

Scheduling in a random environment: stability and asymptotic optimality

U. Ayesta^{1,2,3}, M. Erausquin⁴, M. Jonckheere⁵, I.M. Verloop^{2,6}

¹CNRS-LAAS, Toulouse, France

²Université de Toulouse; INP, LAAS, Toulouse, France

³IKERBASQUE, Basque Foundation for Science, Bilbao, Spain

⁴UPV/EHU, University of the Basque Country, Bilbao, Spain

⁵CONICET, Buenos Aires, Argentina

⁶CNRS-IRIT, Toulouse, France

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 [8], Proportionally Best [1] or Potential Improvement [4] are stable under the maximum stability conditions, whereas the opportunistic scheduler Relative-Best [9] 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. We will refer to such tie-breaking rule as *myopic*. (3) We derive the growth rates of the number of users in the system in overload settings under various policies, which give additional insights on the performance. (4) We conclude that simple priority-index policies with the myopic tie-breaking rule, are stable and asymptotically optimal. All our findings are validated with extensive numerical experiments.

Index Terms—Wireless system, stability, asymptotic fluid optimality, opportunistic scheduling, $c\mu$ -rule.

I. 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 [9]). A policy is called opportunistic if it takes advantage of the channel fluctuations by serving a user whose channel condition is in a good state 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 [33], [2], [32], [17], [3], [25], [29] 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 [8], [21], [9], [26], [22], [1], [4], [30]. For surveys on flow-level modeling we refer to [24] and [10]. In [23], 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 [9], [11], [12], [29], [1], [30]) 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 [11] that any utility-based scheduling policy is stable in a flow-level model. The authors of [1] make the same assumption when they discuss rate-stability for priority-index policies. In the context of optimal control, in [29], [30] scheduling mechanisms were introduced and evaluated. In [4] optimal control is studied without the time-scale separation assumption. The Lagrangian-relaxation method 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 [8], the Relative Best (RB) scheduler [6] 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 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 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 investigate the system under the so-called fluid scaling. 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 much simpler to describe while they provide us with crucial insights on the behavior of the non-scaled version of the process.

First of all, we characterize the maximum stability conditions (the weakest possible conditions on the traffic parameters such that there exists a scheduling policy that makes the system stable) 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. These policies will be referred to as Best-Rate (BR) policies. Such a characterization was previously given for rate stability [1], but stochastic stability was still an open issue.

Second, for a large class of scheduling policies we determine the exact stability conditions and conclude that many known opportunistic scheduling policies like SB, PB, or PI are stable under the maximum stability conditions, whereas the opportunistic scheduler RB or $c\mu$ -rule are not.

Third, we demonstrate the importance of the choice of the tie-breaking rule. Until now, literature proposed to break ties at random [6], [8], [9], [1]. We instead propose to give priority to the user with highest instantaneous departure probability when there are multiple users in their best channel conditions, which we refer to as the myopic tie-breaking rule. We prove that BR policies with this tie-breaking rule are asymptotically optimal and our numerical experiments further illustrate that the myopic tie-breaking rule significantly improves the performance. This in turn shows that simple priority-index policies that balance opportunistic gains with size-based information, will be both maximum stable and asymptotically optimal.

Fourth, our convergence result allows to compare the performance of the various policies in an overload setting. More precisely, we determine the growth rates of the number of users in the various classes and find that BR policies with a myopic tie-breaking rule minimize the total growth rate.

The paper is organized as follows. In Sections II and III we present the model and the scheduling policies. In Section IV we derive fluid limits for a large class of policies. This allows to obtain our stability results as presented in Section V. In Section VI we characterize asymptotically optimal policies, both in normal regime and in overload, and discuss the importance of the tie-breaking rule. In Section VII we perform numerical experiments to validate our theoretical findings.

II. MODEL DESCRIPTION

We consider a time-slotted system serving one user in each time slot. This models for instance a CDMA 1xEV-DO system

as explained in Remark 2. 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 $\mathbb{E}(A_k) = \lambda_k$ and $\mathbb{E}(A_k^2) < \infty$. 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 are geometric (see Remark 2). W.l.o.g. we assume that the channel conditions are ordered such that $0 \leq \mu_{k,1} \leq \mu_{k,2} \leq \dots \leq \mu_{k,N_k} \leq 1$, and $q_{k,N_k} \mu_{k,N_k} \neq 0, \forall k$. 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 t , a scheduler/policy f decides which user is served. Because of the Markov property of the system, we focus on policies that base decisions on the current number of users present in the various classes and on their current channel states. For a given scheduling policy f , let $X_k^f(t)$ denote the number of class- k users present 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 $X^f(\cdot)$, the process X^f is Markov. Hence, it is sufficient to focus on the Markovian description in terms of the number of users in each *class*, $X^f(\cdot)$, instead of the number of users in each *channel state*.

Let us introduce some more notation. We denote by $\|x\|$ the l_1 norm of a vector x . The notation $x \leq y$ is used for the coordinate-wise ordering: $x_i \leq y_i, \forall i$. Finally, we denote by u.o.c. the uniform convergence on compact sets.

Performance criteria: Our performance criteria are stability and long-run average number of users.

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

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

the long-run time-average holding cost, with $c_k > 0$ the holding cost incurred per time slot for having a class- k user in the system. When $c_k = 1, \forall k$, this is equivalent to minimizing the mean sojourn time (cf. Little’s law).

Remark 1: From a modeling perspective, the key assumption needed for the subsequent fluid-limit analysis is that the




Fig. 1. Classification of schedulers.

channel state evolution is independent across users. Moreover, we assume that in each time slot a user's channel state is i.i.d.. However, we believe that a similar study could be carried out under a weaker assumption on a users' channel state process: stationary ergodicity. In this case, the steady-state probabilities of the channel process would play the role of the $q_{k,n}$'s in the current model description. In particular all the results on stability and asymptotic optimality that will be obtained later on would hold in this more general setting. For the sake of clarity and to avoid technicalities in the proofs we have restricted our attention to i.i.d. channel states.

Remark 2 (Modeling of a wireless data network): Our 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 [5] and the OFDM-based LTE systems [31]. 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 [5]).

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 . For the CDMA 1xEV-DO system, the amount of bits transferred in one slot is $\Delta = s_{k,n} \cdot t_c$. Hence, the departure probability of a class- k user under channel condition n 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 $t_c = 1.67ms$

in the CDMA 1xEV-DO system). In Section VII we perform numerical experiments with departure probabilities as obtained from a practical setting using (2).

III. POLICIES

In this section we introduce scheduling policies that will be used throughout the paper. Most of these policies are opportunistic (with the exception of the $c\mu$ -rule), meaning that they take advantage of channel fluctuations by serving a user whose channel condition is currently in a good state, in some sense, with respect to its own statistical behavior.

We first introduce priority-index policies, which 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, a user that has the highest index among all users present is 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 [8], [6], [1]). We define the *myopic tie-breaking rule* as the rule that among the users with the highest index, it selects the one with highest value for $c_k \mu_{k,N_k}$, $\forall k$ (the c_k 's refer to the holding cost introduced in Section II.) One of our main contributions will be that the choice for the tie-breaking rule is crucial for the performance of the system and that the myopic tie-breaking rule is close to optimal (this will be further developed in Sections VI and Section VII).

In [8] 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 $c_k \mu_{k,n} / \sum_{\tilde{n}>n} q_{k,\tilde{n}} (\mu_{k,\tilde{n}} - \mu_{k,n})$, and the tie-breaking rule is the myopic tie-breaking rule. 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 $\omega_k \mu_{k,n}$, with ω_k a class dependent weight.

Examples of weight-based policies are: the $c\mu$ -rule ($\omega_k = c_k$, with c_k the holding cost), Relative Best (RB) [6] ($\omega_k = 1/\sum_{n=1}^{N_k} q_{k,n}\mu_{k,n}$), and Proportionally Best (PB) [1] ($\omega_k = 1/\mu_{k,N_k}$). For all these policies, ties are broken at random.

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 (state N_k) such a user is served.

Definition 5 (Best Rate Priority (BRP) policies): The BRP policies are BR policies with a myopic tie-breaking rule.

As a consequence of our main results, we will obtain that the classes of policies BR and BRP have desirable properties: In Section V we prove that any BR policy is stable under the maximum stability conditions and in Section VI we derive that BRP policies are asymptotically optimal.

In Figure 1 we have summarized the various (classes of) policies. Note that 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. RB and the traditional $c\mu$ rule (i.e., giving in each time slot strict preemptive priority to the user having the highest $c_k\mu_{k,n}$) however do not belong to BR policies since, depending on the set of parameters, the index value of a class- k user in state $n \neq N_k$ might be larger than the index value of a class- l user in state N_l . Remark also that PI is the only BRP policy.

