



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

PNA

Probability, Networks and Algorithms



Probability, Networks and Algorithms

Flow-level stability of channel-aware scheduling algorithms

S.C. Borst, M. Jonckheere

REPORT PNA-R0606 MARCH 2006

Centrum voor Wiskunde en Informatica (CWI) is the national research institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organisation for Scientific Research (NWO). CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2006, Stichting Centrum voor Wiskunde en Informatica
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3711

Flow-level stability of channel-aware scheduling algorithms

ABSTRACT

Channel-aware scheduling strategies provide an effective mechanism for improving the throughput performance in wireless data networks by exploiting channel fluctuations. The performance of channel-aware scheduling algorithms has mainly been examined at the packet level for a static user population, often assuming infinite backlogs. Recently, some studies have also explored the flow-level performance in a scenario with user dynamics governed by the arrival and completion of random service demands over time. Although in certain cases the performance may be evaluated by means of a Processor-Sharing model, in general the flow-level behavior has remained largely intractable, even basic stability properties. In the present paper we derive simple necessary stability conditions, and show that these are also sufficient for a wide class of utility-based scheduling policies. This contrasts with the fact that the latter class of strategies generally fail to provide maximum-throughput guarantees at the packet level.

2000 Mathematics Subject Classification: 60K25, 68M20

Keywords and Phrases: Channel-aware scheduling, flow-level performance, Lyapunov function, stability condition, utility-based scheduling, wireless data networks

Flow-Level Stability of Channel-Aware Scheduling Algorithms

Sem Borst^{*,†,*}, Matthieu Jonckheere^{*,1}

^{*}CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

[†]Bell Laboratories, Lucent Technologies
P.O. Box 636, Murray Hill, NJ 07974-0636, USA

^{*}Department of Mathematics & Computer Science
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands

Abstract—Channel-aware scheduling strategies provide an effective mechanism for improving the throughput performance in wireless data networks by exploiting channel fluctuations. The performance of channel-aware scheduling algorithms has mainly been examined at the packet level for a static user population, often assuming infinite backlogs. Recently, some studies have also explored the flow-level performance in a scenario with user dynamics governed by the arrival and completion of random service demands over time. Although in certain cases the performance may be evaluated by means of a Processor-Sharing model, in general the flow-level behavior has remained largely intractable, even basic stability properties. In the present paper we derive simple necessary stability conditions, and show that these are also sufficient for a wide class of utility-based scheduling policies. This contrasts with the fact that the latter class of strategies generally fail to provide maximum-throughput guarantees at the packet level.

I. INTRODUCTION

Channel-aware scheduling strategies provide an effective mechanism for improving the throughput performance in wireless data networks by harnessing channel variations. A prime example is the Proportional Fair algorithm, which has been adopted as the default scheduler for the CDMA 1xEV-DO system [7], [14], [16], and is also considered for implementation in HSDPA. The performance of channel-aware scheduling algorithms has mainly been investigated at the packet level for a static user scenario, sometimes including packet-scale dynamics [4], [28], but often assuming infinite backlogs [1], [13], [23], see also [24], [25], [31], [33] for related results. The assumption of a static user population is a reasonable modeling convention because of the separation of time scales: the scheduling mechanism operates at the packet level on which the user population evolves only relatively slowly. However, when examining the throughput performance, and in particular comparing the throughput allocation among elastic traffic users under various strategies, it is not satisfactory

to assume that the user population is independent of the specific properties of the scheduler. For example, a scheduling strategy that provides high throughput to users with favorable channel conditions, will tend to satisfy the service demands of these users sooner. As a result, the scheduler would be left facing a user population with a higher fraction of users with poor channel conditions. Conversely, a scheduling strategy that grants reasonable throughput to users with poor channel conditions, should to a certain degree benefit from that by seeing fewer of these users.

Recently, some studies have explored the flow-level performance in a scenario as described above with user dynamics governed by the arrival and completion of random service demands over time. In particular, it has been shown that in certain cases the flow-level performance may be described by means of a multi-class Processor-Sharing model where the total service rate varies with the total number of active users [12]. The latter model provides explicit formulas for the distribution of the number of active users of the various classes, the mean response times, the blocking probabilities, and the throughputs. Similar Processor-Sharing type models have been proposed for various types of wireless systems [10], [22], [27], [32].

Despite the above results, the flow-level performance in general has remained largely elusive, even basic stability properties. In the present paper, we derive simple necessary stability conditions, and show that these are in fact also sufficient for a wide class of utility-based scheduling policies. Of course, (in)stability is to a certain extent a theoretical concept that cannot occur in an actual system due to admission and flow control mechanisms and the inherent finiteness of buffers. However, it is plausible that instability effects will be reflected in poor performance in terms of long delays in practical circumstances as well.

