OPERATIONS RESEARCH
Vol. 53, No. 1, January–February 2005, pp. 12–25
ISSN 0030-364X | EISSN 1526-5463 | 05 | 5301 | 0012

DOI 10.1287/opre.1040.0156

On the Asymptotic Optimality of the Gradient Scheduling Algorithm for Multiuser Throughput Allocation

Alexander L. Stolyar
Bell Laboratories, Lucent Technologies, 600 Mountain Ave., 2C-322, Murray Hill, New Jersey 07974, stolyar@research.bell-labs.com

We consider the model where N queues (users) are served in discrete time by a generalized switch. The switch state is random, and it determines the set of possible service rate choices (scheduling decisions) in each time slot. This model is primarily motivated by the problem of scheduling transmissions of N data users in a shared time-varying wireless environment, but also includes other applications such as input-queued cross-bar switches and parallel flexible server systems.

The objective is to find a scheduling strategy maximizing a concave utility function $H(u_1, \dots, u_N)$, where u_ns are long-term average service rates (data throughputs) of the users, assuming users always have data to be served.

We prove asymptotic optimality of the gradient scheduling algorithm (which generalizes the well-known proportional fair algorithm) for our model, which, in particular, allows for simultaneous service of multiple users and for discrete sets of scheduling decisions. Analysis of the transient dynamics of user throughputs is the key part of this work.

Subject classifications: communications; queues: algorithms, networks, optimization; networks: stochastic.
Area of review: Telecommunications and Networking.
History: Received January 2003; accepted June 2003.

1. Introduction

We consider the model of Stolyar (2004), where N queues (users) are served in discrete time by a generalized switch. The switch state is a random ergodic process, and it determines the set of possible service rate choices (scheduling decisions) in each time slot. This model is motivated primarily by the problem of scheduling transmissions of N data users over a shared time-varying wireless environment, which naturally arises, for example, in modern wireless data technologies such as HDR (CDMA2000 1xEV-DO) (Bender et al. 2000, Jalali et al. 2000). It also includes as special cases other models and applications such as *input-queued cross-bar switches* (McKeown et al. 1996, Mekkittikul and McKeown 1996), a discrete time version of the *parallel server system* (Harrison 1998, Williams 2000), and a *parallel-processing system* (Gans and van Ryzin 1998). (Models of this type have long history, going back to the work of Tassiulas and Ephremides 1992, 1993. See Stolyar 2004 for a further review of the model history and applications.)

In this paper, we consider the "saturated system" in which each user has infinite amount of data to be served (transmitted), and we are concerned with optimizing the vector of average service rates (throughputs) $u = (u_1, \dots, u_N)$ to maximize some concave utility function $H(u)$. In the wireless data scheduling context, this problem was first considered by Tse (2004) (see also Viswanath et al. 2003), in which the utility function $\sum \log(u_n)$, and the corresponding *proportional fair* (PF) scheduling algorithm (defined below) were introduced for the model in which (as in HDR) one user can be served (transmit data) at a time.

The gradient algorithm (defined below) is a natural generalization of the PF algorithm in that it applies to any concave utility function and to the systems where multiple users can be served at a time. The goal of this paper is to formally prove (asymptotic) optimality of the gradient algorithm for our rather general model, which in particular allows for (a) simultaneous service of multiple users and (b) the set of scheduling decisions to be discrete (as is the case in HDR and other wireless data technologies).

To describe our main results and connections to previous work, let us first describe the model in more detail. (The formal description will be given in §3.) Queues (users) $n = 1, \dots, N$, are served by a switch in discrete time $t = 0, 1, 2, \dots$. Switch state $m = (m(t),\ t = 0, 1, 2, \dots)$ is a random ergodic process. In each state m, the switch can choose a scheduling decision k from a set $K(m)$. Each decision k has the associated service rate vector $\mu^m(k) = (\mu_1^m(k), \dots, \mu_N^m(k))$. This vector represents the amount of data of each user which will be served (transmitted) in one "time slot" if decision k is chosen. (We emphasize the fact that our model allows for features (a) and (b), specified earlier.)

As mentioned earlier, we consider the situation when each user always has data to be served, and we are interested in optimizing the average service rates $u =$

$(u_1, \ldots, u_N)$. Namely, we seek to find a scheduling algorithm that produces u solving the problem

$$\max \quad H(u), \tag{1}$$

$$\text{subject to} \quad u \in V, \tag{2}$$

where $H(u)$ is a strictly concave smooth utility function, and set V above is the system rate region, i.e., the set of all feasible long-term service rate vectors. It is easily observed that (under some natural nonrestrictive assumptions) V is a convex closed bounded set, and consequently, the solution u^* to the problem (1)–(2) exists and is unique. (As we comment later in the paper, the assumption of strict concavity of utility function H is not essential. See §12.)

The gradient algorithm is defined as follows. If at time t the switch is in state m, the algorithm chooses a (possibly nonunique) decision

$$k(t) \in \arg\max_k \nabla H(X(t)) \cdot \mu^m(k),$$

maximizing the scalar product with the gradient of $H(X(t))$, where the vector $X(t)$ is updated as follows:

$$X(t+1) = (1-\beta)X(t) + \beta\mu^m(k),$$

with arbitrary initial value $X(0)$ and with $\beta > 0$ being a fixed (small) parameter. (We see from this definition that, roughly, vector $X(t)$ represents exponentially smoothed average service rates.)

The PF algorithm is the gradient algorithm with utility function $H(u) = \sum \log(u_n)$. Also, very often the PF algorithm is considered for the "one-at-a-time" case, when service rate vectors $\mu^m(k)$ can have no more than one nonzero component. In other words, in each time slot, only one user can be picked for service; and if user n is picked, it is served at the rate μ_n^m, depending on the switch state m. In this case the PF algorithm is particularly simple (Tse 2004, Jalali et al. 2000)—in each time slot t it serves the user n for which $\mu_n^m / X_n(t)$ is maximal.

In this paper, we prove the following property for our model.

Property 1. (Asymptotic Optimality of the Gradient Algorithm). Let u^β denote the vector of expected average service rates in a system with fixed parameter $\beta > 0$. Then, as $\beta \to 0$, $u^\beta \to u^*$.

The rates u^β are defined in terms of averages over finite time intervals. Formal definitions and corresponding results, Theorem 1 and Corollary 1, are presented in §7. The results show that the convergence in Property 1 holds in a strong sense.

Establishing Property 1 naturally involves the analysis of the (limiting, continuous time) trajectories of the process $X(t/\beta)$, with small fixed $\beta > 0$. Such trajectories $x = (x(t),\ t \geqslant 0)$, which we call *fluid sample paths* (FSPs), satisfy, in particular, the differential inclusion

$$x'(t) = v(t) - x(t), \quad v(t) \in \arg\max_{v \in V} [\nabla H(x(t))] \cdot v. \tag{3}$$

As the key element of proving Property 1, we prove the following attraction property of FSPs.

Property 2. (Attraction Property). The convergence

$$x(t) \to u^*, \quad t \to \infty$$

holds for any FSP x, with any initial state $x(0)$. (In fact, we prove uniform convergence on the initial states $x(0)$ from a bounded set. See Theorem 4 for precise formulation.)

The problem of asymptotic optimality of the gradient algorithm was studied in the recent papers by Kushner and Whiting (2002) and Agrawal and Subramanian (2002). In particular, in those papers the differential inclusion (3) was derived and the attraction Property 2 was proved, for somewhat different models and under certain constraints, which we now discuss in more detail.

In both Kushner and Whiting (2002) and Agrawal and Subramanian (2002) it is required that the arg max in (3) is a single point v, depending on $x(t)$ continuously. This makes (3) a differential equation of the form

$$x'(t) = f(x(t)), \tag{4}$$

with continuous $f(\cdot)$. In essence, this is the condition that rate region V is *strictly* convex. Such a condition does not hold, for example, if the set of all possible scheduling decisions is finite—a situation typical in many applications, including HDR (Bender et al. 2000, Jalali et al. 2000).

The proof of convergence Property 2 in Kushner and Whiting (2002) relies in an essential way on a further assumption that the dynamics described by (4) is cooperative. This roughly means that the vector function $f(u)$ is such that if we decrease all components of vector u except one, say u_i, then $f_i(u)$ will also decrease. This is another condition on the shape of V. It holds in the one-at-a-time case (when the service rate vectors $\mu^m(k)$ can have at most one nonzero component). However, it does not hold in many important cases when multiple users can be served simultaneously in one time slot, as in Viswanathan et al. (2003) (or in cross-bar switch models).

In Agrawal and Subramanian (2002), Property 2 is proved under the additional assumption that the initial state $x(0) \in V$. (In this case $H(x(t))$ is a natural Lyapunov function, which is no longer the case if $x(0)$ is outside of V.) This assumption is restrictive for the following two (related) reasons. First, such a form of Property 2 does not imply the asymptotic optimality Property 1. Second, in applications and especially wireless applications, the stationarity of the switch (channel) state process is only an approximation—the "stationary" distribution of this process is subject to both slow and sudden changes (for example, due to arrival and departure of users). Therefore, the ability of an algorithm to bring system to the optimal state from any initial state is in fact required for algorithm robustness in real applications.

Our proof of Property 2 follows a "geometric" approach (inherited from Stolyar 2004), which does not rely on the cooperative dynamics assumption or on the condition that rate region is strictly convex, or on the utility function H as the sole Lyapunov function. As a result, we are able to establish Property 2, and consequently the asymptotic optimality Property 1, in the desired generality, which allows in particular for features (a) and (b) of the model.