IV. FLUID LIMITS AND CONVERGENCE

In this section we study fluid-scaling limits for a large class of policies. The fluid scaling consists in studying a sequence of systems indexed by r , i.e., for a given policy f we let $X_k^{f,r}(t)$ 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 $r \in \mathbb{N}$, and $X^{f,r}(t) = (X_1^{f,r}(t), \dots, X_K^{f,r}(t))$. We are then interested in the fluid-scaled processes $Y_k^{f,r}(t) := \frac{X_k^{f,r}(\lfloor rt \rfloor)}{r}$, $t \geq 0$, $k = 1, \dots, K$, with $Y^r(0) = x(0)$. We can write

$$Y_k^{f,r}(t) = x_k(0) + \frac{1}{r} \sum_{s=1}^{\lfloor rt \rfloor} A_k(s) - \frac{1}{r} \sum_{n=1}^{N_k} S_{k,n}(T_{k,n}^{f,r}(rt)), \quad (3)$$

where $T_{k,n}^{f,r}(t)$ is defined as the cumulative amount of time that was spent on serving class k in state n during the interval $(0, t]$ and $S_{k,n}(t)$ denotes the total number of class- k users that have been completed while receiving service for a total duration of time t when being in state n .

In order to derive stability and asymptotic optimality results, we will be interested in limits of the fluid-scaled process. In Section IV-A we will characterize a generic description of weak fluid limits (usually not unique) of Equation (3), following the same reasoning as in [15]. In Section IV-B we focus instead on a special class of policies for which we can prove convergence in probability towards a *unique* limit, which will be referred to as the strong fluid limit. (In [20] similar is done but only for a subset of the state space.) We discuss the differences between weak and strong fluid limits in more detail in Remark 5, after having introduced formally both concepts.

A. Convergence towards weak fluid limits

From (3), we obtain the following result that describes the generic characterization of weak fluid limits for a given policy f . The lemma will allow to determine maximum stable policies (Theorem 5.2) and to characterize asymptotically optimal policies (Section IV).

Lemma 4.1: For almost all sample paths ω and any sequence r_k , there exists a subsequence r_{k_l} such that for all $k = 1, 2, \dots, K$, $n = 1, 2, \dots, N_k$, and $t \geq 0$,

$$\lim_{l \rightarrow \infty} Y_k^{f,r_{k_l}}(t) = y_k^f(t), \quad \text{u.o.c., and} \quad (4)$$

$$\lim_{l \rightarrow \infty} \frac{T_{k,n}^{f,r_{k_l}}(t)}{r_{k_l}} = \tau_{k,n}^f(t), \quad \text{u.o.c.,}$$

with $(y_k^f(\cdot), \tau_{k,n}^f(\cdot))$ a continuous function. In addition,

$$y_k^f(t) = x_k(0) + \lambda_k t - \sum_{n=1}^{N_k} \mu_{k,n} \tau_{k,n}^f(t), \quad (5)$$

$y_k^f(t) \geq 0$, $\tau_{k,n}^f(0) = 0$, $\sum_{k,n} \tau_{k,n}^f(t) \leq t$, and $\tau_{k,n}^f(\cdot)$ are non-decreasing and Lipschitz continuous functions.

Proof: The proof follows similarly to that of [13, Proposition 4.12]. Define $A_k(t)$ as the number of class- k users that arrive in time slot t . We note that $A_k(1), A_k(2), \dots$, are independent and distributed according to A_k , with $\mathbb{E}(A_k) = \lambda_k$. Since $\mu_{k,n}$ is the probability of completing a class- k user when it receives service while being at state n , by the law of large numbers, we obtain that, almost surely,

$$\lim_{r \rightarrow \infty} \frac{1}{r} \sum_{s=1}^{\lfloor rt \rfloor} A_k(s) = \lambda_k t, \quad \text{and} \quad \lim_{r \rightarrow \infty} \frac{1}{r} S_{k,n}(rt) = \mu_{k,n} t. \quad (6)$$

We define $\overline{T}_{k,n}^{f,r}(t) := T_{k,n}^{f,r}(rt)/r$. Since $T_{k,n}^{f,r}(t)$ denotes the cumulative amount of time spent on serving class- k users in state n in time interval $[0, t]$, we get $\overline{T}_{k,n}^{f,r}(rt) - \overline{T}_{k,n}^{f,r}(rs) \leq t - s$, for every $t \geq s$, i.e., $\overline{T}_{k,n}^{f,r}(t)$ is Lipschitz continuous. By the Arzela-Ascoli theorem [27] we obtain that, for almost every sample path w and any subsequence r_k , there exists a subsequence r_{k_l} of r_k such that $\lim_{l \rightarrow \infty} \overline{T}_{k,n}^{f,r_{k_l}}(rt)/r = \tau_{k,n}^f(t)$, u.o.c.. Now, using Equations (3) and (6), it follows that $\lim_{l \rightarrow \infty} Y_k^{f,r_{k_l}}(t) = y_k^f(t)$, with $y^f(t)$ as given in (5). ■

We can now give our definition of weak fluid limits.

Definition 6: We call the processes $\tau^f(t)$ and $y^f(t)$ (as obtained in Lemma 4.1) weak fluid limits for policy f .

Note that, in general, these fluid limits can be different depending on the sample path and the subsequence considered. A policy is said to have a *unique fluid limit* if, for all sample paths and all subsequences, the weak fluid limits coincide.

B. Convergence towards a strong fluid limit

In this subsection we will determine unique fluid limits for a special class of policies. More precisely, we will prove convergence in probability towards a unique limit. The derivation of the strong fluid limit will prove to be very useful: It allows to calculate the exact stability conditions (policy dependent) (see Theorem 5.1 and the numerical Section VII). In addition, the

exact characterization of the strong fluid limit provides crucial insights into the performance of the system, such as the growth rates of the number of users over time and monotonicity results with respect to the tie-breaking rule.

Obtaining exact fluid-limit characterizations will require to deal with averaging phenomena: it may happen that one class of users reaches its stationary regime, i.e., is empty in the fluid scaling, 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 as will be defined in (8) (we refer to this as second-vector fields, following [18]).

We focus on the class of policies that induce partially increasing drift vector fields with uniform limits. In order to describe this class of policies we need first to introduce the drift functions and drift vector fields. 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 [12] 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} and \mathcal{U}^c denotes the complementary set of \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 assume in the following that the drift vector has uniform limits, so that we can define the asymptotic drifts $\delta^\mathcal{U} : \mathbb{N}^{|\mathcal{U}|} \rightarrow \mathbb{R}^K$ as follows:

$$\delta^\mathcal{U}(x_\mathcal{U}) := \lim_{x_k \rightarrow \infty, k \in \mathcal{U}^c} \delta(x). \quad (7)$$

Here, \mathcal{U}^c corresponds to the ‘‘saturated’’ classes for which we let the number of users go to ∞ . We define the stochastic process $X^\mathcal{U}$ 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

$$\bar{\delta}^\mathcal{U} = \sum_{x \in \mathbb{N}^{|\mathcal{U}|}} \delta^\mathcal{U}(x) \pi^\mathcal{U}(x). \quad (8)$$

Finally, following [12] 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 unique strong fluid limit, are verified for many cases of interest, see the next lemma.

Lemma 4.2: 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.

By 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, $i \in \mathcal{U}^c$, grows large, the probability of having in each class in \mathcal{U}^c at least one user in its best state (and hence causing a drift $\delta^\mathcal{U}(x_\mathcal{U})$) converges to 1. Together with the property that the tie-breaking rule does not depend on the number of users, this implies that

$$\delta(x) = \delta^\mathcal{U}(x_\mathcal{U}) \prod_{i \in \mathcal{U}^c} (1 - (1 - q_{i, N_i})^{x_i}) + o(1/|x|), \quad (9)$$

as $x_k \rightarrow \infty, k \in \mathcal{U}^c$, which in turn implies the uniform convergence of $\delta(\cdot)$ to $\delta^\mathcal{U}(\cdot)$ for any \mathcal{U} . ■

We now state the main result of this Section, the description of the strong fluid limit. The proof can be found in Appendix A.

Theorem 4.3: For a given policy f inducing a partially increasing drift vector field with uniform limits, we have

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

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,

$$\frac{dy_k^f(t)}{dt} = \tilde{\delta}_k^{f, \mathcal{U}_l}, \quad t \in [T_l^f, T_{l+1}^f], \quad (10)$$

$$\text{with } T_{l+1}^f = T_l^f + \min_{k \in \mathcal{U}_l^c, \tilde{\delta}_k^{f, \mathcal{U}_l} < 0} \frac{y_k^f(T_l^f)}{-\tilde{\delta}_k^{f, \mathcal{U}_l}}, \quad (11)$$

$$\text{and } \mathcal{U}_{l+1} = \mathcal{U}_l \cup \arg \min_{k \in \mathcal{U}_l^c, \tilde{\delta}_k^{f, \mathcal{U}_l} < 0} \frac{y_k^f(T_l^f)}{-\tilde{\delta}_k^{f, \mathcal{U}_l}},$$

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

We can now give our definition of the strong fluid limit.

Definition 7: For a given policy f inducing a partially increasing drift vector field with uniform limits, we call the process $y^f(t)$ (as obtained in Theorem 4.3) the strong fluid limit for the policy f .

Remark 3 (Calculation of the averaged drifts):

Theorem 4.3 characterizes the strong fluid limit as a piece-wise linear function with slopes $\bar{\delta}^\mathcal{U}$. In practice, the calculations of these slopes involve:

- deriving the asymptotic drifts (see (7)),
- calculating the stationary distributions of $X^\mathcal{U}$,
- averaging the asymptotic drifts with these stationary distributions (see (8)).