It is interesting to observe that the above stability result contrasts with the fact that utility-based scheduling strategies generally fail to provide maximum-throughput guarantees at

¹Work done while second author was with France Telecom R&D, 38-40 rue du Général Leclerc, 92794 Issy-les-Moulineaux, France.

the *packet* level, see for instance [3], [24]. Various simple queue-length-based strategies on the other hand do achieve stability at the packet level whenever possible [4], [17], [30], [31]. In order to reconcile these paradoxical facts, it is worth observing that while utility-based strategies operate agnostically of the queue lengths at the packet level, they *do* respond to congestion that occurs at the flow level. Thus, from a stability perspective, the behavior of a utility-based strategy at the flow level shows resemblance to that of a queue-length-based strategy at the packet level. A somewhat related finding in the context of ‘imperfect’ scheduling in multi-hop wireless networks is reported in [21].

However, a crucial distinction is that at the *packet* level channel fluctuations give rise to *random* time-varying service rates for the various users, which are *independent* of the number of packets stored in the buffer. In contrast, the feasible service rates for the various classes at the *flow* level are *deterministic* as the channel fluctuations ‘average out’, but they *vary* with the number of users because the scheduling gains increase with the degree of multi-user diversity. In that sense, the stability results in the present paper differ from those for queue-length-based strategies at the packet level [4], [30], [31] and extend results for state-independent rate sets as obtained in [5], [6], [8].

The remainder of the paper is organized as follows. In Section II we recapitulate some relevant results for a static user population and state some preliminary facts. We extend the model to accommodate a dynamic user configuration in Section III. In Section IV we establish necessary stability conditions, and prove that these are in fact also sufficient for a wide class of utility-based scheduling strategies. In Section V we discuss some extensions.

II. STATIC USER POPULATION

We first review some relevant results for a system with a static population of M data users. For convenience, we assume that the system operates in a time-slotted fashion. The feasible rates for the various users vary over time according to some stationary and ergodic discrete-time stochastic process $\{R(t)\}$, with $R(t) \subseteq \mathcal{R}$ representing the set of feasible rate vectors to be selected from in time slot t and $\mathcal{R} \subseteq \mathbb{R}_+^M$ denoting the set of all possible rate vectors. To avoid technicalities, we assume that the set \mathcal{R} is finite. With minor modifications, the results extend to scenarios where the set \mathcal{R} is infinite. Let R be a random set with as distribution the stationary distribution of the set of feasible rate vectors, and let $p(S) := \mathbb{P}\{R = S\}$ be the stationary probability that the set of feasible rate vectors is $S \in \mathcal{S}$, with $\mathcal{S} := 2^{\mathcal{R}}$.

Denote by $\mathcal{A} \subseteq \mathbb{R}_+^M$ the set of all achievable (long-term) throughput vectors.

The next proposition provides a characterization of the set \mathcal{A} [4], [13], [15], [26].

Proposition 2.1: The set of achievable throughput vec-

tors \mathcal{A} may be characterized as

$$\mathcal{A} = \{T \in \mathbb{R}_+^M : T \in \sum_{S \in \mathcal{S}} p(S) \text{conv}(S)\},$$

or equivalently,

$$\mathcal{A} = \{T \in \mathbb{R}_+^M : z(T) \geq 1\},$$

where $z(T)$ is the optimal value of the linear program

$$\begin{aligned} \max \quad & z \\ \text{sub} \quad & z \leq z_i = \sum_{S \in \mathcal{S}} \sum_{r \in S} p(S) x(r, S) r_i / T_i \quad i = 1, \dots, M \\ & \sum_{r \in S} x(r, S) \leq 1 \quad S \in \mathcal{S} \\ & x(r, S) \geq 0 \quad r \in S, S \in \mathcal{S}. \end{aligned}$$

The variable $x(r, S)$ in the above linear program may be interpreted as the fraction of time that the rate vector $r \in S$ is selected given that the set of feasible rate vectors is $S \in \mathcal{S}$. Thus, $\sum_{r \in S} x(r, S) r \in \text{conv}(S)$ is the average rate vector selected given that the set of feasible rate vectors is $S \in \mathcal{S}$, and $\sum_{S \in \mathcal{S}} \sum_{r \in S} p(S) x(r, S) r$ is the throughput vector. In particular, the term $\sum_{S \in \mathcal{S}} \sum_{r \in S} p(S) x(r, S) r_i$ represents the throughput received by user i , and the variable z_i measures the throughput as a fraction of the target throughput T_i .

The next proposition provides a characterization of the optimal solution to the above linear program based on the complementary slackness conditions [4], [13].