We also note that the proofs of Property 2 in Kushner and Whiting (2002) and Agrawal and Subramanian (2002) exclude the utility functions of the type $\sum \log(u_n)$, which take value $-\infty$ when some $u_n = 0$. Due to the wide use of such (somewhat less "benign") utility functions in applications, we include them in our results. To do this, for such utility functions we prove and use additional properties of FSPs, beyond the basic differential inclusion (3).

To summarize, the main contributions of this paper are as follows:

The key attraction property (Property 2) of FSPs under the gradient algorithm is proved for arbitrary initial state, and moreover, uniformly on the initial states from a bounded subset;

Property 2 is proved under very general (but common in applications) assumptions on the model and utility function;

The asymptotic optimality (Property 1) of the gradient algorithm is proved, based on the Property 2.

Kushner and Whiting (2002) and Agrawal and Subramanian (2002) also consider versions of the gradient algorithm, such that the parameter β is not fixed but rather is a decreasing vanishing function of time t. (A similar parameter update procedure was also used in earlier work (Liu et al. 2001, 2003) for the algorithms considered there.) Such versions of the algorithm are asymptotically optimal in the sense that long-term average throughputs converge to the exact optimal values u^*, as time $t \to \infty$. However, the smaller the initial value $\beta(0)$, the greater the convergence time. The algorithm with fixed β (which we study in this paper) is different in that it brings user throughputs into a (small) neighborhood of u^*. However, as we show, the convergence is within the time of the order of $O(1/\beta)$, uniformly on essentially any initial state. (This is because, in applications, usually there exists some a priori known upper bound a on the maximum possible service rate of any user at any time. As a result, as long as all $X_n(0) \leqslant a$, the components of $X(t)$ will always stay bounded by a.) This means that the algorithm is able to adjust to the "environment" changes without resetting the algorithm parameters. Such robustness of the algorithm with fixed β is very desirable in applications (as already mentioned earlier).

The rest of this paper is organized as follows. In §2, we introduce basic notations and conventions used in this paper. We introduce the formal model in §3, and in §4 define the system rate region V. The optimization problem for the average user rates is formulated in §5, and the gradient scheduling algorithm is defined in §6. The asymptotic optimality results (Theorem 1, Corollary 1, and Theorem 2), which formalize and strengthen Property 1, are presented in §7. Section 8 contains formal definition of FSPs, formulations of their basic properties, and the process level convergence result (Theorem 3). The main result on uniform attraction of FSPs (Theorem 4), which formalizes Property 2, is formulated and proved in §9. Section 10 contains proofs of the basic properties of FSPs (formulated in §8). In §11, we prove Theorems 1, 2, and 3. Finally, §12 contains concluding remarks, in particular, a discussion of our model assumptions and techniques.

2. Notation

We will use standard notations R and R_+ for the sets of real and real nonnegative numbers, respectively. Corresponding N-times product spaces are denoted R^N and R^N_+. The space R^N is viewed as a standard vector space, with elements $x \in R^N$ being row vectors $x = (x_1, \ldots, x_N)$. The scalar product of $x, y \in R^N$ is

$$x \cdot y \doteq \sum_{n=1}^{N} x_n y_n,$$

and the norm of x is

$$\|x\| \doteq \sqrt{x \cdot x}.$$

Sometimes we use notation

$$\rho(x, y) \doteq \|x - y\|$$

for the distance between vectors x and y, and notation

$$\rho(x, V) \doteq \inf_{y \in V} \rho(x, y)$$

for the distance between vector x and a set $V \subseteq R^N$.

Suppose that x is an N-dimensional vector with nonnegative components x_n, and V is a finite closed subset of R^N. Then,

$$\arg\max_{v \in V} x \cdot v$$

denotes the subset of vectors $v \in V$ with the maximum value of $x \cdot v$. We do not exclude the case when $x_n = +\infty$ for some n. For this case we use the convention that the product $(+\infty)\ c$ is equal to $-\infty$, 0, and $+\infty$, if c is negative, zero, and positive, respectively. Thus, if $x_n = +\infty$ and $V \subseteq R^N_+$, then any $u \in V$ with $u_n > 0$ belongs to $\arg\max_{v \in V} x \cdot v$.

For a scalar function $h(t)$ of a real variable t, we use the following notations for the lower and upper right derivative number:

$$\frac{d_+}{dt}h(t) = \liminf_{dt\downarrow 0} \frac{h(t+dt)-h(t)}{dt},$$
$$\frac{d^+}{dt}h(t) = \limsup_{dt\downarrow 0} \frac{h(t+dt)-h(t)}{dt}.$$

The abbreviation *u.o.c.* means *uniform on compact sets* convergence of functions.

3. The Model

We consider the following queueing system. (This system is same as in Stolyar 2004, where it is called generalized switch.) There is a finite set $N = \{1, 2, \ldots, N\}$ of queues (and corresponding "customer types") served by a switch. (We will use the same symbol N for both the set and its cardinality.) In this paper, we assume that queues always have sufficient "supply" of customers to serve.

The system operates in discrete time $t = 0, 1, 2, \ldots$. By convention, we will identify an (integer) time t with the unit time interval $[t, t+1)$, which will sometimes be referred to as the time slot t.

The switch has a finite set of switch states M. In each time slot, the switch is in one of the states $m \in M$, and the sequence of states $m(t)$, $t = 0, 1, 2, \ldots$, forms an irreducible (finite) Markov chain with stationary distribution $\{\pi_m,\ m \in M\}$,

$$\pi_m > 0 \quad \forall\, m \in M, \qquad \sum_{m\in M} \pi_m = 1.$$

When the switch is in state $m \in M$, a finite number of scheduling decisions can be made, which form a finite set $K(m)$; if a decision $k \in K(m)$ is chosen at time t, then the "amount" $\mu_n^m(k) \geqslant 0$ of "customers" of type $n \in N$ are served and depart the system at time $t+1$. We will denote by $\mu^m(k) \doteq (\mu_1^m(k), \ldots, \mu_N^m(k)) \in R_+^N$ the corresponding vector of service rates, and make a nondegeneracy assumption that for any type $n \in N$, there exists at least one state $m \in M$ and a decision $k \in K(m)$ such that $\mu_n^m(k) > 0$. (The values of $\mu_n^m(k)$ are real numbers, not necessarily integers. For exposition purposes, we sometimes refer to $\mu_n^m(k)$ as "number of customers.")

We denote by $\bar{\mu}$ the maximum of the values of $\mu_n^m(k)$ over all possible triples (n, m, k), and by $\underline{\mu} > 0$ the smallest strictly positive value of $\mu_n^m(k)$ over all (n, m, k).

4. System Rate Region

In this section, we define the system rate region $V \subseteq R_+^N$. Elements $v \in V$ represent all possible vectors of long-term average service rates, which can be provided by the system to the set of queues.

Suppose, for each of the switch states $m \in M$, a probability distribution $\phi_m = (\phi_{mk},\ k \in K(m))$ is fixed, which means that $\phi_{mk} \geqslant 0$ for all $k \in K(m)$, and $\sum_k \phi_{mk} = 1$. For such a set of distributions $\phi \doteq (\phi_m,\ m \in M)$, consider the following vector:

$$v(\phi) = \sum_{m\in M} \pi_m \sum_{k\in K(m)} \phi_{mk}\mu^m(k).$$

If we interpret ϕ_{mk} as the long-term average fraction of time slots when scheduling decision $k \in K(m)$ is chosen among the slots when the switch state is m, then $v(\phi)$ is the corresponding vector of long-term average service rates. (The *static service split* (SSS) scheduling rule from Stolyar 2004, parameterized by ϕ, can serve as an example of a rule that provides average service rate vector equal to $v(\phi)$; when switch is in state m, the SSS rule chooses one of the scheduling decisions $k \in K(m)$ randomly, according to the distribution ϕ_m.)

The system rate region V is defined as the set of all (average service rate) vectors $v(\phi)$ corresponding to all possible ϕ. Obviously, V is a closed bounded convex set in R_+^N, as a linear image of the closed bounded convex set of possible values of ϕ. Rate region V may turn out to be degenerate (i.e., have dimension less than N).

For future reference, note that the following equality holds for any fixed vector $\zeta \in R^N$:

$$\max_{v\in V} \zeta \cdot v = \sum_{m\in M} \pi_m \max_{k\in K(m)} \zeta \cdot \mu^m(k). \tag{5}$$

5. Optimization Problem for the Rate Allocation

Consider the following optimization problem:

$$\max \quad H(u) \tag{6}$$
$$\text{subject to} \quad u \in V, \tag{7}$$

where $H(u)$ is a strictly concave utility function. We will consider two types of utility functions, defined as follows.

Type (I) Utility Function. $H(u)$ is a continuous strictly concave function on R_+^N. Moreover, $H(u)$ is continuously differentiable, i.e., the gradient ∇H is finite and continuous everywhere in R_+^N.

Type (II) Utility Function. $H(u) = \sum_n H_n(u_n)$, where each $H_n(u_n)$ is a strictly concave continuously differentiable function, defined for all $u_n > 0$, and such that $H_n(u_n) \downarrow -\infty$ as $u_n \downarrow 0$. (The definition implies $H_n'(u_n) \uparrow +\infty$ as $u_n \downarrow 0$. We adopt the conventions that $H_n(0) = -\infty$ and $H_n'(0) = +\infty$.)

Remark 1. The assumption of strict concavity of utility function H is not essential for the main results of this paper. See §12 for a more detailed comment.