For instance, assume $K = 2, N_1 = 2, N_2 = 1$, 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 (8).

In the specific case of BR policies with a priority-type tie-breaking rule, we can in fact explicitly derive the strong fluid limit by making use of rate-conservation arguments. This will be used later on to obtain asymptotic optimality statements.

Proposition 4.4: Consider a BR policy with a priority-type tie-breaking rule. Let us reorder the classes according to the priority ordering. The averaged drift vectors are

$$\tilde{\delta}^{\emptyset} = (\lambda_1 - \mu_{1,N_1}, \lambda_2, \dots, \lambda_K), \quad (12)$$

if $T_1 < \infty$, then

$$\tilde{\delta}^{\{1\}} = (0, \lambda_2 - \mu_{2,N_2} \left(1 - \frac{\lambda_1}{\mu_{1,N_1}}\right), \dots, \lambda_K), \quad (13)$$

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

$$\tilde{\delta}^{\mathcal{U}_{k-1}} = (0, \dots, 0, \lambda_k - \mu_{k,N_k} \left(1 - \sum_{j=1}^{k-1} \frac{\lambda_j}{\mu_{j,N_j}}\right), \dots, \lambda_K). \quad (14)$$

Proof: Using Lemma 4.2, the drift $\delta(\cdot)$ associated to a BR policy is partially increasing with uniform limits, hence the strong fluid limit is given by Theorem 4.3. 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 (12). Note that $T_1 < \infty$ if and only if $\frac{\lambda_1}{\mu_{1,N_1}} < 1$ (see definition of T_1 in Theorem 4.3). 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 and there are x_1 class-1 users when all the other classes are saturated. With a slight abuse of notation, let us denote by $\pi^{\{1\}}(A_{1,x_1})$ the probability of event A_{1,x_1} under the stationary distribution of $X^{\{1\}}$. Since $X^{\{1\}}$ is ergodic, by rate stability we have the following rate-conservation equation, $\sum_{x_1} \pi^{\{1\}}(A_{1,x_1})\mu_{1,N_1} = \lambda_1$, see (44). This gives that $\sum_{x_1} \pi^{\{1\}}(A_{1,x_1}^c) = 1 - \frac{\lambda_1}{\mu_{1,N_1}}$. Since $\mathcal{U}_1 = \{1\}$ (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 $\mu_{2,N_2} \sum_{x_1} \pi^{\{1\}}(A_{1,x_1}^c) = \mu_{2,N_2} \left(1 - \frac{\lambda_1}{\mu_{1,N_1}}\right)$ which gives (13).

Consider now the case where $\mathcal{U} := \{1, \dots, k-1\}$ (assuming as before that $\sum_{j=1}^{k-1} \frac{\lambda_j}{\mu_{j,N_j}} < 1$). Let $A_{j,x_{\mathcal{U}}}$ be the event that class $j \in \mathcal{U}$ is served and there are x_i users of class i , $i \in \mathcal{U}$. By rate-conservation arguments, see (44), 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, $\forall j \in \mathcal{U}$, this implies that $\sum_{x_{\mathcal{U}}} \pi^{\mathcal{U}}(\cup_{j \in \mathcal{U}} A_{j,x_{\mathcal{U}}}) = \sum_{j \in \mathcal{U}} \frac{\lambda_j}{\mu_{j,N_j}}$. Since class k is only served when no class- i users are being served, $i \in$

\mathcal{U} , there is a class- k departure with probability $\mu_{k,N_k}(1 - \sum_{x_{\mathcal{U}}} \pi^{\mathcal{U}}(\cup_{j \in \mathcal{U}} A_{j,x_{\mathcal{U}}}))$. Hence, we obtain Equation (14). ■

Remark 4: For all BR policies where the scheduler chooses with probability $\alpha_k^{\mathcal{U}}$ to serve class k 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:

$$\delta^{\emptyset} = (\lambda_1 - \alpha_1^{\{1, \dots, K\}} \mu_{1,N_1}, \lambda_2 - \alpha_2^{\{1, \dots, K\}} \mu_{2,N_2}, \dots, \lambda_K - \alpha_K^{\{1, \dots, K\}} \mu_{K,N_K}). \quad (15)$$

However, in general we cannot explicitly derive the second-vector fields. An exception is the case of 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)

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

Remark 5 (Weak and strong fluid limits): Though a quite subtle technical point, it is worth emphasizing the conceptual difference between the notion of weak fluid limits and the notion of strong fluid limit, introduced in Sections IV-A and IV-B, respectively. Note that “weak” versus “strong” refers to accumulation points versus unique limit. The names do however not take into account the mode of convergence.

Weak limits are a powerful tool for stability if one can characterize that they all vanish after a finite amount of time, as will be used in Theorem 5.2 for the set of BR policies. However, in general weak limits might not capture the precise asymptotic behavior of the process (see [13]). On the contrary, when having a unique strong fluid limit, the asymptotic behavior of the scaled process is completely described, allowing for example to obtain the policy-dependent stability conditions for a large class of policies (see Theorem 5.1).

V. STABILITY ANALYSIS

The derivation of the weak and strong fluid limits in Section IV allows us to conclude the stochastic stability results. The first Theorem derives the stability conditions for any policy having a partially increasing drift with uniform limits, using the strong fluid limit as obtained in Theorem 4.3.

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 4.3.

Proof: If $T_l^f < \infty$ for all l , the strong fluid limit described in Theorem 4.3 is equal to 0 for t large enough, i.e., $Y^{f,r}(t)$ 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 $x_k(0)$ plus the users that have arrived until time $\lfloor rt \rfloor$ divided by r , which is uniform integrable, see [15, Lemma 4.5].

The convergence in probability to 0 and the uniform integrability together imply that $\lim_{r \rightarrow \infty} \mathbb{E}(Y_k^{f,r}(t)) = 0$, for t large enough, $\forall k$, see [7, Theorem 3.5]. Using an extended Foster-Lyapunov criterion as expressed in [18] or [28, Corollary 9.8], this implies the positive recurrence of $X^f(\cdot)$. ■

In the following theorem we state the maximum stability condition, and prove that any BR policy achieves maximum

stability. The proof is based on the weak fluid limit characterization as given in Section IV and can be found in Appendix B.

Theorem 5.2: The maximum stability condition is

$$\sum_{k=1}^K \frac{\lambda_k}{\mu_{k,N_k}} < 1. \quad (16)$$

In addition, any BR policy is maximum stable.

Condition (16) was recognized as the maximal rate stability condition in [1] and as the maximum stability condition under a time-scale separation assumption in [9].

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 VII where we numerically obtain the stability conditions for RB and the $c\mu$ -rule, making use of Theorem 5.1.

Remark 6 (Overload): When $\sum_{k=1}^K \lambda_k / \mu_{k,N_k} > 1$ the system is said to be in overload. That is, there does not exist any policy that can make the system stable. Theorem 4.3 is however still applicable, providing us with the rates at which the number of users in the different classes grow: given $x(0) = 0$, the growth rate of the number of users over time is given by $X^{f,r}(r)/r = Y^{f,r}(1)$, which in the limit is equal to $\delta^{f,\mathcal{U}}$. This gives us a mean to compare the performance of various policies in overload (see as well Sections VI and VII).

VI. 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, i.e.,

$$\min_u \sum_{k=1}^K c_k x_k^u(t), \text{ for all } t \geq 0, \text{ subject to} \quad (17)$$

$$x_k^u(t) = x_k(0) + \lambda_k t - \sum_{n=1}^{N_k} \mu_{k,n} \int_0^t u_{k,n}(v) dv, \quad (18)$$

$$x_k^u(t) \geq 0, \quad k = 1, \dots, K, \quad (19)$$

$$\sum_{k=1}^K \sum_{n=1}^{N_k} u_{k,n}(v) \leq 1, \quad u_{k,n}(v) \geq 0, \quad \forall k, n, v \geq 0, \quad (20)$$

and the control functions $u_{k,n}(v)$ being integrable. Here $x_k^u(t)$ represents the amount of fluid in class k under control $u(\cdot)$.

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 $x_k(t)$ should be interpreted as a limit of the fluid-scaled process. Hence, when $x_k(t) > 0$ this implies that there are infinitely many class- k users so that with probability 1 there are class- k users in each of the channel state conditions (this being independent of the exact values of the $q_{k,n} > 0$'s).