Proposition 2.2: There exists a vector $w^* \in \mathbb{R}_+^M$ such that any optimal solution $x^*(r, S)$ to the above linear program satisfies

$$x^*(r, S) \left[\sum_{i=1}^M w_i^* r_i - \max_{s \in S} \sum_{i=1}^M w_i^* s_i \right] = 0,$$

for all $r \in S, S \in \mathcal{S}$.

The above proposition shows that any feasible (non-dominated) throughput vector can be achieved by some weight-based strategy which selects in time slot t the rate vector $r \in R(t)$ identified by

$$\sum_{i=1}^M w_i^* r_i = \max_{s \in R(t)} \sum_{i=1}^M w_i^* s_i,$$

augmented with a suitable tie-breaking rule. In particular, any component-wise increasing function of the throughput vector is maximized by some weight-based strategy. (This may also be deduced from the observation that the achievable throughput region is convex by a simple time-sharing argument.) However, determining the weights that maximize a given objective function directly is a formidable task, requiring detailed information on the channel statistics. Instead, one may maximize a given objective function by using adaptive weights, without the need for explicit knowledge of the rate characteristics. Specifically, let $U_i(\cdot)$ be the strictly concave utility function of user i . A gradient-based strategy selects in

slot t a rate vector $r \in R(t)$ that maximizes $\sum_{i=1}^M r_i U'_i(W_i(t))$, i.e., it selects in time slot t a rate vector $r \in \hat{R}(t)$ identified by

$$\sum_{i=1}^M r_i U'_i(W_i(t)) = \max_{r \in R(t)} \sum_{i=1}^M s_i U'_i(W_i(t)), \quad (1)$$

where $W_i(t)$ denotes the exponentially smoothed throughput of user i at time t . The exponentially smoothed throughput is updated in each time slot as $W_i(t+1) = (1-\delta)W_i(t) + \delta X(r,t)r_i$, with δ a smoothing coefficient and $X(r,t)$ a 0–1 variable indicating whether rate vector r is selected in time slot t or not. A prime example of a gradient-based strategy is the Proportional Fair algorithm for the CDMA 1xEV-DO system [7], [14], [16] which involves a utility function $U_i(\cdot) \equiv \log(\cdot)$ as the name suggests.

Note that the above selection rule may be interpreted as a stochastic gradient ascent scheme. It has been shown in [2], [20], [29], [34] that under mild assumptions the gradient-based strategy indeed maximizes the aggregate utility as the time constant $T = 1/\delta$ in the exponential smoothing grows large. Specifically, as $\delta \downarrow 0$, the exponentially smoothed throughputs approach a limit point $T^* = (T_1^*, \dots, T_M^*)$ which maximizes the aggregate utility over the achievable throughput region \mathcal{A} , i.e., $\sum_{i=1}^M U_i(T_i^*) \geq \sum_{i=1}^M U_i(T_i)$ for all $T \in \mathcal{A}$. In addition, the limiting vector T^* satisfies the set of fixed-point equations

$$T_i = \sum_{S \in \mathcal{S}} \sum_{r \in S} r_i p(S) \mathbb{I}_{\{\sum_{i=1}^M U'_i(T_i) r_i = \max_{s \in \mathcal{S}} \sum_{i=1}^M U'_i(T_i) s_i\}}, \quad (2)$$

which follow directly from the above selection rule.

We conclude the section with a brief discussion of the special case where only a single user is served in each time slot. In that case, each of the sets $S \in \mathcal{S}$ consists of M vectors $r^{1,S}, \dots, r^{M,S} \in \mathbb{R}_+^M$ where $r_i^{m,S} = 0$ for all $i \neq m$. It will be convenient to denote by $R_i(t) := r_i^{i,R(t)}$ the feasible rate of user i in time slot t . The selection criterion (1) then implies that time slot t is allocated to user i^* identified by

$$R_i(t) U'_i(W_{i^*}(t)) = \max_{j=1, \dots, M} R_j(t) U'_j(W_j(t)),$$

and the set of fixed-point equations (2) take the form

$$T_i = \sum_{S \in \mathcal{S}} \sum_{r \in S} r_i p(S) \mathbb{I}_{\{U'_i(T_i) r_i = \max_{j=1, \dots, M} U'_j(T_j) r_j\}}.$$

III. DYNAMIC USER CONFIGURATION

We now extend the model to accommodate a dynamic configuration of users. The user dynamics result from finite-size service demands that arrive randomly over time. We assume that the duration of the time slots is short relative to the size and arrival frequency of the service demands. Thus, the scheduling mechanism operates on a relatively fast time scale compared to the user dynamics, making it natural to analyze the user-level performance in continuous rather than discrete time, and assume that the users are served simultaneously rather than in a time-slotted fashion. The continuous-time model naturally inherits its service characteristics from the

discrete-time model. Specifically, we assume that the set of feasible instantaneous service rate vectors in the continuous-time context for a given user population coincides with the set of achievable long-term throughput vectors for that user population in a discrete-time setting. The above assumptions ignore the discrete nature of the time slots and neglect the transient fluctuations in the throughput. However, the law of large numbers suggests that these effects should be negligible in some suitable asymptotic sense in a limiting regime where the duration of the time slots shrinks relative to the time scale of the user dynamics.