Type (I) utility functions are more "benign." The examples of interest include $H(u) = \sum_n \log(c_n + u_n)$ or $H(u) = \sum_n c_n u_n$, or a "mixture" of the two (with some log and some linear components), with $c_n > 0$ being fixed constants. (As noted above, strict concavity of H is not essential.) The system dynamics under the gradient algorithm with a type (I) utility function is rather simple. Type (II) functions have additive form, but the components are not bounded and have infinite derivative at 0, which somewhat complicates analysis. The PF utility function $\sum_n \log(u_n)$, widely used in applications, is of type (II). This is one of the reasons why we do not want to exclude such utility functions from our analysis.

Proposition 1. *Suppose that the utility function H is of either type* (I) *or type* (II). *Then, solution u^* to problem* (6)–(7) *exists and is unique.*

Proof. The proof is trivial for type (I) utility function H, because V is a bounded convex compact set and H is strictly concave. For type (II) utility function H, its domain needs to be first restricted to a compact set $V \cap [\eta, \infty)^N$, with $\eta > 0$ small enough so that the supremum of H is certainly attained within this set. □

Remark 2. Our definition of the utility function does not require that H is nondecreasing in each argument, so the optimal point u^* is not necessarily a point on the outer ("north-east") boundary of V. However, in many applications of interest this would be the case.

6. Gradient Scheduling Algorithm

Now consider the random process describing behavior of the system under the following scheduling algorithm which is called the *gradient algorithm*.

Gradient Algorithm

If at time t switch is in state m, choose a scheduling decision

$$k(t) \in \arg\max_{k \in K(m)} [\nabla H(X(t))] \cdot \mu^m(k),$$

where $X(t) = (X_1(t), \ldots, X_N(t))$ is the vector of the average service rate estimates, which is updated as follows:

$$X(t+1) = (1-\beta)X(t) + \beta\mu^{m(t)}(k(t)),$$

with $\beta > 0$ being a (small) parameter. The initial vector $X(0) \in R_+^N$ is arbitrary.

Let us denote

$$S(t) \doteq (X(t), m(t)).$$

Our assumptions imply that $S = \{S(t),\ t = 0, 1, \ldots\}$ is a discrete time Markov chain with the state space $R_+^N \times M$.

7. Asymptotic Optimality of the Gradient Algorithm

This section presents the results showing that, as $\beta \downarrow 0$, the average user service rates converge to the optimal solution u^* of problem (6)–(7). Moreover, this convergence is uniform (in the sense specified in Theorem 1).

First, we need to define the asymptotic regime. From this point on in the paper, we consider a sequence of processes S^β, indexed by the value of parameter β, with $\beta \downarrow 0$ along a sequence $\mathcal{B} = \{\beta_j,\ j = 1, 2, \ldots\}$ such that $\beta_j > 0$ for all j. The initial state $S^\beta(0) = (X^\beta(0), m^\beta(0))$ is fixed for each $\beta \in \mathcal{B}$. (Here and below, the processes and variables pertaining to a fixed parameter β will be supplied the upper index β. Expression $\beta \downarrow 0$ means that β converges to 0 along the sequence $\mathcal{B}$, unless otherwise specified.)

The probability law of the Markov chain $m^\beta(\cdot)$ describing the switch state process is the same for each β. Let us denote by $U^\beta(l_1, l_2)$ the vector of average service rates in the interval $[l_1, l_2]$. Namely, for a pair of integers $1 \leqslant l_1 \leqslant l_2$,

$$U^\beta(l_1, l_2) \doteq \frac{1}{l_2 - l_1 + 1} \sum_{j=l_1}^{l_2} D^\beta(j),$$

where $D^\beta(l) = \mu^{m(l-1)}(k(l-1))$ is the vector of numbers of customers that were served in slot $l-1$ (with $k(l-1) \in K(m(l-1))$ being the scheduling decision chosen in slot $l-1$).

Theorem 1. *Let A be a bounded subset of R_+^N. Then, for any $\epsilon > 0$ there exist $T > 0$ and $T^* > 0$* (*both depending on ϵ and A*), *such that*

$$\lim_{\beta\downarrow 0} \sup_{X^\beta(0)\in A,\, l_1 > T/\beta,\, l_2 - l_1 > T^*/\beta} \|EU^\beta(l_1, l_2) - u^*\| < \epsilon. \tag{8}$$

Corollary 1. *Suppose that for each β we consider a stationary version of the process S^β, in which case $ED^\beta(1) = EU^\beta(l_1, l_2)$* (*for any integer $1 \leqslant l_1 \leqslant l_2$*) *is simply the average service rate vector. Then,*

$$\lim_{\beta\downarrow 0} ED^\beta(1) = u^*. \tag{9}$$

Proof. The proof of Theorem 1 is presented in §11. It is obtained using the following uniform convergence in probability result for X^β.

Theorem 2. *Let A be a bounded subset of R_+^N. Then, for any $\epsilon > 0$ there exists $T > 0$* (*depending on ϵ and A*) *such that*

$$\lim_{\beta\downarrow 0} \sup_{X^\beta(0)\in A,\, t > T/\beta} P\{\|X^\beta(t) - u^*\| > \epsilon\} = 0. \tag{10}$$

Proof. Presented in §11.

8. Fluid Scaled Processes and Fluid Sample Paths

Consider the sequence of systems introduced in the previous section, with parameter $\beta \downarrow 0$. In this section, for each β we consider the random processes (describing system evolution) under fluid scaling and study asymptotic behavior of the sequence of fluid-scaled processes as $\beta \downarrow 0$.

In the rest of this section, we will specify the processes describing system evolution, formally define fluid-scaled processes, and introduce the notion of FSPs, which are, roughly speaking, possible limits of realizations of fluid-scaled processes. The section is concluded by the process level convergence result, Theorem 3, which (roughly) shows that fluid-scaled processes converge to processes with realizations being FSPs.

8.1. Fluid Scaled Processes

Let us extend the definition of $X^\beta(t)$ to continuous time $t \in R_+$ by adopting the convention that $X^\beta(t)$ is constant within each time slot $[l, l+1)$. For each β, let us define some additional random functions, associated with the system evolution. We define them for continuous time $t \in R_+$ as well, although they (just as $X^\beta(t)$) are constant within each time slot $[l, l+1)$. Let

$$\hat{F}_n^\beta(t) \doteq \sum_{l=1}^{\lfloor t \rfloor} D_n^\beta(l)$$

denote the number of type n customers that were served by time $t \geqslant 0$ (i.e., in the interval $[0, t]$), where $D_n^\beta(l) = \mu_n^{m(l-1)}(k(l-1))$ is the number of type n customers served in slot $l-1$, and $k(l-1) \in K(m(l-1))$ is the scheduling decision chosen in slot $l-1$. Denote by $G_m^\beta(t)$ the total number of time slots by (and including) time $t-1$, when the server was in state m; and by $\hat{G}_{mk}^\beta(t)$ the number of time slots by (and including) time $t-1$, when the server state was m and the scheduling decision $k \in K(m)$ was chosen.

Obviously, for any $n \in N$, $m \in M$, and $k \in K(m)$, we have

$$\hat{F}_n^\beta(0) = 0, \qquad G_m^\beta(0) = 0, \qquad \hat{G}_{mk}^\beta(0) = 0,$$

and we have the following relations:

$$G_m^\beta(t) = \sum_{k \in K(m)} \hat{G}_{mk}^\beta(t), \quad t \geqslant 0,$$

$$\hat{F}_n^\beta(t) = \sum_{m \in M} \sum_{k \in K(m)} \mu_n^m(k) \hat{G}_{mk}^\beta(t), \quad t \geqslant 0.$$

Recall that for any $n \in N$ and all integer $l \geqslant 1$,

$$X_n^\beta(l) = \beta D_n^\beta(l) + (1-\beta) X_n^\beta(l-1). \tag{11}$$

We define the process Z^β, describing system evolution, as

$$Z^\beta = (X^\beta, \hat{F}^\beta, G^\beta, \hat{G}^\beta),$$

where

$$X^\beta = (X^\beta(t) = (X_1^\beta(t), \ldots, X_N^\beta(t)),\ t \geqslant 0),$$
$$\hat{F}^\beta = (\hat{F}^\beta(t) = (\hat{F}_1^\beta(t), \ldots, \hat{F}_N^\beta(t)),\ t \geqslant 0),$$
$$G^\beta = ((G_m^\beta(t),\ m \in M),\ t \geqslant 0),$$
$$\hat{G}^\beta = ((\hat{G}_{mk}^\beta(t),\ m \in M,\ k \in K(m)),\ t \geqslant 0).$$

For each β, consider the following process z^β, which is a fluid-scaled version of process Z^β:

$$z^\beta = (x^\beta, \hat{f}^\beta, g^\beta, \hat{g}^\beta),$$

where x^β is obtained by time scaling only

$$x^\beta(t) \doteq X^\beta(t/\beta),$$

and the other components by time and space scaling

$$\hat{f}^\beta(t) \doteq \beta \hat{F}^\beta(t/\beta), \qquad g^\beta(t) \doteq \beta G^\beta(t/\beta),$$
$$\hat{g}^\beta(t) \doteq \beta \hat{G}^\beta(t/\beta).$$

Note that the component functions of z^β are piecewise constant with a "time slot" of the length β.

