } s\}}{p_j}] + \sum_{l=1}^t i(l) - \tau I_{\{1, \dots, N\}}.$$

Combining the above with equations (20) and (22) gives

$$\sum_{j=1}^N W_j(t) = \tilde{M}(t \times \tau) + \sum_{l=1}^t i(l) - t\tau I_{\{1, \dots, N\}}, \quad (30)$$

where $\tilde{M}(t)$ is the martingale with difference sequence as

$$\begin{aligned} \tilde{m}(s) &= \sum_{j=1}^N \mathbb{1}\{\text{packet of user } j \text{ is attempted in slot } s\} \\ &\quad - \frac{\mathbb{1}\{\text{packet of user } j \text{ is delivered in slot } s\}}{p_j} \\ &\triangleq \sum_{j=1}^N \tilde{m}_j(s). \end{aligned} \quad (31)$$

Note that in equation (30), we have decomposed the total workload into the sum of two martingales. Now we can apply the LIL for the two martingales to infer an upper bound on the total workload.

Theorem 7: For any throughput optimal non-idling policy, we have

$$\limsup_{t \rightarrow \infty} \tilde{W}(t) \leq \sum_{j=1}^N \frac{q_j}{p_j} v_j + \sigma_I, \quad (32)$$

where $v_j = \frac{1-p_j}{p_j}$ is the conditional variance $E[\tilde{m}_j^2(s) | \text{user } j \text{ is attempted in slot } s]$, and σ_I^2 is the variance of the idle-time random variable.

Proof: Using equation (30) we have

$$\begin{aligned} \limsup_{t \rightarrow \infty} \tilde{W}(t) &\leq \limsup_{t \rightarrow \infty} \frac{\tilde{M}(t \times \tau)}{\phi(t)} \\ &\quad + \limsup_{t \rightarrow \infty} \frac{\sum_{l=1}^t (i(l) - \tau I_{\{1, \dots, N\}})}{\phi(t)}. \end{aligned}$$

The second term in the RHS of the above inequality is the same as σ_I by Kolmogorov's LIL [11], while the first term is the same as $\sum_{j=1}^N \frac{q_j}{p_j} v_j$ by using LIL for martingale $\tilde{M}(t)$ and noting that

$$\lim_{t \rightarrow \infty} \frac{\sum_{s=1}^t \mathbb{1}\{\text{user } j \text{ is attempted in slot } s\}}{t} = \frac{q_j}{p_j}. \quad \blacksquare$$

Also notice that the representation (30) can be used to compare the difference between the scaled total workloads of two scheduling policies if we assume that the idle time in frame t , $i(t)$, is the same for all non-idling policies. This can be achieved if we compare the policies on a common probability space which is equipped with the processing times of each user's packet in a frame. If $\gamma_j(t), t = 1, 2, \dots$, is the number of time-slots required to deliver a packet of user j in frame t , then we can assume that while comparing the performance of the two policies, the time taken to process client j in frame t is the same for both the policies and is equal to $\gamma_j(t)$. Since the idle time in frame t for any non-idling policy is equal to $(\tau - \sum_j \gamma_j(t))^+$, it is the same for any two policies being compared. This gives us

Theorem 8: Let $\tilde{W}_1(t)$ and $\tilde{W}_2(t)$ be the scaled workloads of two feasibility optimal, non-idling policies Π_1 and Π_2 . Then

$$\limsup_{t \rightarrow \infty} (\tilde{W}_1(t) - \tilde{W}_2(t)) \leq 2 \sum_{j=1}^N \frac{q_j}{p_j} v_j.$$

Proof: Denote by \tilde{W}^1 and \tilde{W}^2 the scaled workloads of the two policies that we wish to compare. Using the probability space described above, and using the representation (30), the scaled workloads of any two policies differ by only the first term in the RHS of (30), i.e.,

$$W^1(t) - W^2(t) = \tilde{M}^1(t\tau) - \tilde{M}^2(t\tau),$$

where the superscript on the RHS indexes the policies. Scaling the above by $\frac{1}{\phi(t)}$ and using $\limsup \sum \leq \sum \limsup$ yields the desired result. ■

C. Direct Computation of Variations.

Next, we provide an approach for analyzing the workload that relies on using the relations (20) and (22) relating the total workload to the martingale difference sequence $m(s)$. The order $k = 1$ will denote the order of serving the clients in the order of increasing channel reliabilities, i.e., if after reindexing the clients we have $p_1 \leq p_2 \leq \dots \leq p_N$, then the order $k = 1$ indicates that the clients are served in the order $1, 2, \dots, N$. Similarly the order $k = N!$ refers to serving the clients in decreasing order of channel reliabilities.