We consider a scenario with K traffic classes. A traffic class represents a category of statistically identical users in terms of service requests and rate characteristics. Class- k users arrive as a Poisson process of rate λ_k , and have exponentially distributed service requirements with mean $1/\mu_k$. Denote by $\rho_k := \lambda_k/\mu_k$ the traffic intensity associated with class k , and define $\rho := (\rho_1, \dots, \rho_K)$.

For a given user population (n_1, \dots, n_K) , let $A(n_1, \dots, n_K) \subseteq \mathbb{R}_+^K$ be the set of feasible service rate vectors for the various classes. Formally, a service rate vector $a \in \mathbb{R}_+^K$ belongs to the set $A(n_1, \dots, n_K)$ if there exists a vector $T = (T_{1,1}, \dots, T_{1,n_1}, \dots, T_{K,1}, \dots, T_{K,n_K})$ that belongs to the achievable throughput region $\mathcal{A} \subseteq \mathbb{R}_+^M$ with $M = \sum_{k=1}^K n_k$ as defined in the previous section for a static scenario with n_k class- k users, $k = 1, \dots, K$, such that $a_k = \sum_{i=1}^{n_k} T_{k,i}$.

The users are served according to a utility-based scheduling strategy. Thus, for a given user population $n = (n_1, \dots, n_K)$, the throughputs for the various users are selected so as to maximize the aggregate weighted utility $\sum_{k=1}^K \sum_{i=1}^{n_k} \theta_k U(T_{i,k})$, with $U(\cdot)$ some concave utility function and $\theta \in \mathbb{R}_+^K$. Because of the concavity, equal weights, and symmetry among users within classes, optimality requires that $T_{k,i} = T_k$ for all $i = 1, \dots, n_k$. Thus, the service rates $s(n) := (s_1(n), \dots, s_K(n))$ for the various classes are selected such that $\sum_{k=1}^K \theta_k n_k U(s_k(n)/n_k) \geq \sum_{k=1}^K \theta_k n_k U(s_k/n_k)$ for all $s := (s_1, \dots, s_K) \in A(n_1, \dots, n_K)$.

The sets $A(n_1, \dots, n_K)$ have two crucial properties. First of all, a simple time-sharing argument implies that each set $A(n_1, \dots, n_K)$ is convex as observed before. Second, the sets $A(n_1, \dots, n_K)$ are monotone increasing in the user population, i.e., if $(m_1, \dots, m_K) \leq (n_1, \dots, n_K)$, then $A(m_1, \dots, m_K) \subseteq A(n_1, \dots, n_K)$. The latter property follows from the fact that additional users may simply be excluded from service without affecting the feasible service rates of the remaining users. Since the sets $A(n_1, \dots, n_K)$ are monotone increasing, we may define $A^* := \lim_{n_1, \dots, n_K \rightarrow \infty} A(n_1, \dots, n_K) = \sup_{(n_1, \dots, n_K) \in \mathbb{N}^K} A(n_1, \dots, n_K)$. It is easily verified that the set A^* is convex.

While the set A^* may have a complicated structure in general, it has a rather simple form in the special case where only a single user is served in each time slot. Define $R_k^* := \inf\{r : \mathbb{P}\{R_k > r\} = 0\}$ as the maximum possible value of the rate of class- k users. Then, assuming independence among users, for any $r < R_k^*$, it is easily verified that $re_k \in A(n_1, \dots, n_K)$ for n_k large enough, with e_k denoting the k -th unit vector, which implies $re_k \in A^*$. By convexity, it follows that $A^* = \text{conv}(\{R_1^*e_1, \dots, R_K^*e_K\}) = \{x \in \mathbb{R}_+^K : \sum_{k=1}^K \frac{x_k}{R_k^*} \leq 1\}$.

If in addition the relative variations around the time-average rates are identically distributed for all classes, i.e., $R_k \stackrel{d}{=} C_k Y$ for all $k = 1, \dots, K$, with C_k representing the time-average rate of class- k users, then $R_k^* = C_k G^*$ with $G^* := \inf\{y : \mathbb{P}\{Y > y\} = 0\}$. In that case, $A^* = \{x \in \mathbb{R}_+^K : \sum_{k=1}^K \frac{x_k}{C_k} \leq G^*\}$.