Recall that the probability law of the Markov chain $m^\beta(\cdot)$ describing the switch state process is the same for each β, and therefore for any $m \in M$ and any fixed $0 \leqslant t_1 \leqslant t_2 < \infty$, we have the law of large numbers:

$$g_m^\beta(t_2) - g_m^\beta(t_1) \to \pi_m(t_2 - t_1) \quad \text{in probability.} \tag{12}$$

8.2. Fluid Sample Paths Under the Gradient Algorithm

We now define FSPs, which are fixed trajectories arising as possible limits of sequences (on β) of z^β realizations, given the realizations of g^β satisfy the functional law of large numbers (see (13) below).

Definition. A fixed set of functions $z = (x, \hat{f}, g, \hat{g})$ we will call a FSP if there exists a sequence $\mathscr{B}_0$ of positive values of β such that $\beta \downarrow 0$, and a sequence of sample paths (of the corresponding processes) $\{z^\beta\}$ such that (as $\beta \downarrow 0$ along sequence $\mathscr{B}_0$)

$$z^\beta \to z \quad u.o.c.,$$

and in addition

$$\|x(0)\| < \infty,$$
$$(g_m^\beta(t),\ t \geqslant 0) \to (\pi_m t,\ t \geqslant 0) \quad u.o.c.\ \forall\, m \in M. \tag{13}$$

Remark. A sequence $\mathscr{B}_0$ whose existence is required in the above definition may be completely unrelated to the sequence $\mathscr{B}$ we introduced earlier.

If $x = (x(t),\ t \geqslant 0)$ is a component of at least one FSP, this function x in itself we will often call an FSP.

The following two lemmas describe some properties of FSPs. We note that these properties may not characterize the family of FSPs completely. However, they are the only ones we will need in the rest of the paper.

Lemma 1. *For any FSP z, all its component functions are Lipschitz continuous in $[0, \infty)$, with the Lipschitz constant $C + \|x(0)\|$, where $C > 0$ is a fixed constant depending only on the system parameters. In addition, the functions $g_m(\cdot)$, $\hat{g}_{mk}(\cdot)$, and $\hat{f}(\cdot)$ are nondecreasing, and satisfy the following relations*:

$$g_m(t) = \pi_m t, \quad t \geqslant 0,\ m \in M, \tag{14}$$

$$g_m(t) = \sum_{k \in K(m)} \hat{g}_{mk}(t), \quad t \geqslant 0,\ m \in M, \tag{15}$$

$$\hat{f}(t) = \sum_{m \in M} \sum_{k \in K(m)} \mu^m(k) \hat{g}_{mk}(t), \quad t \geqslant 0. \tag{16}$$

Proof. The proof is in §10.

Because all component functions of an FSP are Lipschitz, they are absolutely continuous; therefore almost all points $t \in R_+$ (with respect to Lebesgue measure) are such that all component functions of z have proper derivatives (that is, both right and left derivatives exist and are equal). We will call such points *regular*.

Lemma 2. *The family of FSPs z satisfies the following additional properties*:

(i) *For almost all $t \geqslant 0$ (with respect to Lebesgue measure), we have*

$$x'(t) = v(t) - x(t), \tag{17}$$

where

$$v(t) \doteq (d/dt)\hat{f}(t) = \sum_{m \in M} \sum_{k \in K(m)} \hat{g}'_{mk}(t) \mu^m(k) \tag{18}$$

satisfies the condition

$$v(t) \in \arg\max_{v \in V} (\nabla H(x(t))) \cdot v. \tag{19}$$

(ii) *"Boundary condition." Suppose that the utility function H is of type* (II). *Suppose, for some $t \geqslant 0$ and some subset $B \subseteq N$, $x_i(t) = 0$ if $i \in B$ and $x_i(t) > 0$ if $i \in N \backslash B$. Then,*

$$\frac{d_+}{dt} \sum_{i \in B} x_i(t) \geqslant c > 0,$$

where $c > 0$ is a fixed constant depending only on the system parameters.

(iii) *"Compactness." If a sequence of FSPs $z^{(j)} \to z$ u.o.c. as $j \to \infty$, then this z is also an FSP. Moreover, if a sequence $x^{(j)} \to x$ u.o.c. as $j \to \infty$, where $x^{(j)}$ are components of some FSPs $z^{(j)}$, then x is a component of some FSP z.*

Proof. The proof is in §10.

Property (i) in the above lemma says that any FSP x is a solution to the differential inclusion

$$x'(t) = v(t) - x(t), \quad v(t) \in \arg\max_{v \in V} [\nabla H(x(t))] \cdot v. \tag{20}$$

This is the only FSP property needed to prove our main results for type (I) utility functions. To deal with type (II) utility functions, we will use properties (ii) and (iii) as well.

8.3. Process Level Convergence

We will view random processes x^β as processes with realizations in the Skorohod space $D_{R^N}[0, \infty)$ of functions with domain $[0, \infty)$, taking values in R^N, which are right-continuous and have left limits. The Skorohod topology and corresponding Borel σ-algebra on $D_{R^N}[0, \infty)$ are defined in the usual way. (See Ethier and Kurtz 1986 for the definition.)

Theorem 3. *Consider the sequence of processes $\{x^\beta\}$ with $\beta \downarrow 0$ such that $x^\beta(0) \to x(0)$, where $x(0) \in R^N_+$ is a fixed vector. Then, the sequence $\{x^\beta\}$ is relatively compact and any weak limit of this sequence (i.e., a process obtained as a weak limit of a subsequence of $\{x^\beta\}$) is a process with sample paths being FSPs with probability* 1.

Theorem 3 is analogous to the corresponding process convergence results in Kushner and Whiting (2002) and Agrawal and Subramanian (2002). Note, however, that Theorem 3 is stronger in the sense that it claims that the limiting process is concentrated on the family of FSPs, which is "narrower" than simply the family of solutions of the differential inclusion (20). This subtlety is important because we use this theorem in conjunction with Theorem 4, which holds for the FSPs but (in the case of Type (II) utility function) may or may not hold for a wider class of solutions of (20). This also affects the proof.

Proof. The proof is in §11.

9. Uniform Attraction of Fluid Sample Paths

The following FSP uniform attraction theorem is the key result of this paper.

Theorem 4. *Suppose that H is a utility function of either type* (I) *or type* (II). *Then, the family of FSPs x has the following property*: *For any bounded subset $A \subset R^N_+$,*

$$x(t) \to u^*, \quad t \to \infty, \tag{21}$$

uniformly on $x(0) \in A$.

Note that Theorem 4 can be viewed as a "deterministic analog" of Theorem 2. In fact, as we will see, the proof of Theorem 2 easily follows from Theorem 4 along with Theorem 3.

PROOF. The proof of Theorem 4 is the subject of the rest of this section. Throughout this proof, without loss of generality, we assume that the set A is closed, and therefore compact.

9.1. Uniform Convergence to the Set V

Lemma 3 formalizes a simple, but nevertheless important, general observation that if derivative $x'(t)$ is always a vector from $x(t)$ to an arbitrary point within a convex set V, then the distance from $x(t)$ to V can only decrease. Note that in the differential inclusion (22) in Lemma 3 it is required only that $v(t) \in V$. (There is no condition $v(t) \in \arg\max[\nabla H(x(t))] \cdot v$.)

LEMMA 3. *Suppose that V is a convex bounded closed subset of R^N. Suppose that a vector function $(x(t),\ t \geqslant 0)$, taking values in R^N, is Lipschitz continuous, satisfying the following differential inclusion for almost all $t \geqslant 0$:*

$$x'(t) = v(t) - x(t), \quad v(t) \in V. \tag{22}$$

Then, the distance $\rho(x(t), V)$ between $x(t)$ and the set V is a Lipschitz continuous nonincreasing function, and moreover, for almost all $t \geqslant 0$,

$$\frac{d}{dt}\rho(x(t), V) \leqslant -\rho(x(t), V), \tag{23}$$

which implies that

$$\rho(x(t), V) \leqslant \rho(x(0), V)e^{-t}.$$

As we see, under the conditions of this lemma, the functions satisfying (22) converge to set V uniformly on the initial states from a bounded set. For ease of reference we record this fact as Corollary 2 below.

For $\epsilon \geqslant 0$, let us denote by V_ϵ, the ϵ-thickening of a subset $V \subseteq R^N$, namely

$$V_\epsilon \doteq \{y \in R^N \mid \rho(y, V) \leqslant \epsilon\}.$$

COROLLARY 2. *Let arbitrary bounded set $A \subset R^N$ and arbitrary $\epsilon > 0$ be fixed. Then, there exists $T_1 = T_1(\epsilon, A)$ such that for any $(x(t),\ t \geqslant 0)$ satisfying the conditions of Lemma 3 and $x(0) \in A$,*

$$x(t) \in V_\epsilon \quad \forall\, t \geqslant T_1.$$

PROOF OF LEMMA 3. Lipschitz continuity of $\rho(x(t), V)$ follows from the fact that $x(t)$ is Lipschitz continuous. To prove the rest of the statement, it suffices to show that (23) holds for $t = \theta$, where $\theta > 0$ is a regular point such that $\rho(x(\theta), V) > 0$. Consider the point ξ which is the (unique) point of V closest to $x(\theta)$, i.e., $\rho(x(\theta), \xi) = \rho(x(\theta), V)$

Figure 1. Proof of Lemma 3.

