

Stability and asymptotic optimality of opportunistic schedulers in wireless systems*

U. Ayesta^{1,2}, M. Eraisquin^{1,3}, M. Jonckheere⁴, I.M. Verloop¹

¹BCAM – Basque Center for Applied Mathematics, Derio, Spain

²IKERBASQUE, Basque Foundation for Science, Bilbao, Spain

³UPV/EHU, University of the Basque Country, Bilbao, Spain

⁴CONICET, Buenos Aires, Argentina

ABSTRACT

We investigate the scheduling of a common resource between several concurrent users when the feasible transmission rate of each user varies randomly over time. Time is slotted and users arrive and depart upon service completion. This may model for example the flow-level behavior of end-users in a narrowband HDR wireless channel (CDMA 1xEV-DO). As performance criteria we consider the stability of the system and the mean delay experienced by the users. Given the complexity of the problem we investigate the fluid-scaled system, which allows to obtain important results and insights for the original system: (1) We characterize for a large class of scheduling policies the stability conditions and identify a set of maximum stable policies, giving in each time slot preference to users being in their best possible channel condition. We find in particular that many opportunistic scheduling policies like Score-Based [9], Proportionally Best [1] or Potential Improvement [4] are stable under the maximum stability conditions, whereas Relative-Best [10] or the $c\mu$ -rule are not. (2) We show that choosing the right tie-breaking rule is crucial for the performance (e.g. average delay) as perceived by a user. We prove that a policy is asymptotically optimal if it is maximum stable *and* the tie-breaking rule gives priority to the user with the highest departure probability. In particular, we show that simple priority-index policies with a myopic tie-breaking rule, are stable and asymptotically optimal. All our findings are validated with extensive numerical experiments.

*Research partially supported by grant MTM2010-17405 (Ministerio de Ciencia e Innovación, Spain) and grant PI2010-2 (Department of Education and Research, Basque Government). M. Eraisquin's Ph.D. is supported by grants ECO2008-00777 and FPU AP2008-02014, both of the MEC (Spain). The work of M. Jonckheere was carried out thanks to the BCAM visiting-fellow program.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VALUETOOLS 2011, May 16-20, Paris, France

Copyright © 2011 ICST 978-1-936968-09-1

DOI 10.4108/icst.valuetools.2011.245741

Categories and Subject Descriptors

D.4.8 [Performance]: Queueing Theory

Keywords

wireless system, stability, asymptotically optimal, opportunistic scheduling, $c\mu$ -rule

1. INTRODUCTION

Next generation wireless networks are expected to support a wide variety of data services. Due to fading and interference effects, for each user, the quality of a downlink channel, and hence its transmission rate, fluctuates over time. This has triggered a large amount of work aiming at understanding the performance of channel-aware scheduling policies. It is by now accepted that so-called “opportunistic schedulers” have many desirable properties (see for example [10]). A policy is called opportunistic if it takes advantage of the channel fluctuations by serving a user whose channel condition is “good” in some sense with respect to its own statistical behavior. With the objective of minimizing mean users’ delay, there arises a key tradeoff in the design of scheduling mechanisms between making full use of the opportunistic gains (hence ensuring a stable system) and prioritizing users having small residual service sizes.

Broadly speaking, researchers have explored scheduling in wireless systems both at the packet level and at the flow level. In packet-level models it is typically assumed that there exists a finite number of permanent users. The focus of the scheduler is on the number of packets in the queue of each user. We refer for example to [31, 2, 30, 17, 24, 3, 27] for this line of research. In a flow-level model instead, users arrive randomly to the system and leave after receiving their finite-sized service demands. This allows to capture the performance as perceived by the end-users, see for example [9, 20, 10, 25, 21, 1, 4, 28]. For a survey paper on flow-level modeling we refer to [23] and [11]. In [22], hybrid models are studied.

The performance evaluation and optimization of wireless networks at the flow level has proved to be extremely challenging. One of the most successful approaches has been the so-called time-scale separation argument (see [10, 12, 27, 1, 13]) where it is assumed that at the flow scale the dynamics of the channel fluctuations can be averaged out. Under this time-scale assumption, it was shown in [12] that any utility-based scheduling policy is stable in a flow-level model. The authors of [1] make the same assumption when they dis-

Discuss rate-stability for priority-index policies. Another approach is the Lagrangian-relaxation method (introduced in [32]) used in [4]. This allowed the authors of [4] to construct the Potential Improvement (PI) scheduling policy, which is optimal for a relaxed optimization problem. In addition, several other policies have been proposed and numerically investigated in the literature, among others the Proportional Fair [14] discipline, the Score-Based (SB) algorithm [9], the Relative Best (RB) scheduler [7] and Proportionally Best (PB) [1].

To sum up, without a time-scaling separation argument, which is a rather strong assumption, the performance of opportunistic schedulers, regarding stability and performance perceived by the users, is not well understood. In order to gain better insight into the latter issue, in this paper we will study a flow-level model without the time-scale separation assumption. More precisely, we assume that data users arrive randomly in time and have a finite amount of data to download. Time is slotted and the quality of the channel condition of each user varies per time slot. In every time slot at most one user may be served. We are interested in stability and optimization of the system. Given the complexity of the problem we first prove convergence of the fluid scaled system towards a unique fluid limit. We note that the precise characterization of the fluid limit involves averaging phenomena of the scaled system which is not grasped by the usual description of the weak limits.

The fluid-limit description allows us to obtain several important results and insights for the original wireless system. First of all, we can characterize the maximum stability conditions (the weakest possible conditions on the traffic parameters such that there exists a scheduling policy that makes the system positive recurrent) and show that the set of policies that are stable under the maximum stability condition have a very simple characterization: whenever there are users present that are currently in their best channel condition, only such users are served. Such a characterization was previously given for rate stability [1], but, to the best of our knowledge, stochastic stability was still an open issue.

Second, for a large class of scheduling policies we determine the stability conditions and conclude that many known opportunistic scheduling policies like Score-Based [9], Proportionally Best [1] or Potential Improvement [4] are stable under the maximum stability conditions, whereas Relative-Best [7, 10] or $c\mu$ -rule are not.

We furthermore demonstrate the importance for the choice of the tie-breaking rule when the goal is to optimize the performance. Until now, the literature proposed to break ties at random, see for example [7, 9, 10, 1]. We propose the myopic tie-breaking rule, i.e., give priority to the user with highest instantaneous departure probability when there are multiple users in their best channel condition. We prove that the myopic tie-breaking rule is asymptotically optimal and our numerical experiments further illustrate that the myopic tie-breaking rule significantly improves the performance. This in turn shows that one can use simple priority-index policies that will be both maximum stable and asymptotically optimal.