IV. STABILITY CONDITION

The next proposition provides a simple necessary stability condition.

Proposition 4.1: No scheduling strategy achieves stability for $\rho \notin A^*$.

Proof

If r is the long-term mean service rate vector, then it must be the case that $r \in \text{conv}(A^*) = A^*$. In order for the system to be stable, we must have $r = \rho$. Hence, $\rho \notin A^*$ precludes stability of the system. \square

For any $n = (n_1, \dots, n_K)$, define $G_n(s_1, \dots, s_K) := \sum_{k=1}^K \theta_k n_k^\eta s_k^\gamma$, with $\theta, \eta, \gamma > 0$. A scheduling strategy that selects service rate vector $s(n) = (s_1(n), \dots, s_K(n))$ as function of the user population $n = (n_1, \dots, n_K)$ is said to be B -strong if for all $\delta > 0$, there exists a finite set $V_\delta \subseteq \mathbb{N}^K$ such that $G_n(s(n)) \geq (1 - \delta) \sup_{s \in B} G_n(s)$ for all $n \notin V_\delta$.

Note that a scheduling strategy that selects the service rate vector $s(n)$ so as to maximize $G_n(s)$ over $A(n_1, \dots, n_K)$ is A^* -strong, which may be checked as follows. Since $A^* = \lim_{n_1, \dots, n_K \rightarrow \infty} A(n_1, \dots, n_K)$, there exists for all $\delta > 0$ a finite set V_δ such that $(1 - \delta)^{1/\gamma} s \in A(n_1, \dots, n_K)$ for all $s \in A^*$, $n \notin V_\delta$. Hence, $G_n(s(n)) = \sup_{s \in A(n_1, \dots, n_K)} G_n(s) \geq \sup_{s \in A^*} G_n((1 - \delta)^{1/\gamma} s) = (1 - \delta) \sup_{s \in A^*} G_n(s)$ for all $n \notin V_\delta$.

In particular, a utility-based scheduling strategy with $U(x) = x^\gamma$ is A^* -strong, $\eta = 1 - \gamma$. A queue-based scheduling strategy is A^* -strong as well, with $\eta = 1$, $\gamma = 1$.

The next proposition (taking $B = A^*$) shows that the condition identified in Proposition 4.1 is also (nearly) sufficient for stability.

Proposition 4.2: A B -strong scheduling strategy achieves stability for $\rho \in \text{int}(B)$ and $\theta, \eta > 0$, $0 < \gamma \leq 1$.

Proof

The proposition may be established in a similar fashion as the stability of α -fair strategies in wireline bandwidth-sharing networks [8]. The main difference is the fact that the rate region varies with the user population and only approaches the stated capacity region as the user population grows suitably large.

Specifically, let $n(t) := (n_1(t), \dots, n_K(t))$, with $n_k(t)$ denoting the number of class- k users at time t , and consider the set of fluid limits defined by

$$N(t) = \lim_{\omega \rightarrow \infty} \frac{n(\omega t)}{\omega},$$

with $\sum_{k=1}^K n_k(0) = \omega$.

Given an initial state $N(0)$, it follows from the strong law of large numbers that the evolution of the fluid limits is uniquely described by the set of differential equations:

$$\frac{d}{dt} N_k(t) = \lambda_k - \mu_k S_k(t) \quad (3)$$

for all k, t such that $N_k(t) > 0$, where $S(t) := \lim_{\omega \rightarrow \infty} S(\omega t)$, assuming the limit exists, and $S(\omega t) := s(n(\omega t))$.

We now prove that $G_{N(t)}(S(t)) \geq \sup_{s \in B} G_{N(t)}(s)$.

It may be assumed that $N(t) \neq 0$ because otherwise the assertion is trivially true.

Then there exists some ω_δ such that $n(\omega t) \notin V_\delta$ for all $\omega > \omega_\delta$.

Hence, $G_{n(\omega t)}(s(n(\omega t))) \geq (1 - \delta) \sup_{s \in B} G_{n(\omega t)}(s)$ for all $\omega > \omega_\delta$, which implies $G_{N(t)}(S(\omega t)) \geq (1 - \delta) \sup_{s \in B} G_{N(t)}(s)$ for all $\omega > \omega_\delta$.

Since $\delta > 0$ is arbitrary, it readily follows that $G_{N(t)}(S(t)) \geq \sup_{s \in B} G_{N(t)}(s)$ as asserted.

The concavity of the function $G_{N(t)}(\cdot)$ then implies that $G'_{N(t)}(s)(s - S(t)) \leq 0$ for any $s \in B$, which gives

$$\sum_{k=1}^K \theta_k s_k^{\gamma-1} N_k(t)^\eta (s_k - S_k(t)) \leq 0.$$

Since $\rho \in \text{int}(B)$, it follows that $\rho(1 + \epsilon) \in \text{int}(B)$ for some $\epsilon > 0$.