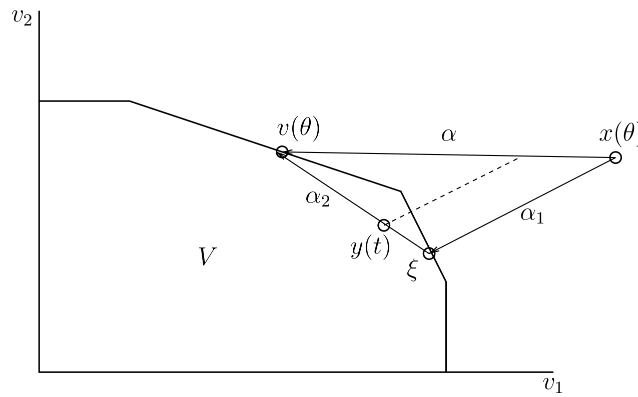

(see Figure 1). Consider the vectors

$$\alpha_1 = \xi - x(\theta), \qquad \alpha_2 = v(\theta) - \xi,$$

and

$$\alpha = v(\theta) - x(\theta) = \alpha_1 + \alpha_2.$$

Consider the auxiliary linear function

$$y(t) = \xi + \alpha_2(t - \theta), \quad t \geqslant \theta,$$

and note that $y(t) \in V$ for all sufficiently small increments $(t - \theta) \geqslant 0$. Then, we can write

$$\frac{d^+}{dt}\rho(x(t), V) \leqslant \frac{d}{dt}\rho(x(t), y(t)) = -\|\alpha_1\| = -\rho(x(\theta), V). \quad \square$$

9.2. Proof of Theorem 4 for Type (I) Utility Function

From the FSP properties described in Lemma 2, this proof uses only property (i), i.e., the differential inclusion (20).

The key idea of this proof is that as long as $x(t)$ is outside of a δ-neighborhood $O_\delta(u^*)$ of point u^*, but is close enough to V (which is guaranteed by Lemma 3, after some initial uniformly bounded time), then $H(x(t)) < H(u^*)$. As a result, the derivative $(d/dt)H(x(t))$ must be strictly positive. This is because

$$\frac{d}{dt}H(x(t)) = [\nabla H(x(t))] \cdot (v(t) - x(t)) \geqslant [\nabla H(x(t))] \cdot (u^* - x(t))$$

(where the inequality follows from (20)), and the right-hand side is strictly positive by concavity of H and inequality $H(x(t)) < H(u^*)$. Additional estimates show that $(d/dt)H(x(t))$ is not only positive but is also bounded away from 0 (for $x(t)$ satisfying the above conditions). The detailed proof is as follows.

Let us fix arbitrary $\delta > 0$. Then, fix $\epsilon > 0$ small enough so that $h^* < H(u^*)$, where h^* is the maximum of $H(y)$ over the subset $U_{\delta,\epsilon} \doteq V_\epsilon \backslash O_\delta(u^*)$. Using concavity and smoothness of H, it is easy to verify that at any point $y \in U_{\delta,\epsilon}$, the directional derivative of H in the direction of point u^* is at least $(H(u^*) - h^*)/\delta$, namely

$$[\nabla H(y)] \cdot (u^* - x) \geqslant \frac{H(u^*) - h^*}{\delta} \|u^* - y\| \geqslant H(u^*) - h^*. \tag{24}$$

By Corollary 2 (from Lemma 3), there exists T_1 such that $x(t) \in V_\epsilon$ for all $t \geqslant T_1$. Then, for any regular point $t \geqslant T_1$ such that $x \in U_{\delta,\epsilon}$ we have

$$\frac{d}{dt} H(x(t)) = [\nabla H(x(t))] \cdot (v(t) - x(t)) \geqslant [\nabla H(x(t))] \cdot (u^* - x(t)) \geqslant H(u^*) - h^*.$$

Because $H(u^*) - h^* > 0$, we see that there exists a constant $T_2 \geqslant T_1$ such that (uniformly on $x(0) \in A$)

$$t_2 = \min\{t \geqslant T_1 \mid x(t) \in \bar{O}_\delta(u^*)\} \leqslant T_2,$$

i.e., $x(t)$ must "hit" the closed neighborhood $\bar{O}_\delta(u^*)$ no later than at time T_2.

For all $t \geqslant t_2$, we must have $H(x(t)) \geqslant h_*$, where h_* is the minimum of $H(x)$ over $x \in V_\epsilon \cap \bar{O}_\delta(u^*)$. (This is because for $t \geqslant T_1$, $H(x(t))$ cannot decrease outside of $\bar{O}_\delta(u^*)$.) Because our $\delta > 0$ could be chosen arbitrarily small, we observe that (uniformly on $x(0) \in A$)

$$\liminf_{t\to\infty} H(x(t)) = H(u^*). \tag{25}$$

Finally, there exists $T_3 \geqslant T_2$ such that (uniformly on $x(0) \in A$) for all $t \geqslant T_3$,

$$x(t) \in V_\epsilon \cap \bar{O}_\delta(u^*),$$

because otherwise (25) would not hold. □

9.3. Proof of Theorem 4 for Type (II) Utility Function

This proof uses all three FSP properties (i)–(iii) established in Lemma 2.

The key idea of the proof is the same as that for type (I) utility function. However, a complication here is that before we can use this idea, we need to establish that, after some finite time, $x(t)$ will not only stay close to V but will also stay away from the boundary of the orthant R_+^N. This "additional work" is done below in Lemmas 4 through 7, and this is where properties (ii)–(iii) of Lemma 2 are used.

As in the proof for type (I) utility function, let us fix arbitrary $\delta > 0$, and then fix $\epsilon > 0$ small enough so that $h^* < H(u^*)$, where h^* is the maximum of $H(x)$ over the subset $U_{\delta,\epsilon} \doteq V_\epsilon \backslash O_\delta(u^*)$. By Corollary 2 (from Lemma 3), there exists T_1' such that $x(t) \in V_\epsilon$ for all $t \geqslant T_1'$. (The above statement and other statements in this proof hold uniformly on $x(0) \in A$, unless specified otherwise.)

The rest of the proof consists of the following sequence of steps.

Lemma 4. *At any regular point $t > 0$, we have $x_i(t) > 0$ for all i.*

Proof. Suppose not. Consider the subset $B \subseteq N$ of types i such that $x_i(t) = 0$. Because t is a regular point, we must have $x_i'(t) = 0$. (Otherwise $x_i(\cdot)$ would be negative just before or just after time t.) Consequently, $\sum_{i\in B} x_i'(t) = 0$. This, however, contradicts Lemma 2(ii). □

Lemma 5. *For all $t > T_1'$, we have $x_i(t) > 0$ for all i.*

Proof. We can choose a regular point $t' > T_1'$, arbitrarily close to T_1'. We have $H(x(t')) > -\infty$ because all $x_i(t') > 0$. Similar to the way it is established in the proof for type (I) utility function, we see that $H(x(t))$ has a strictly positive derivative at any regular point $t \in [T_1', \infty)$ such that $H(x(t)) > -\infty$ and $x(t) \in U_{\delta,\epsilon}$. This implies that $H(x(t)) > -\infty$ for all $t \geqslant t_1'$. □

Lemma 6. *For any $T_1 > T_1'$, there exists $\eta_1 > 0$ such that we have $x_i(T_1) \geqslant \eta_1$ for all i.*

Proof. Suppose not. Then, there exists a sequence of FSPs (with initial states within the compact set A) converging u.o.c. to a function $x = (x(t), t \geqslant 0)$ such that for at least one i, $x_i(T_1) = 0$. By Lemma 2(iii), this function x is also an FSP with $x(0) \in A$. This contradicts Lemma 5. □

Lemma 7. *For any $T_1 > T_1'$, there exists $\eta > 0$ such that we have $x_i(t) \geqslant \eta$ for all i and all $t \geqslant T_1$.*

Proof. Let us fix arbitrary $T_1 > T_1'$ and choose a small $\eta_1 > 0$ such that the statement of Lemma 6 holds. Without loss of generality, we can assume that η_1 is small enough so that

$$H(\eta_1, \ldots, \eta_1) < \min_{x \in \bar{O}_\delta} H(x).$$

Then, again using the fact that $H(x(t))$ has a strictly positive derivative at any regular point $t \in [T_1', \infty)$ such that $H(x(t)) > -\infty$ and $x(t) \in U_{\delta,\epsilon}$, it is easy to see that

$$H(x(t)) \geqslant H(\eta_1, \ldots, \eta_1) \quad \forall\, t \geqslant T_1.$$

This easily implies that all $x_i(t)$ must stay separated from 0 in $[T_1, \infty)$. □

Thus, we have proved that there exists $T_1 > 0$ such that for all $t \geqslant T_1$, $x(t) \in V_{\epsilon,\eta} \doteq V_\epsilon \cap \{x_i \geqslant \eta, \forall\, i\}$. Note that $H(x)$ is bounded on $V_{\epsilon,\eta}$. Given this fact, the rest of the proof is virtually identical to the proof for the type (I) utility function. (Namely, the choice of constants T_2 and T_3, and the proofs of the corresponding properties repeat those in the proof for type (I) utility function verbatim.)

The proof of Theorem 4 for the type (II) function H is complete.

10. Proofs of Lemmas 1 and 2

Proof of Lemma 1. Consider a fixed FSP z and a sequence of paths z^β "defining it." For each β, Z^β is the "unscaled" path from which z^β is obtained.