Theorem 9: For the system described in Section I evolving under the MDF policy, $\limsup_{t \rightarrow \infty} \bar{W}(t)$ is bounded above by

$$-(\tau - \tau I_{1,\dots,N})^2 + \sum_j \left(\frac{1}{p_j} + 2 \left(\sum_{l < j} \frac{1}{p_l} \right) \right) \frac{\pi_j^1}{p_j},$$

and lower bounded by

$$-(\tau - \tau I_{1,\dots,N})^2 + \sum_j \left(\frac{1}{p_j} + 2 \left(\sum_{l < j} \frac{1}{p_l} \right) \right) \frac{\pi_j^{N!}}{p_j}.$$

Proof: We will assume that the clients are renumbered so that under the order k clients are served in the order $1, 2, \dots, N$. We will use the martingale representation of (20). Note that,

$$\begin{aligned} & E [m^2(t) \mid \text{order } k] \\ &= E \left[\left(\tau - \tau I_{1,\dots,N} - \sum_j \frac{\mathbb{1}\{\text{delivery for } j\}}{p_j} \right)^2 \mid \text{order } k \right]. \end{aligned} \quad (33)$$

The expression within the expectation reduces to

$$\begin{aligned} & (\tau - \tau I_{1,\dots,N})^2 + \sum_j \frac{\pi_j^k}{p_j^2} \\ & + 2 \sum_j \left(\sum_{l < j} \frac{1}{p_l} \right) \frac{\pi_j^k}{p_j} - 2(\tau - \tau I_{1,\dots,N}) \left[\sum_j \frac{\pi_j^k}{p_j} \right]. \end{aligned} \quad (34)$$

Denoting the quantities above which do not depend on the order being followed by x , i.e.

$$x \triangleq -(\tau - \tau I_{1,\dots,N})^2, \quad (35)$$

the expression in (34) reduces to

$$x + \sum_j \left(\frac{1}{p_j} + 2 \left(\sum_{l < j} \frac{1}{p_l} \right) \right) \frac{\pi_j^k}{p_j}. \quad (36)$$

Hence by applying Theorem 4, any lower (upper) bound on the above expression would yield a lower (upper) bound on the limiting total scaled workload. Since $\sum_j \frac{\pi_j^k}{p_j} = \tau - \tau I_{\{1,\dots,N\}}$ for all orders k , we note that the expressions,

$$\sum_j \frac{1}{p_j} \times \frac{\pi_j^k}{p_j} \quad \text{and} \quad \sum_j \left(\sum_{l < j} \frac{1}{p_l} \right) \times \frac{\pi_j^k}{p_j},$$

are both maximized under the order k such that $p_1 \leq p_2 \leq \dots \leq p_N$. This gives us the upper bound. The proof of the lower bound is along the same lines. \blacksquare

XI. CONCLUDING REMARKS

We have presented a novel approach to understand the fine performance of scheduling policies utilized for real-time systems with hard deadline constraints. The fluctuation analysis performed shows that the cost associated with the MDF is well defined. The expression for the cost of MDF policy is determined. Moreover the policy optimal with respect to the limiting scaled total workload is described, and upper and lower bounds on the cost of any throughput-optimal policy are given.

XII. APPENDIX

A. Proof of Theorem 1

Proof: Throughout the proof we will use $CM_S(t)$ to represent $CM_S(\mathbf{d}(t))$, and CM_S to represent $CM_S(\mathbf{d})$. Similarly for a Lyapunov function $L(\mathbf{d})$, we will use $L(t)$ to represent $L(\mathbf{d}(t))$.

From inequality (18), we have

$$\begin{aligned} W_n^0 - W_{n+1}^0 &\leq 10c + \mathbb{1}_{\{W_n^0 - W_{n+1}^0 > 10c\}} (CM_{\{1,\dots,n\}} \\ &\quad - CM_{\{n+1,\dots,N\}}), \end{aligned} \quad (37)$$

and hence to bound the process $W_n^0 - W_{n+1}^0$ we will bound the process

$$\mathbb{1}_{\{W_n^0 - W_{n+1}^0 > 10c\}} (CM_{\{1,\dots,n\}}(t) - CM_{\{n+1,\dots,N\}}(t))$$

by considering Lyapunov functions of the form,

$$L_n(\mathbf{d}) \triangleq \mathbb{1}_{\{W_n^0 - W_{n+1}^0 > 10c\}} (CM_{\{1,\dots,n\}}(\mathbf{d}) - CM_{\{n+1,\dots,N\}}(\mathbf{d})), \quad (38)$$

the validity of which as a Lyapunov function can be easily verified.