Hence, in particular we have

$$\sum_{k=1}^K \theta_k \rho_k^{\gamma-1} N_k(t)^\eta (\rho_k(1 + \epsilon) - S_k(t)) \leq 0. \quad (4)$$

Consider the Lyapunov function

$$F(x_1, \dots, x_K) = \sum_{k=1}^K \frac{\theta_k}{\mu_k} \rho_k^{\gamma-1} \frac{x_k^{1+\eta}}{1+\eta}.$$

From (3), we obtain

$$\frac{d}{dt} F(N(t)) = \sum_{k=1}^K \theta_k \rho_k^{\gamma-1} N_k(t)^\eta (\rho_k - S_k(t)). \quad (5)$$

Substituting (4) into (5),

$$\frac{d}{dt}F(N(t)) < -\epsilon \sum_{k=1}^K \theta_k \rho_k^\gamma N_k(t)^\eta.$$

As in [8], it may then be deduced that $F(N(t)) = 0$, and hence $N(t) = 0$, for all $t \geq \tau$ for some $\tau < \infty$, which implies that the system is stable. \square

In the special case where only a single user is served in each time slot so that $A^* = \{x \in \mathbb{R}_+^K : \sum_{k=1}^K \frac{x_k}{R_k^*} \leq 1\}$, the stability condition of Proposition 4.2 reduces to $\sum_{k=1}^K \rho_k / R_k^* < 1$. Thus, an A^* -strong strategy achieves stability as long as the system is stable if every class were served at the maximum possible rate all the time, which is also a necessary condition, as shown by Proposition 4.1. This may be explained by the observation that under an A^* -strong strategy every class will either be served at the maximum possible rate or not at all whenever any of the classes drifts into instability.

If in addition the relative variations around the time-average rates are identically distributed for all classes, so that $A^* = \{x \in \mathbb{R}_+^K : \sum_{k=1}^K \frac{x_k}{C_k} \leq G^*\}$, then the stability condition may be written as $\tilde{\rho} < G^*$ as in [12], with $\tilde{\rho} := \sum_{k=1}^K \tilde{\rho}_k$ and $\tilde{\rho}_k := \rho_k / C_k$.

V. EXTENSIONS

We now discuss some extensions of the stability results. First we consider a scenario where an admission control procedure determines which users are admitted into the system. Users that cannot instantaneously be accommodated are denied access and lost. We make the natural assumption that the admissible region $P \subseteq \mathbb{N}^K$ is monotone, i.e., if $(m_1, \dots, m_K) \leq (n_1, \dots, n_K)$, and $(n_1, \dots, n_K) \in P$, then $(m_1, \dots, m_K) \in P$. Observe that otherwise a service completion might potentially cause the user population to become inadmissible. Denote $N_k^* := \max_{(n_1, \dots, n_K) \in P} n_k$, and define $\mathcal{K}^* := \{1 \leq k \leq K : N_k^* = \infty\}$ as the set of classes which are not subject to admission control. Without loss of generality we assume that the classes are indexed such that $\mathcal{K}^* := \{1, 2, \dots, K^*\}$. Denote $\bar{A}(n_1, \dots, n_{K^*}) := \{x \in \mathbb{R}_+^{K^*} : (x_1, \dots, x_{K^*}, 0, \dots, 0) \in A(n_1, \dots, n_{K^*}, 0, \dots, 0)\}$. Define $\bar{A}^* := \lim_{n_1, \dots, n_K \rightarrow \infty} \bar{A}(n_1, \dots, n_K) = \sup_{(n_1, \dots, n_K) \in \mathbb{N}^K} \bar{A}(n_1, \dots, n_K)$. It is easily verified that the set \bar{A}^* is convex.

Proposition 5.1: An A^* -strong scheduling strategy achieves stability for $(\rho_1, \dots, \rho_{K^*}) \in \text{int}(\bar{A}^*)$ and $\eta > 0$, $0 < \gamma \leq 1$.

Observe that the classes $k \notin \mathcal{K}^*$ which are subject to admission control do not play any role in the above stability condition. The explanation is that the fraction of resources consumed by these classes becomes negligible as the system

saturates with traffic from non-controlled classes.

The stability results also extend to an integrated system with both elastic flows and streaming users as considered in [9], [18], [19]. The explanation is that the holding times of the streaming users are random but independent of the amount of service received. Hence, the population of streaming users is simply that in an infinite-server queue, and not affected by the degree of congestion. The fraction of resources consumed by the streaming users therefore vanishes as the load imposed by the elastic flows approaches the capacity limit.