To prove the Lipschitz property of the components of z, let us recall the $X_n^\beta(\cdot)$ update equation (11):

$$X_n^\beta(l) = \beta D_n^\beta(l) + (1-\beta) X_n^\beta(l-1). \tag{26}$$

Applying this equation iteratively for $l = 1, 2, \ldots,$ we obtain

$$X_n^\beta(l) = \widetilde{X}_n^\beta(l) + (1-\beta)^l X_n^\beta(0), \tag{27}$$

where

$$\widetilde{X}_n^\beta(l) \doteq \sum_{j=0}^{l-1} \beta(1-\beta)^j D_n^\beta(l-j). \tag{28}$$

We see that for any $l \geqslant 1$,

$$\widetilde{X}_n^\beta(l) \leqslant \bar{\mu} = \max_{n,m,k} \mu_n^m(k), \tag{29}$$

because $\widetilde{X}_n^\beta(l)$ is simply the weighted sum of the values of $D_n^\beta(0), \ldots, D_n^\beta(l)$ (with positive weights summing up to at most 1), and each of those values is of course bounded above by $\bar{\mu}$. We can notice (for future reference) that expressions (27)–(29) imply that

$$X_n^\beta(l) \leqslant \max\{\bar{\mu}, X_n^\beta(0)\} \quad \forall\, l \geqslant 0. \tag{30}$$

Let us rewrite (26) as follows:

$$X_n^\beta(l) - X_n^\beta(l-1) = \beta D_n^\beta(l) - \beta X_n^\beta(l-1). \tag{31}$$

We see that

$$|X_n^\beta(l) - X_n^\beta(l-1)| \leqslant \beta(D_n^\beta(l) + \widetilde{X}_n^\beta(l-1) + X_n^\beta(0)) \leqslant \beta(2\bar{\mu} + X_n^\beta(0)). \tag{32}$$

All other components of the process Z^β are nondecreasing, and we trivially have for all integer $l \geqslant 1$ (and any n, m, k),

$$\hat{F}_n^\beta(l) - \hat{F}_n^\beta(l-1) \leqslant \bar{\mu},$$

$$G_m^\beta(l) - G_m^\beta(l-1) \leqslant 1,$$

$$\hat{G}_{mk}^\beta(l) - \hat{G}_{mk}^\beta(l-1) \leqslant 1.$$

These relations along with (32) imply that all components of an FSP z are Lipschitz in $[0, \infty)$ with the specified Lipschitz constant, if we choose C large enough.

The facts that the functions $g_m(\cdot)$, $\hat{g}_{mk}(\cdot)$, and $\hat{f}(\cdot)$ are nondecreasing, and relations (14)–(16) as well, follow directly from the definitions involved. □

Proof of Lemma 2. As in the proof of Lemma 1, consider a fixed FSP z and a sequence of paths z^β "defining it," along with their "unscaled" versions Z^β.

To prove properties (i) and (ii), let us first derive the basic integral equation for $x(\cdot)$ (Equation (35) below). The derivation of this equation is essentially same as in Kushner and Whiting (2002) and Agrawal and Subramanian (2002); we present it for completeness. Let us sum up Equations (31) for $l = 1, \ldots, j$. We obtain

$$X_n^\beta(j) - X_n^\beta(0) = \beta \sum_{l=1}^{j} D_n^\beta(l) - \sum_{l=1}^{j} \beta X_n^\beta(l-1). \tag{33}$$

Switching to scaled processes, for all integer $j \geqslant 0$ we have

$$x_n^\beta(\beta j) - x_n^\beta(0) = \hat{f}_n^\beta(\beta j) - \int_0^{\beta j} x_n^\beta(\xi)\, d\xi. \tag{34}$$

Using the fact that all limiting functions x_n and $\hat{f}_n$ are Lipschitz continuous and we have u.o.c. convergence $z^\beta \to z$, we finally obtain the desired integral equation

$$x_n(t) - x_n(0) = \hat{f}_n(t) - \int_0^t x_n(\xi)\, d\xi, \quad t \in R_+. \tag{35}$$

Because both $x(\cdot)$ and $\hat{f}(\cdot)$ are Lipschitz continuous, Equation (17) holds for every regular point $t > 0$.

Now, let us prove property (ii). (After that we will prove (19), i.e., the rest of property (i).) Suppose that $B \subset N$. (The proof for the case $B = N$ is similar and in fact simpler than that for $B \subset N$.) Because $x(\cdot)$ is continuous, $H_n'(x(t)) = +\infty$ for $n \in B$ and $H_n'(x(t))$ is finite for $n \in B \subset N$, we see that in a small neighborhood of time t, the derivatives $H_n'(x(\xi))$ for $n \in B$ are large compared to those for $n \in B \subset N$. More precisely, the following observation is true:

Observation 1. There exist sufficiently small fixed $\Delta > 0$ and a constant $\eta > 1$ (both depending on the values $x_n(t)$ for $n \in N \backslash B$) such that for any $\xi \in [t-\Delta, t+\Delta]$, any $i \in B$, and any $n \in N \backslash B$,

$$H_i'(x_i(\xi)) > 0$$

and

$$\frac{H_i'(x_i(\xi))}{|H_n'(x_n(\xi))|} \geqslant \eta \frac{N\bar{\mu}}{\underline{\mu}}.$$

(Recall that $\underline{\mu}$ denotes the smallest *strictly positive* value of $\mu_n^m(k)$ over all n, m, k.)

This means that if we consider the "unscaled" paths Z^β, then for all sufficiently small β and any (positive integer) time slot l within interval $[t/\beta, t/\beta + \Delta/\beta]$, we have the following property.

Observation 2. If for the state $m = m(l-1)$ the subset of decisions $k \in K(m)$ such that $\mu_i^m(k) > 0$ for at least one user $i \in B$ is nonempty, then a decision from this subset will be chosen in slot $l-1$; as a result

$$\sum_{i \in B} D_i^\beta(l) \geqslant \underline{\mu}.$$

Let us pick a state m such that $\mu_i^m(k) > 0$ for at least one pair of $k \in K(m)$ and $i \in B$. (Such a state m exists.) By the definition of a FSP, we have the uniform convergence (13) for each $m \in M$. Then, using Observation 2 (and recalling the definition of components $\hat{f}_n$ of an FSP), we obtain estimate

$$\sum_{i\in B} \hat{f}_i(t+\Delta) - \sum_{i\in B} \hat{f}_i(t) \geqslant \pi_m \Delta \underline{\mu} > 0,$$

which implies (because it holds for any small $\Delta > 0$)

$$\frac{d_+}{dt} \sum_{i\in B} \hat{f}_i(t) \geqslant \pi_m \underline{\mu} > 0. \tag{36}$$

However, from the integral equation (35), continuity of $x(\cdot)$, and the fact that $x_i(t) = 0$ for all $i \in B$, we see that

$$\frac{d_+}{dt} \sum_{i\in B} x_i(t) = \frac{d_+}{dt} \sum_{i\in B} \hat{f}_i(t).$$

This (along with (36)) proves

$$\frac{d_+}{dt} \sum_{i\in B} x_i(t) \geqslant \pi_m \underline{\mu} > 0.$$

Because there is only a finite number of subsets $B \subseteq N$, property (ii) has been proved.

Let us prove (19). Consider a fixed regular point $t > 0$. It will suffice to prove that (19) holds for this t. We first note that in the case of a type (II) utility function, we must have $x_n(t) > 0$ for all $n \in N$. Otherwise (because t is regular), we would have $x_n'(t) = 0$ for any n with $x_n(t) = 0$, which is impossible due to property (ii). Thus, for either utility function type, the gradient $\nabla H(x(t))$ is finite, and therefore $\nabla H(y)$ is bounded (and continuous) in a neighborhood of point $y = x(t)$. Given this, the rest of the proof of (19) is analogous to the proof of Lemma 5(ii) in Stolyar (2004). Namely, because $x^\beta(\xi)$ is close to $x(t)$ when $|\xi - t|$ and β are small, the following observation is true.

Observation 3. For any $\epsilon > 0$, there exists a sufficiently small $\Delta > 0$, such that for all "unscaled" paths Z^β with sufficiently small β, we have the following property. For any (positive integer) time slot l within interval $[t/\beta, (t+\Delta)/\beta]$,

$$|\nabla H(X^\beta(l-1)) \cdot D^\beta(l) - a_{m(l-1)}| \leqslant \epsilon,$$

where we use the notation

$$a_m \doteq \max_{k\in K(m)} \nabla H(x(t)) \cdot \mu^m(k), \quad m \in M.$$

From this observation, we have

$$\left| \sum_{t/\beta \leqslant l \leqslant (t+\Delta)/\beta} \nabla H(X^\beta(l-1)) \cdot D^\beta(l) - \sum_{t/\beta \leqslant l \leqslant (t+\Delta)/\beta} a_{m(l-1)} \right| \leqslant \epsilon\Delta/\beta + O(1),$$

where $O(1)$ denotes a term with absolute value bounded above by C as $\beta \to 0$, with $C > 0$ being a fixed constant. From the last display, multiplied by β, it is easy to obtain the following estimate for the fluid-scaled paths:

$$\left| \int_t^{t+\Delta} \nabla H(x^\beta(\xi)) \cdot d\hat{f}^\beta(\xi) - \sum_{m\in M} \int_t^{t+\Delta} a_m dg_m^\beta(\xi) \right| \leqslant \epsilon\Delta + \beta O(1).$$

Taking the limit on $\beta \to 0$ and using (5), we obtain

$$\left| \int_t^{t+\Delta} \nabla H(x(\xi)) \cdot d\hat{f}(\xi) - \sum_{m\in M} a_m \pi_m \Delta \right| = \left| \int_t^{t+\Delta} \nabla H(x(\xi)) \cdot d\hat{f}(\xi) - \left[\max_{v\in V} \nabla H(x(t)) \cdot v\right]\Delta \right| \leqslant \epsilon\Delta.$$

Because Δ can be chosen arbitrarily small (for a given fixed ϵ), $v(t) = \hat{f}'(t)$, and $\nabla H(x(\xi))$ is continuous at $\xi = t$, we have

$$\left| \nabla H(x(t)) \cdot v(t) - \max_{v\in V} \nabla H(x(t)) \cdot v \right| \leqslant \epsilon.$$

Finally, because ϵ can be chosen arbitrarily small, $\nabla H(x(t)) \cdot v(t) = \max_{v\in V} \nabla H(x(t)) \cdot v$, which completes the proof of (19) and with it the proof of property (i).