We will use N different Lyapunov functions to show SSC. Ideally we would like to show that the processes $W_i(t) - W_j(t)$ are stochastically bounded. Instead we show that the processes $W_i^0(t) - W_{i+1}^0(t)$, where $i = 1, \dots, N-1$ are stochastically bounded. From this we infer that $W_1^0(t) - W_N^0(t)$ is stochastically bounded. Since $W_i(t) - W_j(t) \leq W_1^0(t) - W_N^0(t)$ for all i, j and at all times t , we conclude that $W_i(t) - W_j(t)$ is stochastically bounded.

We want to show that $L_n(t)$ is stochastically bounded for all $n = 1, \dots, N-1$. Clearly the drifts of $L_n(t)$ are uniformly bounded at all times t . Hence it suffices to show that the expected value of the drift of $L_n(t)$ is negative when $L_n(t)$ is greater than some threshold, since the result would then follow upon application of Lemma 1. Set this threshold as $\kappa := 20c$.

Since $|W_i(t+1) - W_i(t)| < c$ for all i and at all times t , we have $|L_n(t+1) - L_n(t)| < 2c$ for all n and t . Therefore from the definition of the Lyapunov function $L_n(\mathbf{d})$ (38), if $L_n(t) > 20c$, we have

$$L_n(t) = CM_{\{1,\dots,n\}}(t) - CM_{\{n+1,\dots,N\}}(t) \quad \text{and}$$

$$L_n(t+1) = CM_{\{1,\dots,n\}}(t+1) - CM_{\{n+1,\dots,N\}}(t+1). \quad (39)$$

Define the sets

$$U_t^n \triangleq \{i : W_i(t) \geq W_n^0(t)\}, \quad (40)$$

where U_t^n comprises of clients having the n highest workloads at time t . If $L_n(t) > 20c$, then the sets U_t^n and U_{t+1}^n are the same since $|W_i(t+1) - W_i(t)| < c$. So it suffices to show that

$$E[(CM_{\{1,\dots,n\}}(t+1) - CM_{\{n+1,\dots,N\}}(t+1)) \quad (41)$$

$$- (CM_{\{1,\dots,n\}}(t) - CM_{\{n+1,\dots,N\}}(t)) | L_n(t) > 20c] < -\eta, \quad (42)$$

for some $\eta > 0$, which is the same as showing

$$E[(CM_{\{1,\dots,n\}}(t+1) - CM_{\{1,\dots,n\}}(t)) \quad (43)$$

$$- (CM_{\{n+1,\dots,N\}}(t+1) - CM_{\{n+1,\dots,N\}}(t)) | L_n(t) > 20c] < -\eta.$$

Since the sets U_t^n and U_{t+1}^n are the same, the drift above simplifies to

$$\frac{1}{n} \left[\sum_{j=1}^n \frac{q_j}{p_j} - \sum_{j=1}^n \frac{\pi_j}{p_j} \right] - \frac{1}{N-n} \left[\sum_{j=n+1}^N \frac{q_j}{p_j} - \sum_{j=n+1}^N \frac{\pi_j}{p_j} \right], \quad (43)$$

where we have used π_j to denote the probability of successful transmission of packet of client which is served in the j -th place in the frame t . Now,

$$\begin{aligned} \sum_{j=1}^n \frac{q_j}{p_j} - \sum_{j=1}^n \frac{\pi_j}{p_j} &< \tau - \tau I_{\{1,\dots,n\}} - \sum_{j=1}^n \frac{\pi_j}{p_j} \\ &= \tau - \tau I_{\{1,\dots,n\}} - (\tau - \tau I_{\{1,\dots,n\}}), \\ &= 0, \end{aligned} \quad (44)$$

where the inequality follows from the assumption (5) and the equalities from (5), (8). Similarly,

$$\sum_{j=n+1}^N \frac{q_j}{p_j} - \sum_{j=n+1}^N \frac{\pi_j}{p_j} = - \left[\sum_{j=1}^n \frac{q_j}{p_j} - \sum_{j=1}^n \frac{\pi_j}{p_j} \right] > 0. \quad (45)$$

This shows that the drift is negative with η in equation (41) set to $\eta = 2 \left(\sum_{j=1}^n \frac{\pi_j}{p_j} - \sum_{j=1}^n \frac{q_j}{p_j} \right)$.