The rest of the paper is organized as follows. In Section 2 we present the model. In Section 3 we introduce the policies of interest and define their tie-breaking rules. In Section 4 we derive fluid limits for a large class of policies. In Section 5 we

present our stability results. In Section 6 we characterize the asymptotically optimal policies and discuss the importance of the tie-breaking rule. In Section 7 we perform numerical experiments to validate our theoretical findings.

2. MODEL DESCRIPTION

We consider a time-slotted system serving one user in each time slot. This models for instance a CDMA 1xEV-DO system. There are K classes of users, and in each time slot the number of class- k users arriving to the system, A_k , follows an i.i.d. sequence of random variables, with $\mathbb{E}(A_k) = \lambda_k < \infty$ and $\mathbb{E}(A_k^2) = \gamma_k < \infty$.

For each user the departure probability varies over time as the quality of the channel is changing from slot to slot. The quality of the channel (or state of the channel) for a class- k user is modeled as an i.i.d. sequence of random variables taking values in the finite set $\mathcal{N}_k := \{1, 2, \dots, N_k\}$. For each time slot we let $q_{k,n}$ denote the probability that a class- k user is in channel state $n \in \mathcal{N}_k$. Associated with channel state n is a departure probability $\mu_{k,n}$. This can be used for instance to model a system in which the service requirements of users is geometric (see Remark 1). Without loss of generality we assume that the channel condition labels are ordered such that $0 \leq \mu_{k,1} \leq \mu_{k,2} \leq \dots \leq \mu_{k,N_k} \leq 1$, and $q_{k,N_k} \mu_{k,N_k} \neq 0$ for all k . The channel condition of a class- k user is independent of the channel conditions of all the other users and of the channel quality history.

In each time slot, a scheduler/policy f decides which user is served. Because of the Markov property we can focus on policies that base decisions on the current number of users present in the various classes and on their channel conditions.

Since the channel conditions are independent and identically distributed for each time slot, without loss of generality, we can focus on the number of users in each class instead of the number of users in each state. For a given scheduling policy f , let $X_k^f(t)$ denote the number of class- k users in the various classes at time slot t and $X^f(t) = (X_1^f(t), \dots, X_K^f(t))$. Since the channel conditions are i.i.d. and independent of the process $X^f(\cdot)$, it follows that X^f is a Markov chain.

Performance criteria: Our performance criteria are stability and long-run average number of users. We use the following definition for stability:

DEFINITION 1. *A scheduling policy f is stable if the process X^f is positive recurrent.*

Because of the time-varying channel conditions the system is not work-conserving, and hence it depends strongly on the employed scheduling policy whether the system can be made stable. We define the *maximum stability conditions* as the conditions on the traffic inputs such that there exists a policy that can make the system stable. A *maximum stable policy* is a policy that is stable under the maximum stability conditions. From the performance point of view it is therefore of crucial importance to design a scheduler that is maximum stable.

Besides stability, another important performance measure is the long-run time-average holding cost,

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^K \sum_{t=0}^T c_k \mathbb{E}(X_k^f(t)), \quad (1)$$

with $c_k > 0$ the holding cost incurred per time slot for having

a class- k user in the system. In the case $c_k = 1$, for all k , this is equivalent to minimizing the mean sojourn time (because of Little's law).

REMARK 1 (Modeling of a wireless data network). *The mathematical model, even though simple, captures some of the key properties of wireless communication systems. Time is slotted, as is the case in the CDMA 1xEV-DO [6] and the OFDM-based LTE systems [29]. The available transmission rate of each user fluctuates due to fading effects, and as a consequence, it varies from one slot to another. We note that in real systems the number of feasible transmission rates is finite (see [6]).*

In a wireless system one may classify the users into different classes based on their applications or traffic conditions for example. Let the service requirement (in bits) of a class- k user be a geometric random variable denoted by B_k , and let $\mathbb{E}(B_k)$ denote its expectation. Let Δ denote the amount of bits transferred in one slot under the current channel condition. Note that in practice Δ will vary from slot to slot depending on the channel condition, and the allocation. The probability that a user leaves the system is approximately $\mathbb{P}(b \leq B_k \leq b + \Delta | B_k > b) \approx \Delta / \mathbb{E}(B_k)$, which does not depend on the attained service b (memoryless property of the geometric distribution). This expression becomes asymptotically exact as the ratio $\Delta / \mathbb{E}(B_k)$ goes to 0. Hence, this is the case if the mean service requirement (in bits) of a user is very large compared to the amount of bits that can be served in one slot. Let $s_{k,n}$ denote the transmission rate (in bits per second) of a class- k user when the channel state is n . So taking $\Delta = s_{k,n} \cdot t_c$ we obtain that the departure probability of a class- k user under channel condition n (when it is served at full capacity) can be approximated by

$$\mu_{k,n} := \frac{s_{k,n} \cdot t_c}{\mathbb{E}(B_k)}, \quad (2)$$

where t_c is the length of the slot (for example $t_c = 1.67\text{ms}$ in the CDMA 1xEV-DO system).

3. POLICIES

In this section we introduce scheduling policies that will be used throughout the paper. Most of these policies are opportunistic, meaning that they take advantage of channel fluctuations by serving a user whose channel condition is "good" in some sense with respect to its own statistical behavior.

Priority-index policies are very popular due to their simplicity from an implementation point of view. A priority-index policy is characterized by an index function that assigns an index to each user based solely on its class and its current state.

DEFINITION 2 (Priority-index policy). *In every time slot, users that have the highest index among all users present in the system are served.*

Priority-index policies might need to be augmented with a suitable tie-breaking rule. Such a rule refers to the strategy adopted when there is a tie on the highest index value. A tie means that there are several users present having the highest index value, but these users belong to different classes. In the literature, most of the papers specify to break ties at

random (see for example [9, 7, 1]). One of our main contributions will be to emphasize that the choice for the tie-breaking rule is crucial for the performance of the system (as will be explained in Sections 6 and Section 7).

An important subset of the priority-index policies are the so-called weight-based policies.

DEFINITION 3 (Weight-based policy). *A priority-index policy with index function $\omega_k \mu_{k,n}$. Here ω_k denotes a class dependent weight.*

Important examples of weight-based policies are: the $c\mu$ -rule ($\omega_k = c_k$, with c_k the holding cost), Relative Best (RB) [7] ($\omega_k = 1 / \sum_{n=1}^{N_k} q_{k,n} \mu_{k,n}$), and Proportionally Best (PB) [1] ($\omega_k = 1 / \mu_{k,N_k}$). For all these policies, ties are broken at random.