An optimal control $u^*(\cdot)$ is 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 $u^*(\cdot)$ that solves the fluid control problem is as follows. Let $l = \arg \min \{k : x_k(t) > 0\}$. Then

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

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

Proof: Let us denote $w_k^u(t) = x_k^u(t) / \mu_{k,N_k}$. First, we show that for any feasible control $u(\cdot)$, we have

$$\sum_{k=1}^j w_k^u(t) \leq \sum_{k=1}^j w_k^*(t), \text{ for all } t \geq 0, \quad j = 1, \dots, K. \quad (21)$$

If $\sum_{k=1}^j w_k^*(t) = 0$, then (21) trivially holds. Now assume $\sum_{k=1}^j w_k^*(t) > 0$. By definition of $u^*(\cdot)$ this implies that $\sum_{k=1}^j w_k^*(s) > 0$, for all $s \in [0, t]$, since once all these classes empty under $u^*(t)$, they will remain empty. Since $u^*(t)$ gives full priority to classes 1 until j over classes $j+1$ until K , we have that $\sum_{k=1}^j \int_0^t u_{k,N_k}^*(v) dv = t$. Hence, $\sum_{k=1}^j \int_0^t u_{k,N_k}^*(v) dv = t \geq \sum_{k=1}^j \sum_{n=1}^{N_k} \int_0^t u_{k,n}(v) dv \geq \sum_{k=1}^j \sum_{n=1}^{N_k} \int_0^t \frac{\mu_{k,n}}{\mu_{k,N_k}} u_{k,n}(v) dv$, which implies (21), since $\sum_{k=1}^j w_k^u(t) - \sum_{k=1}^j w_k^*(t) = \sum_{k=1}^j \int_0^t u_{k,N_k}^*(v) dv - \sum_{k=1}^j \sum_{n=1}^{N_k} \int_0^t \frac{\mu_{k,n}}{\mu_{k,N_k}} u_{k,n}(v) dv \geq 0$.

The minimization term $\sum_{k=1}^K c_k x_k^u(t)$ can be written as $\sum_{k=1}^K c_k \mu_{k,N_k} w_k^u(t) = (c_1 \mu_{1,N_1} - c_2 \mu_{2,N_2}) w_1^u(t) + \dots + (c_{K-1} \mu_{K-1,N_{K-1}} - c_K \mu_{K,N_K}) \sum_{k=1}^{K-1} w_k^u(t) + c_K \mu_{K,N_K} \sum_{k=1}^K w_k^u(t)$. Together with (21) and $c_j \mu_{j,N_j} - c_{j+1} \mu_{j+1,N_{j+1}} \geq 0, \forall j$, we obtain that $\sum_{k=1}^K c_k \mu_{k,N_k} w_k^u(t)$ is minimized by $u^*(\cdot)$. ■

The optimal fluid cost serves as a lower bound for the fluid-scaled cost of the stochastic network, see the lemma below.

Lemma 6.2: For any policy f and for almost all sample paths, we have

$$\liminf_{r \rightarrow \infty} \sum_{k=1}^K c_k Y_k^{f,r}(t) \geq \sum_{k=1}^K c_k x_k^*(t), \text{ for all } t \geq 0. \quad (22)$$

Proof: Lemma 4.1 states that for almost all sample paths we have $\liminf_{r \rightarrow \infty} Y_k^{f,r}(t) = y_k^f(t)$, with $y_k^f(t)$ a weak fluid limit for policy f (this follows by considering the subsequence r_l corresponding to the liminf-sequence in Lemma 4.1). Note that a weak fluid limit is an admissible trajectory for the fluid control problem. Hence, $\liminf_{r \rightarrow \infty} \sum_{k=1}^K c_k Y_k^{f,r}(t) = \sum_{k=1}^K c_k y_k^f(t) \geq \sum_{k=1}^K c_k x_k^*(t)$. ■

Since (22) holds almost surely, it follows by Fatou's lemma that the lower bound holds in probability as well, i.e.,

$$\mathbb{P}\left(\sum_{k=1}^K c_k Y_k^{f,r}(t) - \sum_{k=1}^K c_k x_k^{u^*}(t) \geq 0\right) \rightarrow 1, \quad \text{for all } t \geq 0. \quad (23)$$

We define a policy to be *asymptotically optimal* when the lower bound is obtained in probability, i.e., $\forall \epsilon > 0$, $\lim_{r \rightarrow \infty} \mathbb{P}(|\sum_{k=1}^K c_k (Y_k^{f,r}(s) - x_k^{u^*}(s))| \geq \epsilon) = 0$.

The next Theorem characterizes a class of policies that is asymptotically optimal.

Theorem 6.3: Any BRP policy is asymptotically optimal.

Proof: We have $\frac{dx_k^{u^*}(t)}{dt} = \lambda_k - u_{k,N_k}^*(t)\mu_{k,N_k}$, with $u^*(\cdot)$ an optimal control as derived in Lemma 6.1. This drift coincides with the drift of the strong fluid limit $y^{BRP}(\cdot)$, see Proposition 4.4, hence $y^{BRP}(t) = x^{u^*}(t)$. Together with Theorem 4.3, we obtain that $\lim_{r \rightarrow \infty} \sum_{k=1}^K c_k Y_k^{BRP,r}(t)$ converges in probability to $\sum_{k=1}^K c_k x_k^{u^*}(t)$. ■

It can be checked that the above implies that any BRP policy minimizes $\liminf_{r \rightarrow \infty} \mathbb{E}(\int_0^\infty \sum_{k=1}^K c_k Y_k^{f,r}(t) dt)$. Unfortunately, this does not give any performance guarantee in terms of the long-run time-average holding cost as in Equation (1). Numerical experiments reported in Section VII indicate however that BRP policies significantly outperform all other policies in terms of the long-run time-average holding cost.

Note that the optimality results described in this section also apply in overload systems. We have the following corollary for the total growth rate.

Corollary 6.4: Any BRP policy minimizes the growth rate of the total cost, i.e., for all $\epsilon > 0$,

$$\lim_{r \rightarrow \infty} \mathbb{P}\left(\frac{\sum_{k=1}^K c_k X_k^{r,f}(r)}{r} - \frac{\sum_{k=1}^K c_k X_k^{r,BRP}(r)}{r} \geq -\epsilon\right) = 1.$$

Proof: Combining (23) with the asymptotic optimality of a BRP policy, the statement is immediate. ■

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. SB and PB will as well become asymptotically optimal when the myopic tie-breaking rule would be applied, showing the importance of the tie-breaking rule.

Remark 7: From Theorem 6.3 we conclude that the myopic tie-breaking rule is crucial in order to obtain an asymptotically optimal scheduling policy. In view of Equation (2), we note that under this myopic rule, higher priority is given to users with smaller service requirements, $\mathbb{E}(B_k)$. Hence, BRP policies appropriately mix size-based information with achieving opportunistic gains. This is in agreement with the findings of [30] where the authors investigate the tradeoff between prioritizing small users and opportunistic scheduling: They show that *if* the opportunistic capacity is upper bounded and increases as $1 - a^x$, with $a \in [0, 1)$ and x the number of users, then a significant improvement of performance can

be achieved by exploiting information on the service time requirement. In our model the capacity has this behavior, see for instance Equation (9), and as will be observed in the numerical results, exploiting size-based information indeed allows to obtain significant improvements.

Remark 8: In the asymptotic regime, the optimal tie-breaking rule is to give priority based on μ_{k,N_k} . However, outside this regime the optimal tie-breaking rule will depend both on the μ_{k,N_k} 's and on the probabilities that users are in their best channel condition, the $q_{N_k,k}$'s. The reason for this is that in the asymptotic regime there are many users in the system, which is why the q_{k,N_k} 's become irrelevant in the limit (as the number of users grows to infinite, the probability of having one best-rate user becomes equal to 1). Note however that our numerical simulations in the next section suggest that outside the asymptotic regime the asymptotically optimal tie-breaking rule performs very well as well.

VII. 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 I (taken from [5]).

We simulate the queueing system as described in Section II. 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 I. 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 expected total number of users in the system, see Equation (1). We will compare the performance of the policies SB, RB, PI, PB and the $c\mu$ rule, which were introduced in Section III.

Before presenting the numerical results we first summarize the main conclusions that we will make in this section:

- For different policies we validate the convergence of the system towards its fluid limit.
- We prove that not all policies obtain maximum stability. More precisely, we calculate the stability conditions under RB and the $c\mu$ -rule and observe that these are much more stringent than the stability condition for BR policies.
- The drifts of the fluid limit, δ^U , which are calculated numerically (and in some cases theoretically), provide very important insights on the performance. In particular, insightful monotonicity results in the tie-breaking rule are obtained with respect to the performance of the system. In addition, the drift analysis allows to show that PI outperforms all other policies in terms of the growth rates.
- Our simulations illustrate that the tie-breaking rule has a big impact on the performance of the system and we find that BRP policies, i.e., policies that combine opportunistic scheduling with the myopic tie-breaking rule, minimize the long-run time-average holding cost (see (1)).

Channel state (n)	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 I
TRANSMISSION RATES AND CHANNEL CONDITION PROBABILITIES IN THE CDMA 1xEV-DO WIRELESS NETWORK, AS REPORTED IN [5].

a) *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 4.3. When both classes are saturated, i.e., $\mathcal{U} = \emptyset$, the drift is

$$\delta^{f,\emptyset} = (\lambda_1 - \alpha^f \mu_{1,N_1}, \lambda_2 - (1 - \alpha^f) \mu_{2,N_2}), \quad (24)$$

see Remark 4. 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 user, i.e., there occur no ties, hence one can set $\alpha^f = 1$ in (24) for $f = c\mu, RB$. For PI, SB and PB we do have ties, and we set $\alpha^{PI} = 1$ (since PI applies the myopic tie-breaking rule) and $\alpha^{SB} = \alpha^{PB} = 1/2$ (since SB and PB apply a random tie-breaking rule). In Table II we present the so-obtained values for $\delta^{f,\emptyset}$. 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 4.3 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 depends on the policy. From Proposition 4.4 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 3, see Table II. In particular we observe that these drifts are positive for the latter two policies, which implies instability of the system, as can be seen in Figure 2 b). 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.