APPENDIX

We now present an alternative proof of Proposition 4.2 which uses Foster's criterion with a Lyapunov function in terms of the original process rather than the fluid limits.

Proof

Consider the Lyapunov function

$$F(n_1, \dots, n_K) = \sum_{k=1}^K \frac{\theta_k}{\mu_k} \rho_k^{\gamma-1} \frac{n_k^{1+\eta}}{1+\eta}.$$

When the user population is $n = (n_1, \dots, n_K)$, transitions may occur to state $n + e_k$ at rate λ_k or to state $n - e_k$ at rate $\mu_k s_k(n)$.

Hence, given that the current user population is $n = (n_1, \dots, n_K)$, the drift in the Lyapunov function is

$$\Delta(n) = \sum_{k=1}^K \theta_k \rho_k^{\gamma-1} n_k^\eta (\rho_k - s_k(n)) + o\left(\sum_{k=1}^K n_k^\eta\right). \quad (6)$$

It thus suffices to show that $\Delta(n) < 0$ whenever $n \notin W$ for some finite set $W \subseteq \mathbb{N}^K$.

For all $\delta > 0$, there exists a finite set $V_\delta \subseteq \mathbb{N}^K$ such that $G_n(s(n)) \geq (1 - \delta) \sup_{s \in B} G_n(s)$ for all $n \notin V_\delta$.

The concavity of the function $G_n(\cdot)$ then implies that $G'_n(s)(s - S(n)) \leq F(\delta)$ for any $s \in B$, which gives

$$\sum_{k=1}^K \theta_k s_k^{\gamma-1} n_k^\eta (s_k - s_k(n)) \leq F(\delta)$$

for all $n \notin V_\delta$.

Since $\rho \in \text{int}(B)$, it follows that $\rho(1 + \epsilon) \in \text{int}(B)$ for some $\epsilon > 0$.

Hence, in particular we have

$$\sum_{k=1}^K \theta_k \rho_k^{\gamma-1} n_k^\eta (\rho_k(1 + \epsilon) - s_k(n)) \leq F(\delta) \quad (7)$$

for all $n \notin V_\delta$.

Substituting (7) into (6),

$$\Delta(n) < -\epsilon \sum_{k=1}^K \theta_k \rho_k^{\gamma-1} n_k^\eta + F(\delta) + o\left(\sum_{k=1}^K n_k^\eta\right),$$

for all $n \notin V_\delta$, and thus $\Delta(n) < 0$ whenever $n \notin W$ for some finite set W , which completes the proof. \square