In [9] the Score-Based (SB) policy is introduced. SB is a priority-index policy where the index value of a class- k user in state n is given by $\sum_{\tilde{n}=1}^n q_{k,\tilde{n}}$, and ties are broken at random. In [4] the Potential Improvement (PI) policy is introduced. PI is a priority-index policy with as index value

$$\frac{c_k \mu_{k,n}}{\sum_{\tilde{n}>n} q_{k,\tilde{n}} (\mu_{k,\tilde{n}} - \mu_{k,n})}$$

and the tie-breaking rule is the myopic-rule, that is, among the users with the highest index, select the one with highest value for $c_k \mu_{k,N_k}$, $k = 1, \dots, K$. The c_k 's refer to the holding cost introduced in Section 2.

It will be convenient to define the following two classes of policies, which play an important role in the results on the stability analysis and asymptotic optimality.

DEFINITION 4 (Best Rate (BR) policies). *The BR policies are such that whenever there are users present that are currently in their best channel condition, i.e., in state N_k , such a user is served.*

DEFINITION 5 (Best Rate Priority (BRP) policies). *The BRP policies are such that whenever there are users present that are currently in their best channel condition, i.e., in state N_k , among those users the one with the highest value for $c_k \mu_{k,N_k}$ is served.*

The collection of BRP policies is a subset of the BR policies. In our main results we will show that the classes of policies BR and BRP have desirable properties: In Section 5 we show that any BR policy is stable under the maximum stability conditions and in Section 6 we show that BRP policies are asymptotically optimal.

In Figure 1 we have depicted a diagram in which we have summarized the various (classes of) policies. From the policies introduced above, SB, PB and PI are BR policies. This follows since the highest possible index value is 1 for SB and PB, and ∞ for PI, and these indices can only be obtained whenever a user is in its best possible channel condition. For RB and the $c\mu$ rule it can be that the index value of a class- k user in state $n < N_k$ is larger than the index value of a class- l user in state N_l , hence RB is not a BR policy.

From the policies introduced above, PI is the only BRP policy. This results from the fact that instead of a random tie-breaking rule, the myopic tie-breaking rule is used in PI. Hence, whenever there are users present in their best channel condition (i.e., having index ∞), the user having the highest value for $c_k \mu_{k,N_k}$ is chosen.




Figure 1: Classification of schedulers.

4. FLUID LIMITS AND CONVERGENCE

In this section we study the fluid-scaling limits for a large class of policies. Fluid scaling or time-space scalings, corresponding to “zooming out” the trajectories, have been used extensively to study stochastic processes with complex dynamics [16]. The limiting processes are usually much simpler to describe while they provide crucial insights on the behavior of the non-scaled version of the process. However, most results for stochastic networks are concerned with a generic description of the weak limits (usually not unique) [15] or convergence in probability to a unique limit but only for a subset of the state space [19]. To be able to characterize both stability and asymptotic optimality for a wide range of policies, we need to prove convergence in probability towards a unique fluid limit and characterize the whole trajectory of this fluid limit. This requires to deal with averaging phenomena: for instance, it may happen that one class of users reaches its stationary regime before the other classes do. In this case, the drift of the other classes needs to be averaged with the stationary distribution of this class. Hence, a description of the fluid limit will involve averaged drifts (we refer to this as second-vector fields, following [18]).

For the stochastic process $X^f(t)$ associated with a policy f , we define the drift function by $\delta^f(x) := (\delta_1^f(x), \dots, \delta_K^f(x))$, $x \in \mathbb{N}^K$, with

$$\delta_i^f(x) := \mathbb{E}(X_i^f(1) - x_i | X^f(0) = x).$$

(We will drop the superscript f when it is clear that we consider a unique policy.) We say that a vector field $v : \mathbb{N}^K \rightarrow \mathbb{R}^K$ has uniform limits [13] if for any $\mathcal{U} \subset \{1, \dots, K\}$, there exists a function $v^{\mathcal{U}} : \mathbb{N}^{|\mathcal{U}|} \rightarrow \mathbb{R}^K$ (constant when $\mathcal{U} = \emptyset$) such that

$$\lim_{R \rightarrow \infty} \sup_{x \in \mathbb{N}^K : |x_{\mathcal{U}^c}| > R} |v(x) - v^{\mathcal{U}}(x_{\mathcal{U}})| = 0,$$

where $x_{\mathcal{U}}$ denotes the restriction of the vector x to indices in the subclass \mathcal{U} . Intuitively, this means that the drift vector-field has limits when we make the number of users of some of the classes go to infinity, and that we can interchange the order of the coordinates when taking these limits. We shall suppose in the following that the drift vector has uniform limits and thus that we can define the asymptotic drifts $\delta^{\mathcal{U}} : \mathbb{N}^{|\mathcal{U}|} \rightarrow \mathbb{R}^K$ as follows:

$$\delta^{\mathcal{U}}(x_{\mathcal{U}}) := \lim_{x_k \rightarrow \infty, k \in \mathcal{U}^c} \delta(x). \quad (3)$$

Here, \mathcal{U}^c (the complementary set of \mathcal{U}) corresponds to the “saturated” classes for which we make the number of users go

to infinity. We define the stochastic processes $X^{\mathcal{U}}$ as the \mathcal{U} -dimensional stochastic process corresponding to the original process seeing an infinite number of users of class $k \in \mathcal{U}^c$ and let $\pi^{\mathcal{U}}$ denote its stationary measure assuming it exists. We define the averaged drift vectors by

$$\tilde{\delta}^{\mathcal{U}} = \sum_{x \in \mathbb{N}^{|\mathcal{U}|}} \delta^{\mathcal{U}}(x) \pi^{\mathcal{U}}(x). \quad (4)$$

Finally, following [13] we say that a vector field v is partially increasing if $v_i(x)$ is increasing in x_j for all $j \neq i$. These assumptions, which are crucial to prove the convergence towards the fluid limit, are verified for many cases of interest, see the next lemma.

LEMMA 4.1. *A priority index policy or a BR policy with non-state dependent tie-breaking rule (i.e., independently of the numbers of users) induces a partially increasing drift vector field with uniform limits.*

Proof: We prove the lemma for BR policies, the other case being similar. When increasing the number of users of one class only, the probability that this class has at least one user in its best possible state is increased. Hence, given that the tie-breaking rule does not depend on the number of users, the probability that this class is served is increased while the probability that a user of another class is served decreases. This implies that the drift vector field is partially increasing.