We note that the slopes of the strong fluid limit indeed coincide with the slopes of the original process, $Y^r(t)$ ($r = 10000$), as plotted in Figure 2.

A direct consequence of the drift function is that SB, PB, and PI (in fact all BR policies) empty the system at the same time (under the maximum stability condition), i.e., T_2^f is the same (this can be seen directly from Equation (43) for example). 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.

f	$\delta^{f,\emptyset}$		$\delta^{f,\{1\}}$	
	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 II
DRIFT OF THE FLUID LIMIT.

f	$\delta^{f,\emptyset}$		$\delta^{f,\{1\}}$	
	Class 1	Class 2	Class 1	Class 2
PI	-0.16	0.05	0	0.01
$c\mu$ -rule	-0.16	0.05	0	0.036
PB/SB	0.04	0	-	-
RB	-0.16	0.05	0	0.029

TABLE III
DRIFT OF THE FLUID LIMIT IN OVERLOAD SYSTEM.

b) *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 numerically calculated as follows: First, we calculate the stationary distribution of class-1 users once the stationary regime is reached. Then, we calculate the average drift $\bar{\delta}^{f,\{1\}}$ according to Definition (8). Policy f is stable if $\bar{\delta}^{f,\{1\}} < 0$, which reduces to $\rho < 0.79$ for the $c\mu$ -rule and to $\rho < 0.84$ for RB. In Figure 3 a) we plot the mean number of users for different values of ρ and we observe that the mean number of users (and hence the mean delay) for these policies grows to infinity as the load ρ approaches the critical value.

c) *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 (Recall that the parameter α is the probability, in case of a tie, that class 1 is favoured over class 2). 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) over PI as we vary α . We observe that the degradation in the mean performance is decreasing as α




Fig. 2. (a) Scaled number of class-1 users, (b) Scaled number of class-2 users, (c) Scaled total number of users




Fig. 3. (a) Mean number of users and stability thresholds, (b) PI under different tie-breaking rules: relative degradation (in %) over PI with $\alpha = 1$, (c) Scaled total number of users in overload, $\rho = 1.1$.

increases. In the fluid-scaling system a similar observation can be made: from (24) and Remark 4 it follows that the drift of the total number of users of the strong fluid limit (as obtained in Theorem 4.3) is decreasing (or constant) in α , for all time t .

The results show that the myopic tie-breaking rule, which was proven to be asymptotically optimal in the fluid limit (see Theorem 6.3), is in practice indeed 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%.

d) Overload: In Figure 3 c) we plot a trajectory of the total scaled number of users $Y_1^f(t) + Y_2^f(t)$ (for $r = 100$) when $\lambda_1 = 0.240$ and $\rho = 1.1$, so all policies are unstable. In Table III we give the values for the drifts (growth rates) of the fluid limit $y^f(t)$. In this example, the worst performance is under SB/PB and the best performance is for PI. This is in contrary to the stable regime where SB/PB is maximum stable with a performance strictly better than the $c\mu$ -rule and RB, see Figure 3 a). This implies that the performance of this policy can differ very much between stable and overload regimes. In addition, the total growth rate of the system is always minimized under any BRP policy, that is a best-rate policy with the myopic tie-breaking rule, as proved in Corollary 6.4.

VIII. 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 $c\mu$ 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.

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 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] 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.
- [6] F. Berggren and R. Jäntti. Asymptotically fair transmission scheduling over fading channels. *IEEE Transactions on Wireless Communications*, 3(1):326–336, 2004.
- [7] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1999.
- [8] T. Bonald. A score-based opportunistic scheduler for fading radio channels. In *Proceedings of European Wireless*, pages 283–292, 2004.

- [9] 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.
- [10] 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.
- [11] S.C. Borst and M. Jonckheere. Flow-level stability of channel-aware scheduling algorithms. In *Proceedings of WiOpt*, September 2006.
- [12] 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.
- [13] M. Bramson. *Stability of queueing networks*. Springer-Verlag, 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] W. Feller. *An Introduction to Probability Theory and Its Application*, volume 1. Wiley & sons, 1961.
- [20] 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.
- [21] 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.
- [22] P. Lassila and S. Aalto. Combining opportunistic and size-based scheduling in wireless systems. In *Proceedings of ACM MSWiM*, pages 323–332, 2008.
- [23] S. Liu, L. Ying, and R. Srikant. Throughput-optimal scheduling in the presence of flow-level dynamics. In *Proceedings of IEEE INFOCOM*, 2010.
- [24] X. Liu, E. Chong, and N. Shroff. Optimal opportunistic scheduling in wireless networks. In *Proceedings of VTC*, 2003.
- [25] 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.
- [26] 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.
- [27] M. Renardy and R.C. Rogers. *An introduction to partial differential equations*. Springer-Verlag, 2004.
- [28] Ph. Robert. *Stochastic Networks and Queues*. Springer-Verlag, 2003.
- [29] B. Sadiq, S.-J. Baek, and G. de Veciana. Delay-optimal opportunistic scheduling and approximations: The log rule. In *Proceedings of IEEE INFOCOM*, 2009.
- [30] B. Sadiq and G. de Veciana. Balancing SRPT prioritization vs opportunistic gain in wireless systems with flow dynamics. In *Proceedings of ITC 22*, 2010.
- [31] B. Sadiq, R. Madan, and A. Sampath. Downlink scheduling of multiclass traffic in LTE. *EURASIP Journal on Wireless Communications and Networking*, 2009.
- [32] 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.
- [33] 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.
- [34] D. Williams. *Probability with martingales*. Cambridge University Press, 1991.

APPENDIX A: PROOF OF THEOREM 4.3

We first derive a property on the drift vector.

Lemma 8.1: Suppose $\delta^f(\cdot)$ has uniform limits. Then for any \mathcal{U} and any compact set $\mathcal{C} \subset \mathbb{R}_{++}^{|\mathcal{U}^c|}$ there exists an $L > 0$ and

ϵ_r with $\lim_{r \rightarrow \infty} \epsilon_r = 0$ such that

$$\begin{aligned} |\delta^f(x) - \delta^{\mathcal{U},f}(x_{\mathcal{U}})| &\leq \epsilon_r + L \left| \frac{x_{\mathcal{U}^c}}{r} - y \right|, \\ \forall x \in \mathbb{R}_+^K, x_{\mathcal{U}^c} \in \mathbb{R}_{++}^{|\mathcal{U}^c|}, y \in \mathcal{C} \subset \mathbb{R}_{++}^{|\mathcal{U}^c|}. \end{aligned} \quad (25)$$

Proof: Consider $\epsilon > 0$, and $y \in \mathcal{C}$ such that $y_i > \epsilon, \forall i \in \mathcal{U}^c$. Define $\mathcal{B}_{ry, r\epsilon} \subset \mathbb{R}_{++}^{|\mathcal{U}^c|}$ the ball of center ry and radius $r\epsilon$. Assume that the drift function $\delta(\cdot)$ has uniform limits. Hence, $\epsilon_r = \sup_{y \in \mathcal{C}} \sup_{x_{\mathcal{U}^c} \in \mathcal{B}_{ry, r\epsilon}} |\delta^f(x) - \delta^{\mathcal{U},f}(x_{\mathcal{U}})|$, converges to 0, as $r \rightarrow \infty$. Since $|\frac{x_{\mathcal{U}^c}}{r} - y| \geq \epsilon$ when $x_{\mathcal{U}^c} \notin \mathcal{B}_{ry, r\epsilon}$, we obtain that $|\delta^f(x) - \delta^{\mathcal{U},f}(x_{\mathcal{U}})| \leq \sup_{x_{\mathcal{U}^c} \in \mathcal{B}_{ry, r\epsilon}} |\delta^f(x) - \delta^{\mathcal{U},f}(x_{\mathcal{U}})| + \sup_{x_{\mathcal{U}^c} \notin \mathcal{B}_{ry, r\epsilon}} |\delta^f(x) - \delta^{\mathcal{U},f}(x_{\mathcal{U}})| \leq \epsilon_r + L \left| \frac{x_{\mathcal{U}^c}}{r} - y \right|$, with $L := \frac{2 \max_{k,n} (\lambda_k, \mu_{k,n})}{\epsilon} \geq \frac{\sup_{x_{\mathcal{U}^c} \notin \mathcal{B}_{ry, r\epsilon}} |\delta^f(x) - \delta^{\mathcal{U},f}(x_{\mathcal{U}})|}{\epsilon}$. ■

In what follows, we will use the martingale decomposition of the process (see [34]), i.e., we define $M(\cdot)$ such that

$$X(t) = X(0) + \sum_{s=0}^{t-1} \delta(X(s)) + M(t), \quad (26)$$

where M is a martingale, and $|M(0)| = 0$. The following lemma gives a useful property for the martingale $M(t)$.