The compactness property (iii) of the family of FSPs is analogous to the compactness properties described in Lemmas 5.1–5.3 in Stolyar (1995). It follows directly from the construction of an FSP as a limit. Namely, for each FSP $z^{(j)}$, consider a defining sequence $z^{(j),\beta}$, $\beta \to 0$, such that $z^{(j),\beta} \to z^{(j)}$ u.o.c. For each j, consider $z^{(j),\beta_j}$, which is within distance $\epsilon_j > 0$ from $z^{(j)}$ in the uniform metric in the interval $[0, j]$, where $\epsilon_j \downarrow 0$. Then, it is easy to see that the sequence $z^{(j),\beta_j}$ defines FSP z. The second claim of (iii) follows from the first one. Indeed, from any sequence of FSPs $z^{(j)}$ such that $x^{(j)} \to x$ u.o.c., we can always choose a subsequence along which $z^{(j)} \to z = (x, \hat{f}, g, \hat{g})$ u.o.c., implying that this z is an FSP. □

11. Proofs of Theorems 1–3

Proof of Theorem 3. It suffices to prove a more general fact that any weak limit of the sequence of processes z^β is a process z concentrated on FSPs with probability 1. (In this case all processes are in the Skorohod space $D_{R^L}[0,\infty)$, where L is the total number of scalar component functions of z^β.)

Each sequence $\{g_m^\beta\}$ satisfies the law of large numbers, namely,

$$g_m^\beta(t_2) - g_m^\beta(t_1) \to \pi_m(t_2 - t_1) \quad \text{in probability } \forall\, 0 \leqslant t_1 \leqslant t_2, \tag{37}$$

and the sequence $\{z^\beta\}$ is "asymptotically Lipschitz", namely,

$$P\{\|z^\beta(t_2)-z^\beta(t_1)\| \leqslant C^*(t_2-t_1)+\epsilon^*\} \to 1 \qquad \forall\, 0 \leqslant t_1 \leqslant t_2,\ \epsilon^* > 0,$$

for some fixed $C^* > 0$.

Using the above two facts, the proof is completely analogous to the proof of Theorem 7.1 in Stolyar (1995).

We note that, alternatively, the proof can also be obtained more directly using Skorohod representation (see Ethier and Kurtz 1986). Namely, asymptotic Lipschitz property implies that the sequence $\{z^\beta\}$ is relatively compact. Then, any converging subsequence can be constructed on a probability space such that the convergence is with probability 1. Then, because property (37) holds, property (13) holds with probability 1. This, by the definition of an FSP, implies that z^β converges to an FSP with probability 1. □

Proof of Theorem 2. First, let us choose $a > 0$ to be fixed and large enough so that $a > \bar{\mu}$, $V \subseteq [0, a]^N$, and $A \subseteq [0, a]^N$. Then, without loss of generality, it will suffice to prove the theorem for the set A rechosen to be $A = [0, a]^N$.

For a given $\epsilon > 0$ and $\delta > 0$, by Theorem 4, we can choose $T > 0$ (depending on A and ϵ), so that for any FSP with $x(0) \in A$, we have

$$\sup_{t\in[T,\,T+\delta]} \|x(T)-u^*\| \leqslant \epsilon/2.$$

Using this, Theorem 3, and the continuous mapping theorem (see Billingsley 1968), it is easy to obtain

$$\lim_{\beta\downarrow 0} \sup_{x^\beta(0)\in A} P\left\{\sup_{t\in[T,\,T+\delta]} \|x^\beta(t)-u^*\| > \epsilon\right\} = 0. \tag{38}$$

Note, however, that (by (30)) $x_n^\beta(t)$ remains bounded by the maximum of $\bar{\mu}$ and $x_n^\beta(0)$ for all t. This means that $x^\beta(t) \in A$ for all t and all β. Applying (38) to the processes restarted at times $\tau \geqslant 0$, we can rewrite (38) in a stronger form:

$$\lim_{\beta\downarrow 0} \sup_{x^\beta(0)\in A} \sup_{\tau\geqslant 0} P\left\{\sup_{t\in[\tau+T,\,\tau+T+\delta]} \|x^\beta(t)-u^*\| > \epsilon\right\} = 0,$$

which implies the desired result. □

Proof of Theorem 1. Both $U^\beta(l_1, l_2)$ and $X^\beta(l_2)$ are differently computed averages of the values of $D^\beta(l)$: The former one is the average over $l_1 \leqslant l \leqslant l_2$ with equal weights $1/(l_2-l_1+1)$, and the latter one is (roughly) the average of $D^\beta(l_2-j)$ with "exponential" weights $\beta(1-\beta)^j$. We know from Theorem 2 that $EX^\beta(l)$ is close to u^* for large l, and therefore, when l_1 is large,

$$\frac{1}{l_2-l_1+1} \sum_{l_1\leqslant l\leqslant l_2} EX^\beta(l) \tag{39}$$

is close to u^*, uniformly on $l_2 \geqslant l_1$. In this proof, we show that when both l_1 and $l_2 - l_1$ are sufficiently large, the expected value $EU^\beta(l_1, l_2)$ is close to (39).

Let us choose $a > 0$ the same way as in the proof of Theorem 2. So, without loss of generality, we can assume that $A = [0, a]^N$. Then, again as observed in the proof of Theorem 2, for all $l \geqslant 0$ and all β, $X^\beta(l) \in A$, and consequently $\|X^\beta(l)\| \leqslant b \doteq \sqrt{N}a$.

Let us fix (small) $\epsilon > 0$. Theorem 2 along with the fact that the values of $X^\beta(l)$ are uniformly bounded implies that we can choose $T > 0$ such that for all sufficiently small β, uniformly on $l \geqslant T/\beta$, we have

$$\|EX^\beta(l)-u^*\| < \epsilon. \tag{40}$$

Without loss of generality, we can always choose T to be large enough so that e^{-T} is arbitrarily small. We note that

$$\sum_{j\geqslant T/\beta} \beta(1-\beta)^j = e^{-T} + O(\beta),$$

where (here and below) $O(\beta)$ denotes a term with absolute value bounded above by $C\beta$ as $\beta \to 0$, with $C > 0$ being a fixed constant.

In what follows, l_1, l_2 are functions of β, such that $l_1 \geqslant T/\beta$, $l_2 - l_1 \geqslant T'/\beta$, and $T' > T$. Also, we will use a simplified notation $d(p) \doteq ED^\beta(p)$. Then, we can write

$$Q \doteq \frac{1}{l_2-l_1+1} \sum_{l_1\leqslant l\leqslant l_2} EX^\beta(l) \tag{41}$$

$$= \frac{1}{T'/\beta} \sum_{l=l_1}^{l_2} \sum_{j=0}^{l-1} \beta(1-\beta)^j d(l-j) + O(\beta) \tag{42}$$

$$= Q_1 + Q_2 + Q_3 + O(\beta),$$

where Q_1, Q_2, and Q_3 break down the summation in (42), into the summations over

$$\sum_{l=l_1}^{l_2} \sum_{0\leqslant j<l-l_1}, \qquad \sum_{l_1\leqslant l<l_1+T/\beta} \sum_{l-l_1\leqslant j\leqslant l-1}, \qquad \sum_{l_1+T/\beta\leqslant l\leqslant l_2} \sum_{l-l_1\leqslant j\leqslant l-1},$$

respectively. Because we have $\|d(p)\| \leqslant b$, we obtain the following estimates:

$$\|Q_2\| \leqslant bT/T' + O(\beta),$$
$$\|Q_3\| \leqslant be^{-T}(T'-T)/T' + O(\beta).$$

For Q_1, by changing summation indexes from (l, j) to $(p = l-j, j)$, we can write

$$Q_1 = \frac{1}{T'/\beta} \sum_{p=l_1}^{l_2} \sum_{j=0}^{l_2-p} \beta(1-\beta)^j d(p)$$

$$= \frac{1}{T'/\beta} \sum_{p=l_1}^{l_2} \sum_{j=0}^{\infty} \beta(1-\beta)^j d(p) - Q_4 - Q_5$$

$$= EU^\beta(l_1, l_2) + O(\beta) - Q_4 - Q_5,$$

where

$$Q_4 = \frac{1}{T'/\beta} \sum_{l_1 \leqslant p < l_2 - T/\beta} \sum_{j > l_2 - p} \beta(1-\beta)^j d(p),$$

$$Q_5 = \frac{1}{T'/\beta} \sum_{l_2 - T/\beta \leqslant p \leqslant l_2} \sum_{j > l_2 - p} \beta(1-\beta)^j d(p),$$

$$\|Q_4\| \leqslant be^{-T}(T'-T)/T' + O(\beta),$$

$$\|Q_5\| \leqslant bT/T' + O(\beta).$$

We see that

$$\begin{aligned}\|Q - EU^\beta(l_1, l_2)\| &\leqslant \|Q_2\| + \|Q_3\| + \|Q_4\| + \|Q_5\| + O(\beta)\\ &\leqslant 2be^{-T}(T'-T)/T' + 2bT/T' + O(\beta).\end{aligned}$$

Thus, if we fix T large enough so that e^{-T} is small, and then choose T^* such that T/T^* is small, then we have $\|Q - EU^\beta(l_1, l_2)\| \leqslant \epsilon$ for all sufficiently small β, uniformly on $T' > T^*$. Because we also know that $\|Q - u^*\| \leqslant \epsilon$, the result follows. □

12. Concluding Remarks

In this section we discuss our model assumptions and techniques, and their possible generalizations and extensions. To facilitate the discussion, let us first observe that the analysis in this paper can be roughly divided into the following logical steps.