Using further the independence of the channel variations, the probability that class $i \in \mathcal{U}^c$ has at least one user in its best state is $1 - (1 - q_{i, N_i})^{x_i}$, where x_i is the number of class- i users. Hence, when the numbers of class- i users, $i \in \mathcal{U}^c$, go to infinity, the probability of having at least one user in its best state is converging to 1. Together with the property that the tie-breaking rule does not depend on the number of users, this implies that:

$$\delta(x) = \delta^{\mathcal{U}}(x_{\mathcal{U}}) \prod_{i \in \mathcal{U}^c} (1 - (1 - q_{i, N_i})^{x_i}) + o(1/|x|),$$

which in turn implies the uniform convergence of δ . \square

We construct different realizations of the processes, depending on the initial state. To be precise, for a given policy f we let $X^{f,r}(t)$ denote the number of class- k users at time t when the initial state equals $X_k^r(0) = rx_k(0)$, $k = 1, \dots, K$, with $r \in \mathbb{N}$. We are then interested in the fluid-scaled processes $Y^{f,r}(t) := \frac{X^{f,r}(lr t)}{r}$, $t \geq 0$, $k = 1, \dots, K$, with $Y^r(0) = x(0)$.

We now state the convergence towards a unique fluid limit which is the main result of this Section. The proof can be found in the technical report [5].

THEOREM 4.2. *For a given policy f inducing a partially increasing drift vector field with uniform limits, we have*

$$\lim_{r \rightarrow \infty} \mathbb{P} \left(\sup_{0 \leq s \leq t} |Y^{f,r}(s) - y^f(s)| \geq \epsilon \right) = 0, \quad \text{for all } \epsilon > 0,$$

with $y^f(t)$ a piece-wise linear function that can be described recursively as follows. Let $\mathcal{U}_0 = \emptyset$ and $T_0 = 0$. Then we have,

$$\frac{dy_k^f(t)}{dt} = \tilde{\delta}_k^{f, \mathcal{U}_l}, \quad t \in [T_l^f, T_{l+1}^f], \quad (5)$$

$$\text{with } T_{l+1}^f = T_l^f + \min_{k \in \mathcal{U}_l^c, \tilde{\delta}_k^{f, \mathcal{U}_l} < 0} \frac{y_k^f(T_l^f)}{-\tilde{\delta}_k^{f, \mathcal{U}_l}}, \quad (6)$$

$$\text{and } \mathcal{U}_{l+1} = \mathcal{U}_l \cup \arg \min \left\{ \frac{y_k^f(T_l^f)}{-\tilde{\delta}_k^{f, \mathcal{U}_l}}, k \in \mathcal{U}_l^c \right\}, \quad (7)$$

with the assumption that if there exists no $k \in \mathcal{U}_l^c$ with $\tilde{\delta}_k^{f, \mathcal{U}_l} < 0$, then $T_{l+1}^f = \infty$.

REMARK 2 (CALCULATION OF THE AVERAGED DRIFTS). *Theorem 4.2 characterizes the fluid limit as a piece-wise linear function with slopes $\tilde{\delta}^{\mathcal{U}}$. In practice, the calculations of these slopes involve:*

- deriving the asymptotic drifts (see (3)),
- calculating the stationary distributions of $X^{\mathcal{U}}$,
- averaging the asymptotic drifts with these stationary distributions (see (4)).

For instance, assume $K = 2$, $N_1 = 2$ and $N_2 = 1$ and a Bernoulli arrival process. Consider the policy that gives priority to the best class-1 user present in the system and otherwise (i.e., when there is no class 1) serves a class-2 user. Then

$$\delta^{\emptyset} = (\lambda_1 - \mu_{1, N_1}, \lambda_2),$$

$$\delta^{\{1\}}(x_1) = (\lambda_1 - s_1(x_1), \lambda_2 - \mu_{2, N_2} \mathbf{1}_{(x_1=0)}),$$

with $s_1(x_1) = \mu_{1, N_1} (1 - (1 - q_{1, N_1})^{x_1}) + \mu_{1, N_1 - 1} (1 - q_{1, N_1})^{x_1}$. The process $X^{\{1\}}$ is a 1-dimensional Markov chain with stationary distribution

$$\pi^{\{1\}}(x_1) = C \prod_{j=1}^{x_1} \frac{\lambda_1 (1 - s_1(j-1))}{(1 - \lambda_1) s_1(j)},$$

where C is a normalization constant. The average drift can now be computed using (4).

In the specific case of BR policies with a priority-type tie-breaking rule, we can in fact explicitly derive the fluid limit, since we can make use of rate conservation arguments. This will prove to be very useful to obtain both maximal stability and asymptotic optimality statements.

PROPOSITION 4.3. *Consider a BR policy with a priority type tie-breaking rule. Let us reorder the classes according to the priority ordering. The associated asymptotic drift in the interior of the orthant is*

$$\delta^{\emptyset} = (\lambda_1 - \mu_{1, N_1}, \lambda_2, \dots, \lambda_K). \quad (8)$$

If $T_1 < \infty$, then

$$\tilde{\delta}^{\{1\}} = (0, \lambda_2 - \mu_{2, N_2} \left(1 - \frac{\lambda_1}{\mu_{1, N_1}}\right), \dots, \lambda_K), \quad (9)$$

and more generally, if $T_{k-1} < \infty$, then $\mathcal{U}_{k-1} = \{1, \dots, k-1\}$ and

$$\tilde{\delta}^{\mathcal{U}_{k-1}} = (0, \dots, 0, \lambda_k - \left(1 - \sum_{j \in \mathcal{U}} \frac{\lambda_j}{\mu_{j, N_j}}\right) \mu_{k, N_k}, \dots, \lambda_K). \quad (10)$$

Proof: Using Lemma 4.1, the drift δ associated to a best rate policy is partially increasing with uniform limits and Theorem 4.2 holds. When $\mathcal{U} = \emptyset$, there are infinitely many users of each class and hence there is always a class-1 user in its best state, which directly implies (8). Note that $T_1 < \infty$ if and only if $\frac{\lambda_1}{\mu_{1, N_1}} < 1$. In this case the process $X^{\{1\}}$ is ergodic. For $T_1 \leq t \leq T_2$, we can simplify the asymptotic drift $\tilde{\delta}^{\{1\}}$ using the specific properties of the policy and rate conservation arguments: let A_{1, x_1} be the event that class 1 is served given there are x_1 users of class 1. Since $X^{\{1\}}$ is ergodic, we have the following rate conservation equation

$$\sum_{x_1} \pi^{\{1\}}(A_{1, x_1}) \mu_{1, N_1} = \lambda_1,$$

(for details see the technical report [5]), which gives that $\sum_{x_1} \pi^{\{1\}}(A_{1, x_1}^c) = 1 - \frac{\lambda_1}{\mu_{1, N_1}}$. Since $\mathcal{U} = \{1\}$ (so in particular there is still an infinite amount of class-2 users which are exclusively served when there are no class-1 users in their best state) class 2 receives service at rate $\mu_{2, N_2} \sum_{x_1} \pi^{\{1\}}(A_{1, x_1}^c) = \mu_{2, N_2} (1 - \frac{\lambda_1}{\mu_{1, N_1}})$ which gives (9).