Lemma 8.2: It holds that $\mathbb{E}(|M_k(t)|^2) \leq C't, \forall k = 1, \dots, K$, for some constant $C' > 0$.

Proof: Using that the drift is bounded and the arrival increments have a finite variance, it immediately follows from (26) that $\mathbb{E}(|M_k(t)|^2) < \infty, \forall t$. Hence,

$$\begin{aligned} \mathbb{E}(|M_k(t)|^2) & \\ &\leq \mathbb{E}(|M_k(0)|^2) + \sum_{s=1}^t \mathbb{E}(|(M_k(s) - M_k(s-1))|^2), \end{aligned} \quad (27)$$

see [34, Chapter 12.1]. From (26), we have that for all s , $|M_k(t) - M_k(t-1)| \leq |X_k(t) - X_k(t-1)| + |\delta(X_i(t-1))|$. Since the drift is bounded, and the arrival process has finite mean and variance, this implies that $\mathbb{E}(|M_k(t) - M_k(t-1)|^2) \leq C'$, with $C' > 0$ a constant. Thus, the statement follows from (27). ■

We prove Theorem 4.3 for two classes of users, the general case being notationally much more cumbersome but mathematically equivalent. The trajectory of the fluid limit is defined piece-wise. We refer as the *first part of the fluid limit* to the case $t < T_1$, and as *second part of the fluid limit* to the case $T_1 \leq T_2$, and we shall prove the convergence separately for the two parts. The basic techniques used to prove the convergence of the first part of the trajectory are fairly standard and we proceed using the tools of [20]. Let us however emphasize that the proof relies crucially on the extended Lipschitz condition (25). Verifying this condition might be difficult in practice. In Lemma 8.1, combined with Lemma 4.2, we however showed that the conditions are satisfied for various scheduling policies of interest. Proving the convergence of the second part of the trajectory is much more subtle as it involves an averaging phenomena: one of the classes reaches a stationary regime while the other class keeps being macroscopically big. We must therefore rely on

stochastic comparisons and on the ergodic theorem for Markov chains. Such phenomena have been observed and analyzed in [18] for homogeneous random walks in the positive orthant.

First part of the trajectory: Assume strictly positive initial conditions for both classes, i.e., $x_1(0) > 0$ and $x_2(0) > 0$. At time $t = 0$ all classes are saturated, hence in the first part of the trajectory, $t < T_1$, we need to consider $\mathcal{U} = \mathcal{U}_0 = \emptyset$. Consider the deterministic function $y(\cdot)$ as defined in the theorem, $y(t) = x(0) + \delta^\theta t$, and we define an error function $e^r(t) := \mathbb{E}(\sup_{0 \leq s \leq t} |Y^r(s) - y(s)|)$. From (26) we obtain

$$|Y^r(t) - y(t)| \leq \frac{1}{r} \sum_{s=0}^{\lfloor rt \rfloor - 1} \left| \delta(X^r(s)) - \delta^\theta \right| + \left| \frac{M^r(\lfloor rt \rfloor)}{r} \right|. \quad (28)$$

We first control the martingale. Using successively Cauchy-Schwartz, Doob's inequality and Lemma 8.2, we get

$$\begin{aligned} \mathbb{E} \left(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \right)^2 &\leq \mathbb{E} \left(\left(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \right)^2 \right) \\ &\leq 4\mathbb{E} \left(\left| \frac{M^r(\lfloor rt \rfloor)}{r} \right|^2 \right) \leq \frac{C't}{r}, \text{ for a } C' > 0. \end{aligned} \quad (29)$$

We now control the drift part in Equation (28). Since the drift has uniform limits, we obtain from Lemma 8.1 that $|\delta(X^r(s)) - \delta^\theta| \leq \epsilon_r + L \left| \frac{X^r(s)}{r} - y(\frac{s}{r}) \right|$, $\forall s \leq rt$. Hence, bounding the Riemann sum for piece-wise constant functions by an integral, there exists a $\kappa > 0$ such that

$$\begin{aligned} \frac{1}{r} \sum_{s=0}^{\lfloor rt \rfloor - 1} \left| \delta(X^r(s)) - \delta^\theta \right| &\leq \epsilon_r t + \kappa \int_0^t \left| \frac{X^r(rs)}{r} - y(s) \right| ds \\ &\leq \epsilon_r t + \kappa \int_0^t \sup_{0 \leq \tilde{s} \leq s} \left| \frac{X^r(r\tilde{s})}{r} - y(\tilde{s}) \right| ds. \end{aligned} \quad (30)$$

Taking expectations on both sides we obtain

$$\mathbb{E} \left(\frac{1}{r} \sum_{s=0}^{\lfloor rt \rfloor - 1} \left| \delta(X^r(s)) - \delta^\theta \right| \right) \leq \epsilon_r t + \kappa \left(\int_0^t e^r(s) ds \right). \quad (31)$$

Together with (28), (29) and (31), we deduce that

$$\begin{aligned} e^r(t) &\leq \mathbb{E} \left(\frac{1}{r} \sum_{s=0}^{\lfloor rt \rfloor - 1} \left| \delta(X^r(s)) - \delta^\theta \right| \right) + \mathbb{E} \left(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \right) \\ &\leq \epsilon_r t + \kappa \left(\int_0^t e^r(s) ds \right) + \sqrt{\frac{C't}{r}}. \end{aligned}$$

Hence, for any $\epsilon > 0$ there exists r_0 such that, for a $C'' > 0$, $e^r(t) \leq C'' \left(\epsilon t + \int_0^t e^r(s) ds \right)$, for $r \geq r_0$. Using Gronwall's lemma [27], we obtain that for any $r \geq r_0$,

$$e^r(t) \leq C'' \epsilon t \left(1 + e^t \int_0^t e^{-s} ds \right) \leq C'' \epsilon t (1 + te^t),$$

for all $t < T_1$. Hence, $\lim_{r \rightarrow \infty} e^r(t) = 0$ for $t < T_1$. Since $e^r(t) = \mathbb{E}(\sup_{0 \leq s \leq t} |Y^r(s) - y(s)|) \geq \mathbb{P}(\sup_{0 \leq s \leq t} |Y^r(s) - y(s)| > \epsilon') \cdot \epsilon'$, we obtain $\lim_{r \rightarrow \infty} \mathbb{P}(\sup_{0 \leq s \leq t} |Y^r(s) - y(s)| > \epsilon') = 0$, for $t < T_1$.

In case $T_1 = \infty$, the Theorem 4.3 is now proved $\forall t > 0$. In the remainder of the proof we therefore assume that $T_1 < \infty$.

Second part of the trajectory (one stationary class): We decomposed the trajectory of Y^r into a first part ($t < T_1$) and a second part ($t > T_1$). By the Markov property, we can study the second part of the trajectory supposing that $Y^r(T_1) = (x_1, y_2(T_1)r)$. Since $T_1 < \infty$, the process $X_1^{\{1\}}$ is ergodic. Using the monotonicity of the drift, we have that $X_1^r(t) \leq_{st} X_1^{\{1\}}(t)$ for $t > T_1$ (this can be observed by coupling the two processes and noting that in the case of an infinite number of class-2 users, the service rate for class 1 is the lowest possible since there will always be a class-2 user in its best state). This implies that the family $(X_1^r(t))_r$ is tight and that $\frac{X_1^r(t)}{r} \rightarrow 0$, as $r \rightarrow \infty$, for $t > T_1$. At time T_1 class 2 is saturated, i.e., $\mathcal{U}_1 = \{1\}$. In what follows we assume that $T_1 \leq t < T_2$. Consider the deterministic function $y_2(\cdot)$ as defined in the theorem, $y_2(t) = y_2(T_1) + \tilde{\delta}_2^{\{1\}}(t - T_1)$, $T_1 \leq t < T_2$, and define the error function $E_2^r(t) := \sup_{T_1 \leq s \leq t} |Y_2^r(s) - y_2(s)|$. From (26) we obtain

$$Y_2^r(t) = y_2(T_1) + \frac{1}{r} \sum_{s=\lfloor rT_1 \rfloor}^{\lfloor rt \rfloor - 1} \delta_2(X^r(s)) + \frac{M_2^r(\lfloor rt \rfloor)}{r}. \quad (32)$$

Using the triangular inequality we obtain,

$$\begin{aligned} E_2^r(t) &\leq \frac{1}{r} \sum_{s=\lfloor rT_1 \rfloor}^{\lfloor rt \rfloor - 1} \left| \delta_2(X^r(s)) - \delta_2^{\{1\}}(X_1^r(s)) \right| \\ &\quad + \frac{1}{r} \sum_{s=\lfloor rT_1 \rfloor}^{\lfloor rt \rfloor - 1} \left| \delta_2^{\{1\}}(X_1^r(s)) - \tilde{\delta}_2^{\{1\}} \right| + \sup_{0 \leq s \leq t} \left| \frac{M_2^r(\lfloor rs \rfloor)}{r} \right|. \end{aligned} \quad (33)$$