REFERENCES

- [1] Agrawal, R., Bedekar, A., La, R.J., Subramanian, V. (2001). Class and channel condition based weighted proportional fair scheduler. In: *Teletraffic Engineering in the Internet Era, Proc. ITC-17*, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 553–565.
- [2] Agrawal, R., Subramanian, V. (2002). Optimality of certain channel-aware scheduling policies. In: *Proc. 40th Annual Allerton Conf. Commun. Control., Comp.*, 1532–1541.
- [3] Andrews, D.M. (2004). Instability of the Proportional Fair scheduling algorithm for HDR. *IEEE Trans. Wireless Commun.* **3**, 1422–1426.
- [4] Andrews, D.M., Kumaran, K., Ramanan, K., Stolyar, A.L., Vijayakumar, R., Whiting, P.A. (2004). Scheduling in a queueing system with asynchronously varying service rates. *Prob. Eng. Inf. Sc.* **18**, 191–217.
- [5] M. Armony, N. Bambos (1999). Queueing networks with interacting service resources. In: *Proc. 37th Annual Allerton Conf. Commun., Control, Comp.*, 42–51.
- [6] N. Bambos, G. Michailidis (2004). Queueing and scheduling in random environments. *Adv. Appl. Prob.* **36**, 293–317.
- [7] Bender, P., Black, P., Grob, M., Padovani, R., Sindhushayana, N., Viterbi, A. (2000). CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users. *IEEE Commun. Mag.* **38** (7), 70–77.
- [8] Bonald, T., Massoulié, L. (2001). Impact of fairness on Internet performance. In: *Proc. ACM Sigmetrics / Performance 2001 Conf.*, 82–91.
- [9] Bonald, T., Proutière, A. (2004). On performance bounds for the integration of elastic and adaptive streaming flows. In: *Proc. ACM Sigmetrics / Performance 2004 Conf.*, 235–245.
- [10] Bonald, T., Proutière, A. (2003). Wireless downlink data channels: user performance and cell dimensioning. In: *Proc. ACM Mobicom 2003*, 339–352.
- [11] Bonald, T., Proutière, A., Régnié, G., Roberts, J.W. (2001). Insensitivity results in statistical bandwidth sharing. In: *Teletraffic Engineering in the Internet Era, Proc. ITC-17*, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 125–136.
- [12] Borst, S.C. (2003). User-level performance of channel-aware scheduling algorithms in wireless data networks. In: *Proc. Infocom 2003*.
- [13] Borst, S.C., Whiting, P.A. (2001). Dynamic rate control algorithms for HDR throughput optimization. In: *Proc. Infocom 2001*, 976–985.
- [14] Chaponniere, E.F., Black, P.J., Holtzman, J.M., Tse, D.N.C. (2002). Transmitter directed code division multiple access system using path diversity to equitably maximize throughput. US Patent 6,449,490.
- [15] Cruz, R.L., Santhanam, A.V. (2003). Optimal routing, link scheduling and power control in multi-hop wireless networks. In: *Proc. Infocom 2003*.
- [16] Jalali, A., Padovani, R., Pankaj, R. (2000). Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system. In: *Proc. IEEE VTC 2000 Spring Conf.*, 1854–1858.
- [17] Kahale, N., Wright, P.E. (1997). Dynamic global packet routing in wireless networks. In: *Proc. IEEE Infocom '97*, 1416–1423.
- [18] Key, P.B., Massoulié, L., Bain, A., Kelly, F.P. (2003). A network flow model for mixtures of file transfers and streaming traffic. In: *Providing QoS in Heterogeneous Environments, Proc. ITC-18*, Berlin, eds. J. Charzinski, R. Lehnert, P. Tran Gia (North-Holland, Amsterdam), 1021–1030.
- [19] Key, P.B., Massoulié, L., Bain, A., Kelly, F.P. (2004). Fair Internet traffic integration: network flow models and analysis. *Annales des Telecommunications* **59**, 1338–1352.
- [20] Kushner, H.J., Whiting, P.A. (2004). Convergence of Proportional-Fair sharing algorithms under general conditions. *IEEE Trans. Wireless Commun.* **3**, 1250–1259.
- [21] Lin, X., Shroff, N.B. (2005). The impact of imperfect scheduling on cross-layer rate control in wireless networks. In: *Proc. Infocom 2005*.
- [22] Litjens, R., Roijers, F., Van den Berg, J.L., Boucherie, R.J., Fleuren, M.J. (2003). Performance analysis of wireless LAN's: an integrated packet/flow level approach. In: *Providing QoS in Heterogeneous Environments, Proc. ITC-18*, Berlin, eds. J. Charzinski, R. Lehnert, P. Tran Gia (North-Holland, Amsterdam), 931–940.
- [23] Liu, X., Chong, E.K.P., Shroff, N.B. (2003). A framework for opportunistic scheduling in wireless networks. *Comp. Netw.* **41**, 451–474.
- [24] Neely, M.J., Modiano, E., Li, C.-P. (2005). Fairness and optimal stochastic control for heterogeneous networks. In: *Proc. Infocom 2005*.
- [25] Neely, M.J., Modiano, E., Rohrs, C.E. (2002). Power and server allocation in a multi-beam satellite with time-varying channels. In: *Proc. Infocom 2002*, 1451–1460.
- [26] Neely, M.J., Modiano, E., Rohrs, C.E. (2003). Dynamic power allocation and routing for time-varying wireless networks. In: *Proc. Infocom 2003*.
- [27] Prakash, R., Veeravalli, V.V. (2002). A time-scale separation technique for the analysis of random access systems with incremental redundancy. In: *Proc. ISIT 2002*.
- [28] Shakkottai, S., Stolyar, A.L. (2001). Scheduling algorithms for a mixture of real-time and non-real time data in HDR. In: *Teletraffic Engineering in the Internet Era, Proc. ITC-17*, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 793–804.
- [29] Stolyar, A.L. (2005). On the asymptotic optimality of the gradient scheduling algorithm for multi-user throughput allocation. *Oper. Res.* **53**, 12–25.
- [30] Tassiulas, L., Ephremides, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Aut. Contr.* **37**, 1936–1948.
- [31] Tassiulas, L., Ephremides, A. (1993). Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. Inf. Theory* **30**, 466–478.
- [32] Telatar, I.E., Gallager, R.G. (1995). Combining queueing theory with information theory for multi-access. *IEEE J. Sel. Areas Commun.* **13**, 963–969.
- [33] Tsibonis, V., Georgiadis, L., Tassiulas, L. (2003). Exploiting wireless channel state information for throughput maximization. In: *Proc. Infocom 2003*.
- [34] Viswanath, P., Tse, D.N.C., Laroia, R. (2002). Opportunistic beamforming using dumb antennas. *IEEE Trans. Inf. Theory* **48**, 1277–1294.