Consider now the case where the classes in $\mathcal{U} := \mathcal{U}_{k-1}$ are stationary (assuming as in the previous case that $\sum_{j \in \mathcal{U}} \frac{\lambda_j}{\mu_{j, N_j}} < 1$). Let $A_{j, x_{\mathcal{U}}}$ be the event that class $j \in \mathcal{U}$ is served given there are x_i users of class i , $i \in \mathcal{U}$. By similar rate conservation arguments we obtain

$$\sum_{x_{\mathcal{U}}} \pi^{\mathcal{U}}(A_{j, x_{\mathcal{U}}}) \mu_{j, N_j} = \lambda_j, \quad j \in \mathcal{U}.$$

Noting that the sets $A_{j, x_{\mathcal{U}}}$ are disjoint, $\sum_{x_{\mathcal{U}}} \pi^{\mathcal{U}}(\cup_{j \in \mathcal{U}} A_{j, x_{\mathcal{U}}}) = \sum_{j \in \mathcal{U}} \frac{\lambda_j}{\mu_{j, N_j}}$, which gives (10). \square

REMARK 3. *For all BR policies where the scheduler chooses with probability $\alpha_i^{\mathcal{U}}$ to serve class i when a subset \mathcal{U} of classes has at least one user in its best channel condition, the fluid limit in the interior of the orthant has a drift given by:*

$$\delta^{\emptyset} = (\lambda_1 - \alpha_1^{\{1, \dots, K\}} \mu_{1, N_1}, \lambda_2 - \alpha_2^{\{1, \dots, K\}} \mu_{2, N_2} \quad (11)$$

$$, \dots, \lambda_K - \alpha_K^{\{1, \dots, K\}} \mu_{K, N_K}). \quad (12)$$

However, we cannot simplify further the second-vector fields in general. An exception is the case $K = 2$, i.e., two classes. Then, using the rate conservation argument as in the previous proposition, we obtain (assuming w.l.o.g. that class 1 empties first)

$$\tilde{\delta}^{\{1\}} = (0, \lambda_2 - \left(1 - \frac{\lambda_1}{\mu_{1, N_1}}\right) \mu_{2, N_2}).$$

5. STABILITY ANALYSIS

The derivation of fluid limits allows us to conclude for stochastic stability. In particular, we describe the stability

conditions for a large class of policies, we derive the maximal achievable stability region, and in addition characterize the class of policies that will achieve maximum stability.

The next Theorem characterizes the stability of any policy having a partially increasing drift vector field with uniform limits.

THEOREM 5.1. *A policy f inducing a partially increasing drift vector field with uniform limits is stable if $T_l^f < \infty$ for all l , where T_l^f is given by Theorem 4.2.*

Proof: If $T_l^f < \infty$ for all l , the fluid limit described in Theorem 4.2 is equal to 0 for t large enough, i.e., $Y_k^{f,r}(t)$ converges in probability to 0 for t large enough. In addition, the random variable $Y_k^{f,r}(t)$ is uniformly integrable. This can be seen by the fact that $Y_k^{f,r}(t)$ can be upper bounded by $Y_k^r(0)$ plus the total amount of users that have arrived until time rt divided by r , which is uniform integrable, see [15].

The convergence in probability to 0 of the scaled system and the uniform integrability together imply that $\lim_{r \rightarrow \infty} \mathbb{E}(Y_k^{f,r}(t)) = 0$, for t large enough, for all k , see [8, Theorem 3.5]. Using an extended Foster-Lyapunov criterion, as expressed in [18] or [26, Corollary 9.8] this implies the positive recurrence of the Markov process. \square

We now state the maximum stability condition and characterize a set of policies that achieve maximum stability.

THEOREM 5.2. *The maximum stability condition is*

$$\sum_{k=1}^K \frac{\lambda_k}{\mu_{k,N_k}} < 1. \quad (13)$$

In addition, any BR policy is maximum stable.

Proof: We first prove that a BR policy with a non-state dependent tie-breaking rule is stable when (13). For the proof of a BR policy with a general tie-breaking rule we refer the reader to the technical report [5].

Consider the fluid limit $y(t)$ of a BR policy with a non-state dependent tie-breaking rule. We prove that the fluid limit is 0 after a finite time by considering the Lyapunov function $\sum_{k=1}^K \frac{y_k(t)}{\mu_{k,N_k}}$. We have that $\sum_{k=1}^K \frac{dy_k(t)}{dt} \frac{1}{\mu_{k,N_k}} = \sum_{k=1}^K \frac{\delta_k^{\mathcal{U}}}{\mu_{k,N_k}}$. For any BR policy we have that for all \mathcal{U} and for all $x_{\mathcal{U}}$,

$$\sum_{k=1}^K \frac{\delta_k^{\mathcal{U}}(x_{\mathcal{U}})}{\mu_{k,N_k}} = \lim_{x_i \rightarrow \infty, i \in \mathcal{U}^c} \sum_{k=1}^K \frac{\delta_k(x)}{\mu_{k,N_k}} = \sum_{k=1}^K \frac{\lambda_k}{\mu_{k,N_k}} - 1 < 0,$$

because the server spends all its capacity on users in their best state (since $\mathcal{U}^c \neq \emptyset$). Hence, $\sum_{k=1}^K \frac{dy_k(t)}{dt} \frac{1}{\mu_{k,N_k}} < 0$, so the fluid limit empties in finite time, i.e., $T_l < \infty$ for all l . Therefore, by Theorem 5.1 the policy is stable.

When condition (13) is not satisfied, the system is not rate stable, which precludes stability. \square

Condition (13) was recognized as the maximal rate stability condition in [1] and as the maximum stability condition under a time-scale separation assumption in [10].

We note that SB, PI and PB are stable under the maximum stability conditions (they belong to the class of BR policies). The intuition behind Theorem 5.2 is that asymptotically the system under a BR policy behaves as a classical

work-conserving system where class k has departure probability μ_{k,N_k} . On the contrary, other policies, including RB and the $c\mu$ -rule, spend (at the fluid scale) a non-negligible fraction of time serving users that are not in their best states, and are therefore not maximum stable. For an example, we refer to Section 7 where we numerically derive stability conditions for RB and the $c\mu$ -rule.

6. ASYMPTOTIC OPTIMALITY

Besides stability, another important performance measure concerns the long-run average holding cost as given in (1). Deriving an optimal policy with respect to this criterion is difficult and the size of the state space makes the problem intractable. For this reason we introduce a related deterministic control problem, which allows us to prove that any BRP policy is asymptotically optimal for the original stochastic system. This emphasizes the important role of the tie-breaking rule in order to achieve efficient performance of the system.