From Lemma 8.1 we obtain that there exists ϵ_r with $\lim_{r \rightarrow \infty} \epsilon_r = 0$, such that $\forall s = \lfloor rT_1 \rfloor, \dots, \lfloor rt \rfloor$,

$$\left| \delta_2(X^r(s)) - \delta_2^{\{1\}}(X_1^r(s)) \right| \leq \epsilon_r + L \left| \frac{X_2^r(s)}{r} - y_2\left(\frac{s}{r}\right) \right|. \quad (34)$$

We now derive a bound on $|\delta_2^{\{1\}}(X_1^r(s)) - \tilde{\delta}_2^{\{1\}}|$. In order to do that, we define a one-dimensional process \underline{X}_1^r having the same dynamics as our initial process but with the number of class-2 users fixed to $r(y_2(T_1) - \theta)$, $0 < \theta < y_2(T_1)$. Let $\underline{\pi}^{\theta, r}$ be the stationary distribution of \underline{X}_1^r . Adapting Lemma 6 in [12] to our discrete time setting, we obtain that



$$\lim_{r \rightarrow \infty} \underline{\pi}^{\theta, r}(x) = \pi^{\{1\}}(x), \forall x. \quad (35)$$

We define $T^r = \inf\{t : X_2^r(rt) < r(y_2(T_1) - \theta)\}$. Using coupling arguments, we see that $T^r \geq T_1$ a. s., while for $T_1 < t \leq T^r$, $\underline{X}_1^r(t) \leq_{st} X_1^r(t) \leq_{st} X_1^{\{1\}}(t)$. Since $\delta_2^{\{1\}}(\cdot)$ is increasing (follows from monotonicity assumption) we obtain

$$\begin{aligned} \frac{1}{r} \sum_{s=\lfloor rT_1 \rfloor}^{\lfloor rt \rfloor - 1} \delta_2^{\{1\}}(\underline{X}_1^r(s)) &\leq \frac{1}{r} \sum_{s=\lfloor rT_1 \rfloor}^{\lfloor rt \rfloor - 1} \delta_2^{\{1\}}(X_1^r(s)) \\ &\leq \frac{1}{r} \sum_{s=\lfloor rT_1 \rfloor}^{\lfloor rt \rfloor - 1} \delta_2^{\{1\}}(X_1^{\{1\}}(s)) \text{ for } t \leq T^r. \end{aligned} \quad (36)$$

Since $\delta_2^{\{1\}}$ is uniformly upper-bounded, we have that

$$\begin{aligned} \lim_{r \rightarrow \infty} \sum_x \underline{\pi}^{\theta, r}(x) \delta_2^{\{1\}}(x) &= \sum_x \lim_{r \rightarrow \infty} \underline{\pi}^{\theta, r}(x) \delta_2^{\{1\}}(x) \\ &= \sum_x \pi^{\{1\}}(x) \delta_2^{\{1\}}(x) = \tilde{\delta}_2^{\{1\}}. \end{aligned}$$

Fig. 4. Case $T^\infty < \infty$.Fig. 5. Case $T^\infty = \infty$.

Hence, we obtain that for all $\epsilon > 0$ there exists r_0 such that $\sum_{x \geq 0} \pi^{\theta, r_0}(x) \delta_2^{\{1\}}(x) \geq \tilde{\delta}_2^{\{1\}} - \epsilon$, for all $r \geq r_0$, and applying the ergodic theorem for Markov chains we get

$$\lim_{r \rightarrow \infty} \frac{1}{r} \sum_{s=rT_1}^{\lfloor rt \rfloor - 1} \delta_2^{\{1\}}(X_1^{\{1\}}(s)) = (t - T_1) \tilde{\delta}_2^{\{1\}}, \text{ a.s. and}$$

$$\lim_{r \rightarrow \infty} \frac{1}{r} \sum_{s=rT_1}^{\lfloor rt \rfloor - 1} \delta_2^{\{1\}}(\mathbf{X}_1^r(s)) \geq (t - T_1)(\tilde{\delta}_2^{\{1\}} - \epsilon), \text{ a.s..}$$

Together with (36) we obtain that there exists an r_0 such that

$$\left| \frac{1}{r} \sum_{s=\lfloor tT_1 \rfloor}^{\lfloor rt \rfloor - 1} \delta_2^{\{1\}}(X_1^r(s)) - \tilde{\delta}_2^{\{1\}} \right| \leq \epsilon, \text{ a.s., for } t \leq T^r, r \geq r_0. \quad (37)$$

Let us now condition on the set of events $\Omega_1^r(\tilde{\epsilon}) = \{\omega : \sup_{s \leq t} \frac{|M_2^r(\lfloor rs \rfloor)|}{r} \leq \tilde{\epsilon}\}$. From (33), (34) and (37), and using an inequality similar to (30), we obtain that for any sample path in $\Omega_1^r(\tilde{\epsilon})$ (with $\tilde{\epsilon}$ small enough) it holds that

$$E_2^r(t) \leq C''(\epsilon t + \int_0^t E_2^r(s) ds), \text{ for all } t \leq T^r, \quad (38)$$

with r large enough, and $C'' > 0$. Using Gronwall's lemma [27], Equation (38) implies that $E_2^r(t) \leq C'' \epsilon t (1 + te^t)$. Hence, we can conclude that for a given $\epsilon > 0$, there is an $\tilde{\epsilon} > 0$ such that for each sample path in $\Omega_1^r(\tilde{\epsilon})$ it holds that

$$E_2^r(t) \leq \epsilon, \text{ for all } t \leq T^r, \text{ and } r \text{ large enough.} \quad (39)$$

We define $T^\infty := T_1 - \frac{\theta}{\tilde{\delta}_2^{\{1\}}}$ if $\tilde{\delta}_2^{\{1\}} < 0$, and $T^\infty = \infty$ otherwise. Assume $\omega \in \Omega_1^r(\tilde{\epsilon})$. If $T^r < T^\infty < \infty$, then

$$T^\infty - T^r = \frac{E_2^r(T^r)}{|\tilde{\delta}_2^{\{1\}}|} \leq \frac{\epsilon}{|\tilde{\delta}_2^{\{1\}}|}, \text{ so } T^r \geq T^\infty - \frac{\epsilon}{|\tilde{\delta}_2^{\{1\}}|},$$

see also Figure 4. When $T^\infty = \infty$, it is easy to check that $T^r = \infty$ for ϵ small enough: Consider an $\epsilon < \theta$. If $T^r < \infty$, then $|Y_2^r(T^r) - y_2(T^r)| > \theta$. However, since $\omega \in \Omega_1^r(\tilde{\epsilon})$ we also have $|Y_2^r(T^r) - y_2(T^r)| < \epsilon$, which gives a contradiction with the above. See also Figure 5.

Therefore, given t such that $t < T^\infty$, there exists an ϵ small enough such that $t < T^r$ for all r large enough. Hence, from (39) we obtain for $t < T^\infty$,

$$\mathbb{P} \left(\sup_{T_1 \leq s \leq t} |Y_2^r(s) - y_2(s)| \geq \epsilon \right) = \mathbb{P}(E_2^r(t) \geq \epsilon) \leq \mathbb{P}(\Omega/\Omega_1^r(\tilde{\epsilon})). \quad (40)$$

From (29), we obtain $\lim_{r \rightarrow \infty} \mathbb{E} \left(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \right) = 0$. Since $\mathbb{E} \left(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \right) \geq \mathbb{P}(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \geq \tilde{\epsilon}) \cdot \tilde{\epsilon}$, we obtain that $\mathbb{P}(\Omega/\Omega_1^r(\tilde{\epsilon})) = \mathbb{P}(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \geq \tilde{\epsilon}) \leq \frac{1}{\tilde{\epsilon}} \mathbb{E} \left(\sup_{s \leq t} \left| \frac{M^r(\lfloor rs \rfloor)}{r} \right| \right) \rightarrow 0$, as $r \rightarrow \infty$. Together with (40), this implies that for $t < T^\infty$, $\lim_{r \rightarrow \infty} \mathbb{P}(\sup_{T_1 \leq s \leq t} |Y_2^r(s) - y_2(s)| \geq \epsilon) = 0$. Letting $\theta \uparrow y_2(T_1)$, hence $T^\infty \uparrow T_1 + \frac{y_2(T_1)}{\tilde{\delta}_2^{\{1\}}} = T_2$, we obtain that $Y_2^r(\cdot)$ converges in probability to $y_2(\cdot)$ on the set $[T_1, T_2]$.

APPENDIX B: PROOF OF THEOREM 5.2

If condition (16) is not satisfied, the system is not rate stable, see Lemma 8.3 below, which precludes stability. Hence, (16) is a necessary condition for stability. We are now left with proving that BR policies are stable under condition (16). In order to do that, we (i) prove that, under any BR policy, the set of sample paths where service is given to a user that is not in its best state is asymptotically almost surely empty. Then (ii) we prove that the fluid limit will be equal to zero almost surely for some time T large enough, if condition (16) holds. This will allow to conclude the proof.