The statement $\lim_{t \rightarrow \infty} \tilde{W}_i$ is independent of i follows from Lemma 2. ■

B. Proof of Lemma 3

Proof: Applying the law of large numbers for martingales to the martingale difference sequence defined in (27), we have

$$\frac{\sum_{s=1}^t \mathbb{1}\{\text{order } k \text{ is followed in frame } s, \text{ packet of } j \text{ is delivered}\}}{t} - \pi_j^k \alpha_k(t) \rightarrow 0. \quad (46)$$

Summing the above expression over k we obtain,

$$\frac{\sum_{s=1}^t \mathbb{1}\{\text{packet of } j \text{ is delivered in frame } s\}}{t} - \sum_k \pi_j^k \alpha_k(t) \rightarrow 0, \quad (47)$$

where $\sum_k \alpha_k(t) = 1, \forall t$. But since,

$$\liminf_{t \rightarrow \infty} \frac{\sum_{s=1}^t \mathbb{1}\{\text{packet of } j \text{ is delivered in frame } s\}}{t} \geq q_j,$$

this gives

$$\liminf_{t \rightarrow \infty} \sum_k \pi_j^k \alpha_k(t) \geq q_j \text{ which implies } \liminf_{t \rightarrow \infty} \frac{\sum_k \pi_j^k \alpha_k(t)}{p_j} \geq \frac{q_j}{p_j}. \quad (48)$$

Summing the LHS of the last expression above over j , we have

$$\begin{aligned} \sum_j \frac{\sum_k \pi_j^k \alpha_k(t)}{p_j} &= \sum_k \alpha_k(t) \left(\sum_j \frac{\pi_j^k}{p_j} \right) \\ &= \sum_k \alpha_k(t) (\tau - \tau I_{\{1,\dots,N\}}) \\ &= \tau - \tau I_{\{1,\dots,N\}}, \end{aligned} \quad (49)$$

where the first equality follows by interchanging the order of summation, the second step by using the fact $\sum_j \frac{\pi_j^k}{p_j} = \tau - \tau I_{1,\dots,N}$ for all orders k , and the third step follows from the fact $\sum_k \alpha_k(t) = 1, \forall t$ (stated in (8)). Combining (48), (49) and noting that $\sum_j \frac{q_j}{p_j} = \tau - \tau I_{\{1,\dots,N\}}$ yields the result. ■

REFERENCES

- [1] I-Hong Hou, Vivek Borkar and P. R. Kumar, "A theory of qos for wireless," in *In Proc. of IEEE INFOCOM*, 2009.
- [2] I-Hong Hou and P.R. Kumar, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *INFOCOM, 2010 Proceedings IEEE*, 2010, pp. 1–9.
- [3] S. Hormann, "Fluctuation analysis of dependent random processes," Ph.D. dissertation, Oktober 2006.
- [4] A. Stolyar, "Max weight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic," in *Annals of Applied Probability*, vol. 14, no. 1, 2004, pp. 1–53.
- [5] A. Eryilmaz and R. Srikant, "Asymptotically tight steady-state queue length bounds implied by drift conditions," *Queueing Systems*, vol. 72.
- [6] Moshe Shaked, J. George Shanthikumar, *Stochastic Orders*, ser. Springer series in statistics. Springer, New York, 2007.
- [7] Rahul Singh, I-Hong Hou and P. R. Kumar, "Pathwise performance of debt based policies for wireless networks with hard delay constraints," *52nd IEEE Conference on Decision and Control*, 2013, to appear.
- [8] B. Hajek, "Hitting-time and occupation-time bounds implied by drift analysis with applications," *Advances in Applied Probability*, vol. 14, no. 3, pp. pp. 502–525, 1982.
- [9] Y. Chow, "On a strong law of large numbers for martinagles," *The Annals of Mathematical Statistics*, vol. 38, no. 2, pp. 610–610, 1967.
- [10] W. Stout, "A martingale analogue of Kolmogorov's law of the iterated logarithm," *Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 15, no. 4, pp. 279–290, 1970. [Online]. Available: <http://dx.doi.org/10.1007/BF00533299>
- [11] W. Feller, *An Introduction to Probability Theory and Its Applications*. Wiley, January 1968, vol. 1.