We study the following deterministic fluid control model, which arises from the original stochastic model by only taking into account the mean drifts.

$$\min \int_0^D \sum_{k=1}^K c_k x_k(t) dt, \quad (14)$$

$$\text{subject to} \quad (15)$$

$$x_k(t) = x_k(0) + \lambda_k t - \sum_{n=1}^{N_k} \mu_{k,n} \int_0^t u_{k,n}(v) dv, \quad (16)$$

$$x_k(t) \geq 0, \quad k = 1, \dots, K, \quad (17)$$

such that for all $v \geq 0$,

$$\sum_{k=1}^K \sum_{n=1}^{N_k} u_{k,n}(v) \leq 1, \quad u_{k,n}(v) \geq 0, \quad \text{for all } k, n, \quad (18)$$

and the control functions $u_{k,n}(v)$ being integrable.

We remark that though in general the fluid limit of a policy does depend on the distributions of the random environments (i.e., the $q_{k,n}$'s), these do not appear in the above equations of the fluid control model. This is because the fluid trajectory $x_k(t)$ can be interpreted as a limit of a fluid-scaled process (see technical report [5] for details). Hence, when $x_k(t) > 0$ this implies that there are infinitely many class- k users so that with probability 1 there are infinitely many class- k user in each of the channel state conditions (this being independent of the exact values of the $q_{k,n} > 0$'s).

An optimal control $u^*(\cdot)$ and its corresponding trajectory $x^*(\cdot)$ are derived in the following lemma.

LEMMA 6.1. *Assume $c_1 \mu_{1,N_1} \geq c_2 \mu_{2,N_2} \geq \dots \geq c_K \mu_{K,N_K}$. The fluid control that solves the fluid control problem is as follows. Let $l = \arg \min \{k : x_k(t) > 0\}$. Then*

$$u_{k,N_k}^*(t) = \frac{\lambda_k}{\mu_{k,N_k}}, \quad \text{for } k < l, \quad u_{l,N_l}^*(t) = 1 - \sum_{i=1}^{l-1} \frac{\lambda_i}{\mu_{i,N_i}},$$

and $u_{k,n}^*(t) = 0$ otherwise.

Proof: It is immediate that $u_{k,n}^*(t) = 0$ for $n < N_k$, for all $k = 1, \dots, K$. Hence, the fluid control model reduces to a

multi-class server with service rates μ_{k,N_k} , $k = 1, \dots, K$ for which the optimal control is as stated in the lemma. \square

The optimal fluid cost serves as a lower bound for the fluid-scaled cost of the stochastic network, see the lemma below. The proof may be found in the technical report [5].

LEMMA 6.2. *For any policy f and $D > 0$ we have*

$$\liminf_{r \rightarrow \infty} \mathbb{E} \left(\int_0^D \sum_{k=1}^K c_k Y_k^{f,r}(t) dt \right) \geq \int_0^D \sum_{k=1}^K c_k x_k^*(t) dt.$$

We define a policy to be *asymptotically optimal* when the lower bound is obtained. The following Theorem characterizes a class of policies that is asymptotically optimal.

THEOREM 6.3. *Any BRP policy is asymptotically optimal.*

Proof: We have $\frac{dx_k^*(t)}{dt} = \lambda_k - u_{k,N_k}^*(t)\mu_{k,N_k}$, $k = 1, \dots, K$, with $u^*(\cdot)$ the optimal control as derived in Lemma 6.1. This drift coincides with the drift of the fluid limit $y^{BRP}(\cdot)$, see Proposition 4.3, hence $y^{BRP}(t) = x^*(t)$.

The terms $Y_k^{f,r}(t)$ are uniform integrable (see the proof of Theorem 5.1), hence we have that the random variable $\int_0^D \sum_{k=1}^K c_k Y_k^{f,r}(t) dt$ is uniformly integrable. Together with the fact that $Y^r(\cdot)$ converges in probability to $y^{BRP}(t) = x^*(\cdot)$, see Proposition 4.2, we then obtain that [8]

$$\lim_{r \rightarrow \infty} \mathbb{E} \left(\sum_{k=1}^K c_k \int_0^D Y_k^{f,r}(t) dt \right) = \int_0^D \sum_{k=1}^K c_k x_k^*(t) dt.$$

\square

Unfortunately, asymptotic optimality of BRP policies does not give any performance guarantee in terms of the long-run time-average holding cost as expressed in Equation (1). Numerical experiments reported in Section 7 indicate however that BRP policies can significantly outperform all other policies in terms of the long-run time-average holding cost.

To the best of our knowledge, the only policy studied in the literature that belongs to BRP, and hence is both maximum stable and asymptotically optimal, is PI. We recall that PI was derived in [4] as the solution of a relaxed optimization problem. We remark that SB will as well become asymptotically optimal when the myopic tie-breaking rule is used.

7. NUMERICAL EXPERIMENTS

We consider a CDMA 1xEV-DO system with two classes of users ($K = 2$). Time is slotted, with the length of one slot being $t_c = 1.67ms$. In each time slot, one new class- k user arrives with probability λ_k . We choose 10.257 kb as the expected service requirement of both a class-1 and class-2 user. Associated to the state of the channel, we have transmission rates (kb/s), see Table 1 (taken from [6]). We assume that class-1 users have five possible transmission rates while class-2 users have three. The corresponding probabilities $(q_{k,n})$ are given in Table 1. In addition, applying equation (2) we calculate the departure probabilities $(\mu_{k,n})$. We fix $\lambda_2 = 0.05$, so $\lambda_2/\mu_{2,N_2} = 0.5$. We set $c_1 = c_2 = 1$, so that we are interested in minimizing the total number of users in the system.

We compare the performance of the policies SB, RB, PI, PB and the $c\mu$ rule, which were introduced in Section 3. We summarize below the main conclusions of this section:

- The drifts of the fluid limit, δ^U , (which can be calculated numerically (and in some cases theoretically)) provide insightful information on the behavior under the different policies.
- We calculate the stability conditions under RB and the $c\mu$ -rule and observe that these are much more stringent than the one of BR policies (e.g. SB, PB and PI).
- Our simulations illustrate that the tie-breaking rule has a very big impact on the performance of the system and we observe that BRP policies, i.e., policies that combine opportunistic scheduling with the myopic tie-breaking rule, minimize the long-run time-average holding cost as expressed in (1).

Fluid limit.

We first illustrate how the scaled process converges to the fluid limit. We take $r = 10000$, $Y^r(0) = X(0)/r = (1, 1)$ and plot the scaled processes $Y_1^r(t)$, $Y_2^r(t)$, and $Y_1^r(t) + Y_2^r(t)$ for different policies, see Figure 2. In this simulation we set $\lambda_1 = 0.14$, so $\lambda_1/\mu_{1,N_1} = 0.35$.