In the remainder of the proof we focus on BR policies and assume condition (16) to be satisfied. Let us define

$$T := \frac{\sum_{k=1}^K \frac{y_k^f(0)}{\mu_k, N_k}}{\left(1 - \sum_{k=1}^K \frac{\lambda_k}{\mu_k, N_k}\right)} < \infty. \quad (41)$$

We define the random variable $T_\epsilon^r = \inf\{t : \sum_{k=1}^K \frac{Y_k^r(t)}{\mu_k, N_k} \leq \epsilon\}$ for $0 < \epsilon < 1$ and let $\bar{T}_\epsilon^r = \min\{T_\epsilon^r, T\}$. Consider the event A_r which occurs if, during the time interval $[0, \lfloor r\bar{T}_\epsilon^r \rfloor]$, there is at least one user which has been served while not being in its best possible state. Since any BR policy gives absolute priority to users which are in their best possible state, it follows that $\mathbb{P}(A_r) = 1 - \prod_{s=0}^{\lfloor r\bar{T}_\epsilon^r \rfloor} [1 - \prod_{k=1}^K (1 - q_{k, N_k})^r Y_k^r(\frac{s}{r})]$. Since $\bar{T}_\epsilon^r \leq T$ and using basic algebra, it is easy to check that there exists a constant $\xi \in [0, 1)$ such that $\mathbb{P}(A_r) \leq 1 - (1 - \xi^r)^{rT} =: g(r)$. We will now show that

$$\mathbb{P}(\cap_{r=1}^\infty \cup_{\tilde{r}=r}^\infty A_{\tilde{r}}) = 0. \quad (42)$$

Since $\log(1+x) = x + o(x)$ when x is close to 0, we obtain that $(1 - \xi^r)^{rT} = e^{-rT\xi^r + o(r\xi^r)}$ for large values of r . Using Taylor expansions and the fact that $\sum_{r=1}^\infty r\xi^r < \infty$, it follows that, $\sum_{r=1}^\infty g(r) = \sum_{r=1}^\infty (1 - (1 - \xi^r)^{rT}) =$

$\sum_{r=1}^{\infty} 1 - e^{-rT\xi^r + o(r\xi^r)} < \infty$. Hence, we obtain (42) using Borel-Cantelli's lemma [19].

We will now show that at time T the weak fluid limit is almost surely equal to zero. From Equation (42) we obtain that, almost surely, event A_r^c occurs when r is large enough. Hence, only users in their best possible state are served in the interval $[0, r\bar{T}_\epsilon^r]$ when r is large enough, i.e., $T_{k,n}^{BR,r}(t) = 0$, for $n \neq N_k$ and $t \leq r\bar{T}_\epsilon^r$, for r large enough, a.s.. Hence, for almost all ω we have that in the weak fluid limit presentation of Lemma 4.1 it holds that $\sum_{k=1}^K \tau_{k,N_k}^{BR}(t) = t$ and $\tau_{k,n}^{BR}(t) = 0$ for all $n \neq N_k$, $t \leq \liminf_{r \rightarrow \infty} \bar{T}_\epsilon^r$. So we obtain that for $t \leq \liminf_{r \rightarrow \infty} \bar{T}_\epsilon^r$, any weak fluid limit $y^{BR}(\cdot)$ satisfies

$$\begin{aligned} \sum_{k=1}^K \frac{y_k^{BR}(t)}{\mu_{k,N_k}} &= \sum_{k=1}^K \frac{x_k(0)}{\mu_{k,N_k}} + \sum_{k=1}^K \frac{\lambda_k}{\mu_{k,N_k}} t - \sum_{k=1}^K \tau_{k,N_k}^{BR}(t) \\ &= \sum_{k=1}^K \frac{x_k(0)}{\mu_{k,N_k}} - \left(1 - \sum_{k=1}^K \frac{\lambda_k}{\mu_{k,N_k}}\right) t. \end{aligned} \quad (43)$$

Let $T_\epsilon < \infty$ denote the moment that $\sum_{k=1}^K \frac{y_k^{BR}(t)}{\mu_{k,N_k}} = \epsilon$. For a given sample path ω , let r_k be the subsequence corresponding to $\liminf_{r \rightarrow \infty} \bar{T}_\epsilon^r$. By Lemma 4.1 we know that there exists a subsequence r_{k_l} of r_k such that $|\sum_k \frac{y^{BR}(\bar{T}_\epsilon^{r_{k_l}})}{\mu_{k,N_k}} - \sum_k \frac{Y^{BR,r_{k_l}}(\bar{T}_\epsilon^{r_{k_l}})}{\mu_{k,N_k}}| \leq \epsilon'$, for $\epsilon' > 0$ and l large enough. Hence, if $\bar{T}_\epsilon^{r_{k_l}} \leq T_\epsilon$, we have that $T_\epsilon - \bar{T}_\epsilon^{r_{k_l}} \leq \frac{\epsilon'}{1 - \sum_k \lambda_k / \mu_{k,N_k}}$, i.e., $\bar{T}_\epsilon^{r_{k_l}} \geq T_\epsilon - \frac{\epsilon'}{1 - \sum_k \lambda_k / \mu_{k,N_k}}$. Since $\lim_{\epsilon' \downarrow 0} \liminf_{r \rightarrow \infty} \bar{T}_\epsilon^r \geq T_\epsilon$, the description (43) holds for all $t \leq T_\epsilon$. Now letting ϵ go to zero, then $T_\epsilon \rightarrow T$, hence Equation (43) holds for any $t \leq T$, and in particular $|y^f(T)| = 0$.

We conclude that for almost all sample paths, any weak fluid limit converges to 0 at time T , in particular the sequence corresponding to the liminf and the limsup. Hence, one can conclude that $\lim_{r \rightarrow \infty} Y_k^{f,r}(T) = y_k^f(T) = 0$, almost surely. Since $Y_k^{f,r}(T)$ converges almost surely (and therefore, also in probability) to 0, we can use the same argument as in Theorem 5.1, in order to conclude that the system will be stable under policy f , which finishes the proof of Theorem 5.2.

Below we state the rate stability result.

Lemma 8.3: Assume $X(t)$ is rate stable, i.e., $\lim_{t \rightarrow \infty} X_i(t)/t = 0$. Assume $X(t)$ is rate stable, i.e., $\lim_{t \rightarrow \infty} X_i(t)/t = 0$ for all i . Then $\sum_{k=1}^K \frac{\lambda_k}{\mu_k} \leq 1$.

Proof: We first derive rate conservation equations for the (per class) number of users. Using the martingale decomposition of the Markov chain, we obtain that $X_i(t)/t = x_i(0)/t + \sum_{s=1}^t \delta_i(X(s))/t + M_i(t)/t$, where M_i is a square martingale such that $E[M_i(t)^2] \leq Ct$, with $C > 0$ a constant. Hence, $\frac{M_i(t)}{t} \rightarrow 0$, almost surely. So, if $\frac{X_i(t)}{t} \rightarrow 0$ then

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \delta_i(X(s)) = 0. \quad (44)$$

Hence the time average of the Cesaro mean of the service $s_i(\cdot)$ dedicated to class i is equal to λ_i . Given that we consider a system with capacity 1, $\sum_{i=1}^K \frac{s_i(t)}{\mu_{i,N_i}} \leq 1, \forall t$, which, combined with the previous equation, gives that $\sum_{i=1}^K \frac{\lambda_i}{\mu_{i,N_i}} \leq 1$. ■

PLACE
PHOTO
HERE

Urtzi Ayesta is a CNRS researcher working at LAAS, Toulouse, France, and an Ikerbasque research professor at the University of the Basque Country, Spain. He holds a PhD degree from Université de Nice-Sophia Antipolis (France), a MS degree in Electrical Engineering from Columbia University (US) and a BS/MS degree in Telecommunication Engineering from Nafarroako Unibertsitate Publikoa-Universidad Publica de Navarra (Spain).

PLACE
PHOTO
HERE

Martin Erausquin obtained his B.Sc. in Mathematics from the University of the Basque Country (UPV/EHU) in 2007 and he did the Master "Banking and Quantitative Finance" between 2007 and 2010, at the UPV/EHU. As of September 2010, Martin is a Ph.D. student at UPV/EHU under the supervision of Urtzi Ayesta (CNRS) and Eva Ferreira (UPV/EHU).

PLACE
PHOTO
HERE

Matthieu Jonckheere received the MSc degree from ENST Paris and University Paris VI in 2002 and the PhD degree in applied mathematics from the Ecole Polytechnique, France, in 2005, working at France Telecom labs. During 2006-2007, he was a visiting scholar at Centrum voor Wiskunde en Informatica (CWI), Amsterdam and in 2007-2009, he was appointed as an assistant professor in the mathematics and computer science department at Eindhoven University of Technology. He has since 2009 a CONICET research position in the mathematics department of the University of Buenos Aires. His main research interests are in the area of probability theory and performance evaluation of communication networks.

PLACE
PHOTO
HERE

Ina Maria Verloop received the MSc degree in Mathematics from Utrecht University, The Netherlands, in 2005, and the PhD degree from the Mathematics and Computer Science Department of Eindhoven University of Technology, in 2009. Her PhD research was carried at the Center for Mathematics and Computer Science (CWI) in Amsterdam, The Netherlands. From 2009-2011 she was a postdoc at the Basque Center for Applied Mathematics, Derio, Spain. Since October 2011 she is a CNRS researcher at IRIT, Toulouse, France. Her research interests

are in the performance analysis of communication networks, scheduling and queueing theory.