We describe the fluid limit $y^f(t)$ as defined in Theorem 4.2. When both classes are present at the fluid scale, i.e., $\mathcal{U} = \emptyset$, the drift is

$$\delta^{f,\emptyset} = (\lambda_1 - \alpha^f \mu_{1,N_1}, \lambda_2 - (1 - \alpha^f) \mu_{2,N_2}), \quad (19)$$

see Remark 3. Here α^f is a random tie-breaking rule, i.e., in case of a tie, α^f is the probability that class 1 is favoured over class 2. For our set of parameters, the best class-1 user under the $c\mu$ -rule and RB is always preferred over the best class-2 users, i.e., there occur no ties, hence one can set $\alpha^f = 1$ in (19) for $f = c\mu, RB$. For PI, SB and PB we do have ties, and we set $\alpha^{PI} = 1$ and $\alpha^{SB} = \alpha^{PB} = 1/2$. In Table 2 we present the values of $\delta^{f,\emptyset}$. From the drifts it is clear that under all policies class 1 empties before class 2. The moment that this happens, T_1^f , can be derived from Theorem 4.2 and satisfies $T_1^{PI} = T_1^{c\mu} = T_1^{RB} < T_1^{SB} = T_1^{PB}$, see also Figure 2 a).

For $T_1^f < t \leq T_2^f$, the drift of class 1 is 0, whereas the drift of class 2 is going to depend on the policy. From Proposition 4.3 we have that for all BR policies (e.g. PI, SB and PB) $\delta^{f,\{1\}} = (0, \lambda_2 - \mu_{2,N_2}(1 - \lambda_1/\mu_{1,N_1}))$. For the $c\mu$ rule and RB we calculate the drift numerically using Remark 2. In particular we observe that these drifts are positive for the latter two policies, which implies instability of the system. We observe that for $t \leq T_1^f$ the number of class-2 users increases under policies PI, $c\mu$ and RB, while for SB and PB, the drift of class-2 users is negative for $t \leq T_2^{SB, PB}$.

A direct consequence of the drift function is that SB, PB, and PI (in fact all BR policies) empty the system all at the same time, i.e., T_2^f is the same. However, the performance of a policy will depend on the order in which classes are served. In the fluid limit, this is fully determined by the choice of the tie-breaking rule. Note that, as can be seen from Figure 2 c), PI (and hence any BR policy with the myopic tie-breaking rule) minimizes the total number of users at any moment in time.

Channel state	1	2	3	4	5	6	7	8	9	10	11
Transmission rate (kb/s) in CDMA	38.4	76.8	102.6	153.6	204.8	307.2	614.4	921.6	1228.8	1843.2	2457.6
Probabilities in CDMA	0.00	0.01	0.04	0.08	0.15	0.24	0.18	0.09	0.12	0.05	0.04
$q_{1,n}$	0	0	0.05	0	0.23	0	0.42	0	0.21	0	0.09
$q_{2,n}$	0	0	0.15	0	0.33	0	0.52	0	0	0	0
$\mu_{1,n}$	0	0	0.017	0	0.033	0	0.1	0	0.2	0	0.4
$\mu_{2,n}$	0	0	0.017	0	0.033	0	0.1	0	0	0	0

Table 1: The transmission rates and the corresponding channel condition probabilities in the CDMA 1xEV-DO wireless network, as reported in [6].




Figure 2: (a) Scaled number of class-1 users (b) Scaled number of class-2 users (c) Scaled total number of users

f	$\delta^{f,0}$		$\delta^{f,\{1\}}$	
	Class 1	Class 2	Class 1	Class 2
PI	-0.26	0.05	0	-0.015
$c\mu$ -rule	-0.26	0.05	0	0.0096
PB/SB	-0.06	0	0	-0.015
RB	-0.26	0.05	0	0.0004

Table 2: Drift of the fluid limit.

Stability region.

We now vary the value of λ_1 from 0.004 to 0.196, and as a consequence we have that $\rho := \lambda_1/\mu_{1,N_1} + \lambda_2/\mu_{2,N_2}$ varies from 0.51 to 0.99. The policies PI, PB and SB belong to the BR policies, and are hence stable when $\rho < 1$. For the $c\mu$ -rule and RB the stability condition can be calculated (numerically) by setting $\delta_2^{f,\{1\}}$ equal to zero and using Remark 2. In particular, the $c\mu$ -rule is stable if and only if $\rho < 0.79$ and RB is stable if and only if $\rho < 0.84$. In Figure 3 a) we plot the mean number of users, see (1), for different values of ρ and we observe that the number of users for these policies grows to infinity as the load approaches the critical value.

Impact of Tie-Breaking rule.

We study the impact of the tie-breaking rule on the performance of the system. In order to investigate this issue in more depth, we simulate PI under different random tie-breaking rules, i.e., we let the probability α vary from 0 until 1. We emphasize that PI as defined in [4] uses by default the myopic tie-breaking rule, i.e., $\alpha^{PI} = 1$. In Figure 3 b) we plot the relative degradation (in terms of the mean number of users, see (1)) over PI as we vary α . The results show that the myopic tie-breaking rule, which is asymptotically

optimal (see Theorem 6.3), is also optimal when minimizing the mean number of users. In addition, the relative degradation of the tie-breaking rule with $\alpha = 1/2$ (compared to the myopic tie-breaking rule $\alpha^{PI} = 1$) can be very large. For example, for $\rho = 0.8$ the degradation is 29% and for $\rho = 0.9$ it is 45%.

8. CONCLUSION

We have characterized the classes of policies that are maximum stable and asymptotically optimal in a system with random environment. An important conclusion, validated by numerical experiments, is that the tie-breaking rule has a tremendous impact on the performance. Our analysis also shows that simple priority-index policies like PI or SB with a myopic tie-breaking rule, are stable and asymptotically optimal. While in this model we assumed geometric service requirements, we do believe that direct extensions of all our results exist for phase-type distributed service requirements. In particular, we expect that an optimal tie-breaking rule will be of a simple priority-index type.

Acknowledgements We are grateful to S.C. Borst (TU/e and Bell-Labs) and P. Jacko (BCAM) for many fruitful discussions.

9. REFERENCES

- [1] S. Aalto and P. Lassila. Flow-level stability and performance of channel-aware priority-based schedulers. In *Proceedings of NGI, 6th EURO-NF Conference on Next Generation Internet*, 2010.
- [2] M. Andrews. Instability of the proportional fair scheduling algorithm for HDR. *IEEE Transactions on Wireless Communications*, 3(5):1422–1426, 2004.
- [3] M. Andrews, K. Kumaran, K. Ramanan, S. Stolyar, R. Vijayakumar, and P. Whiting. Scheduling in a

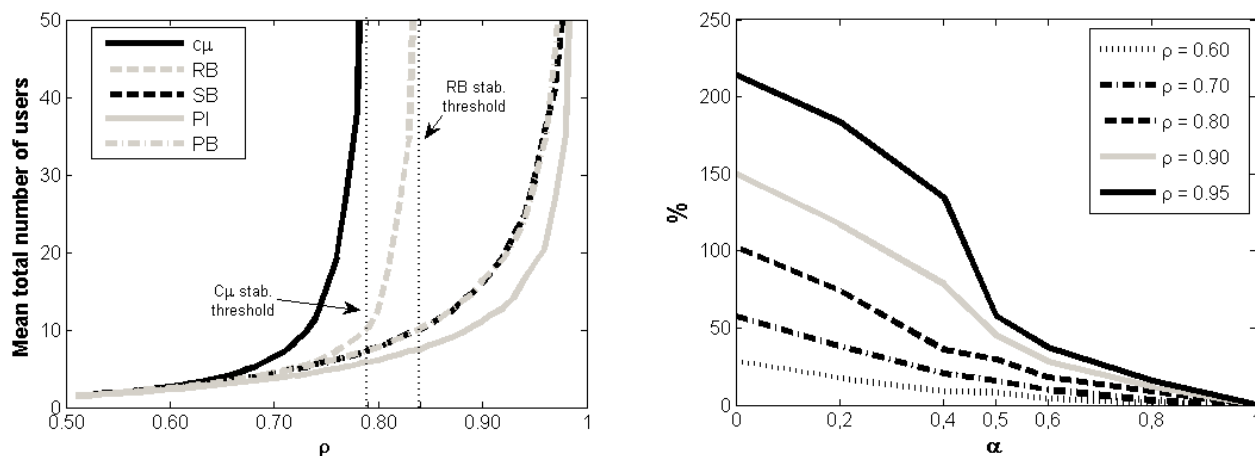


Figure 3: (a) Mean number of users and stability thresholds (b) PI under different tie-breaking rules: relative degradation (in %) over PI with $\alpha = 1$.

- queueing system with asynchronously varying service rates. *Probability in the Engineering and Informational Sciences*, 18(2):191–217, 2004.
- [4] U. Ayesta, M. Erausquin, and P. Jacko. A modeling framework for optimizing the flow-level scheduling with time-varying channels. *Performance Evaluation*, 67:1014–1029, 2010.
- [5] U. Ayesta, M. Erausquin, M. Jonckheere, and I.M. Verloop. Scheduling in a random environment: stability and asymptotic optimality. Available at *ArXiv*, 1101.5794v1, 2011.
- [6] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi. CDMA/HDR: A bandwidth-efficient high-speed wireless data service for nomadic users. *IEEE Communications Magazine*, 38(7):70–77, 2000.
- [7] F. Berggren and R. Jäntti. Asymptotically fair transmission scheduling over fading channels. *IEEE Transactions on Wireless Communications*, 3(1):326–336, 2004.
- [8] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1999.
- [9] T. Bonald. A score-based opportunistic scheduler for fading radio channels. In *Proceedings of European Wireless*, pages 283–292, 2004.
- [10] S.C. Borst. User-level performance of channel-aware scheduling algorithms in wireless data networks. *IEEE/ACM Transactions on Networking*, 13(3):636–647, 2005.
- [11] S.C. Borst, T. Bonald, N. Hegde, M. Jonckheere, and A. Proutiere. Flow-level performance and capacity of wireless networks with user mobility. *Queueing Systems*, 63(1–4):131–164, 2009.
- [12] S.C. Borst and M. Jonckheere. Flow-level stability of channel-aware scheduling algorithms. In *Proceedings of WiOpt*, September 2006.
- [13] S.C. Borst, M. Jonckheere, and L. Leskelä. Stability of parallel queueing systems with coupled service rates. *Discrete Event Dynamic Systems*, 18:447–472, 2008.
- [14] E.F. Chaponniere, P.J. Black, J.M. Holtzman, and D.N.C. Tse. Transmitter directed code division multiple access system using path diversity to equitably maximize throughput. US Patent 6,449,490.
- [15] J.G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability*, 5:49–77, 1995.
- [16] R.W.R. Darling and J.R. Norris. Differential equation approximations for markov chains. *Probability Surveys*, 5:37–79, 2008.
- [17] A. Eryilmaz, R. Srikant, and J. Perkins. Stable scheduling policies for fading wireless channels. *IEEE/ACM Transactions on Networking*, 13(2):411–424, 2005.
- [18] G. Fayolle, V.A. Malyshev, and M.V. Mensikov. *Topics in the Constructive Theory of Countable Markov Chains*. Cambridge University Press, 1994.
- [19] G. Fort, S. Meyn, E. Moulines, and P. Priouret. The ODE method for the stability of skip-free markov chains with applications to MCMC. *Annals of Applied Probability*, 18(2):664–707, 2008.
- [20] M. Hu, J. Zhang, and J. Sadowsky. Traffic aided opportunistic scheduling for wireless networks: algorithms and performance bounds. *Computer Networks*, 46(4):505–518, 2004.
- [21] P. Lassila and S. Aalto. Combining opportunistic and size-based scheduling in wireless systems. In *Proceedings of ACM MSWiM*, pages 323–332, 2008.
- [22] S. Liu, L. Ying, and R. Srikant. Throughput-optimal scheduling in the presence of flow-level dynamics. In *Proceedings of IEEE INFOCOM*, 2010.
- [23] X. Liu, E. Chong, and N. Shroff. Optimal opportunistic scheduling in wireless networks. In *Proceedings of VTC*, 2003.
- [24] M.J. Neely. Order optimal delay for opportunistic scheduling in multi-user wireless uplinks and downlinks. *IEEE/ACM Transactions on Networking*, 16(5):1188–1199, 2008.
- [25] D. Park, H. Seo, H. Kwon, and B. Lee. Wireless packet scheduling based on the cumulative distribution function of user transmission rates. *IEEE Transactions*

- on *Communications*, 53(11):1919–1929, 2005.
- [26] Ph. Robert. *Stochastic Networks and Queues*. Springer-Verlag, 2003.
- [27] B. Sadiq, S.-J. Baek, and G. de Veciana. Delay-optimal opportunistic scheduling and approximations: The log rule. In *Proceedings of IEEE INFOCOM*, 2009.
- [28] B. Sadiq and G. de Veciana. Balancing SRPT prioritization vs opportunistic gain in wireless systems with flow dynamics. In *Proceedings of ITC 22*, 2010.
- [29] B. Sadiq, R. Madan, and A. Sampath. Downlink scheduling of multiclass traffic in LTE. *EURASIP Journal on Wireless Communications and Networking*, 2009.
- [30] A.L. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Annals of Applied Probability*, 14(1):1–53, 2004.
- [31] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, 39(2):466–478, 1993.
- [32] P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*, 25:287–298, 1988.