Step 1. Proving the process level convergence (Theorem 3), namely, the fact that any sequence of scaled processes $\{x^\beta\}$ has subsequences converging to a process with sample paths being FSPs.

Step 2. Establishing basic properties of FSPs (Lemmas 1 and 2), most importantly the differential inclusion (20) but also properties (ii) and (iii) of Lemma 2.

Step 3. Proving the FSP uniform attraction property (Theorem 4), based on the FSP properties established in Step 2. This is the key step.

Step 4. Establishing gradient algorithm asymptotic optimality (Theorem 1, Corollary 1, and Theorem 2), using the results of Steps 1 and 3. These results are our primary goal.

We now proceed with the discussion.

Smooth Nonstrictly Concave Utility Functions. The assumption of strict concavity of utility function H is not essential for the results of this paper. If H is just concave, the set of optimal solutions to problem (6)–(7) is a convex compact set $\bar{u}^*$. All our results and proofs hold virtually verbatim, with u^* replaced by $\bar{u}^*$ and the distances $\|\cdot - u^*\|$ to u^* being replaced by the distances $\rho(\cdot, \bar{u}^*)$ to the set $\bar{u}^*$, as for example in (8) and (10). The convergences (9) and (21) are generalized as $\lim \rho(ED^\beta(1), \bar{u}^*) = 0$ and $\rho(x(t), \bar{u}^*) \to 0$, respectively. We make the strict concavity assumption to simplify the exposition.

More General Type (II) Utility Functions. All our results hold as is for more general type (II) functions, where some of the components H_n are allowed to be (more "benign") finite and continuously differentiable everywhere on R_+, including 0. Extensions of the analysis to this case are transparent.

Generality of the FSP Uniform Attraction Proof. Our Step 3, the proof of the FSP uniform attraction property, is quite general. (In particular, it is strictly more general than the corresponding attraction results in Kushner and Whiting 2002, Agrawal and Subramanian 2002.) For a type (I) utility function (the case considered in Kushner and Whiting 2002 and Agrawal and Subramanian 2002) our proof uses only the differential inclusion (20) and the assumption that rate region V is a convex compact set. (We do not need any further assumptions on V, such as a smoothness or "cooperativeness" Kushner and Whiting 2002 condition.) In addition, we prove the uniform attraction for a type (II) utility function as well. This latter proof, however, uses FSP property Lemma 2(ii), which is in turn derived (in Step 2) using in part the "discrete structure" of our model, which we address next.

Discrete Structure of the Model. Our model assumes that the sets of switch states M and scheduling decisions $K(m)$ in each state m are finite. For such a "discrete structure" of the model, which is very common in applications, the rate region V is in fact a polyhedron (because the set of possible values of ϕ is a polyhedron and V is its linear image). Such nonsmooth rate regions V are not allowed by the results of Kushner and Whiting (2002) and Agrawal and Subramanian (2002), because they cause FSPs to be generally nonsmooth.

Discussion of the Results' Extensions for Different Models. The models in Kushner and Whiting (2002) and Agrawal and Subramanian (2002), although very close, are not within the framework of our model. (Neither is our model within the framework of those references, as discused above.) However, our asymptotic optimality results (Step 4) allow extensions to other models. As an example, consider the model of Agrawal and Subramanian (2002), where state process m is general stationary ergodic and service rate sets corresponding to each m are uniformly bounded compact sets (plus some other conditions), and the utility function is of type (I). The process level convergence (Step 1) and (20) for the FSPs (Step 2) are proved in Agrawal and Subramanian (2002). Then, as discussed above, our Theorem 4 (Step 3) applies. Examination of our proofs (in Step 4) of Theorem 1, Corollary 1, and Theorem 2 shows that neither the discrete structure nor Markov assumption (only stationarity!) are used there. Consequently, Theorem 1, Corollary 1, and Theorem 2 hold for the model of Agrawal and Subramanian (2002). These asymptotic optimality properties do not follow from the results in Agrawal and Subramanian (2002), because their proof requires that FSPs are attracted to the optimal point u^* from any initial state.

Discussion of Nonsmooth Utility Functions. Although our Step 3 can be generalized to the case of a general concave utility function H, the gradient algorithm asymptotic optimality results (of Step 4) do not necessarily hold for such utility functions. To be more specific, suppose that $H(u)$ is a concave function, finite for all $u \in R_+^N$, with finite directional derivatives

$$\frac{\partial}{\partial w} H(u) \doteq \frac{d_+}{d\xi} H(u + w\xi)$$

for all $u \in R_+^N$ and $w \in R^N$, such that $u_n = 0$ implies $w_n \geqslant 0$. (Such directional derivatives are well defined due to the concavity of H.) Consider the following differential inclusion, which generalizes (20):

$$x'(t) = v(t) - x(t),$$
$$v(t) \in \arg\max_{v \in V} \frac{\partial}{\partial(v - x(t))} H(x(t)). \tag{43}$$

Then, it is easy to check that our Theorem 4 and its proof still hold for the family of solutions of (43). Now, consider the corresponding generalization of the gradient algorithm, namely the algorithm choosing a decision

$$k \in \arg\max_{k \in K(m)} \frac{\partial}{\partial(\mu^m(k) - X(t))} H(X(t)).$$

It is not hard to construct examples (with, say, $H(u) = \min_n u_n$) such that FSPs for this algorithm do not satisfy inclusion (43). (Step 2 "does not work.") As a consequence, the asymptotic optimality of the gradient algorithm (Theorem 1, Corollary 1, and Theorem 2) does not necessarily hold for general concave utility functions.

References

Agrawal, R., V. Subramanian. 2002. Optimality of certain channel aware scheduling policies. *Proc. 40th Annual Allerton Conf. Comm., Control, Comput.*, Monticello, IL, 1532–1541.

Bender, P., P. Black, M. Grob, R. Padovani, N. Sindhushayana, A. Viterbi. 2000. CDMA/HDR: A bandwidth efficient high speed wireless data service for nomadic users. *IEEE Comm.* **38**(July) 70–77.

Billingsley. P. 1968. *Convergence of Probability Measures*. John Wiley and Sons, New York.

Ethier, S. N., T. G. Kurtz. 1986. *Markov Process: Characterization and Convergence*. John Wiley and Sons, New York.

Gans, N., G. van Ryzin. 1998. Optimal dynamic scheduling of a general class of parallel-processing queueing systems. *Adv. Appl. Probab.* **30** 1130–1156.

Harrison. J. M. 1998. Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete review policies. *Ann. Appl. Probab.* **8** 822–848.

Jalali, A., R. Padovani, R. Pankaj. 2000. Data throughput of CDMA-HDR, a high efficiency—high data rate personal communication wireless system. *Proc. IEEE Semiannual Vehicular Tech. Conf.*, VTC2000-Spring. Tokyo, Japan, 1854–1858.

Kushner, H., P. Whiting. 2002. Asymptotic properties of proportional fair sharing algorithms. *Proc. 40th Annual Allerton Conf. Comm., Control, Comput.*, Monticello, IL, 1051–1059.

Liu, X., E. K. P. Chong, N. B. Shroff. 2001. Opportunistic transmission scheduling with resource-sharing constraints in wireless networks. *IEEE J. Selected Areas in Comm.* **19**(10) 2053–2064.

Liu, X., E. K. P. Chong, N. B. Shroff. 2003. A framework for opportunistic scheduling in wireless networks. *Comput. Networks* **41** 451–474.

McKeown, N., V. Anantharam, J. Walrand. 1996. Achieving 100% throughput in an input-queued switch. *Proc. INFOCOM'96*, San Francisco, CA, 296–302.

Mekkittikul, A., N. McKeown. 1996. A starvation free algorithm for achieving 100% throughput in an input-queued switch. *Proc. ICCCN'96*, Rockville, MD, 226–231.

Stolyar, A. L. 1995. On the stability of multiclass queueing networks: A relaxed sufficient condition via limiting fluid processes. *Markov Processes Related Fields* **1**(4) 491–512.

Stolyar, A. L. 2004. MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.* **14**(1) 1–53.

Tassiulas, L., A. Ephremides. 1992. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Auto. Control* **37** 1936–1948.

Tassiulas, L., A. Ephremides. 1993. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. Inform. Theory* **39** 466–478.

Tse, D. 2004. Multiuser diversity and proportionally fair scheduling. In preparation.

Viswanath, P., D. Tse, R. Laroia. 2002. Opportunistic beamforming using dumb antennas. *IEEE Trans. Inform. Theory* **48**(6) 1277–1294.

Viswanathan, H., S. Venkatesan, H. C. Huang. 2003. Downlink capacity evaluation of cellular networks with known interference cancellation. *IEEE J. Selected Areas in Comm.* **21**(5) 802–811.

Williams, R. J. 2000. On dynamic scheduling of a parallel server system with complete resource pooling. *Fields Institute Communications*, Vol. 28. American Mathematical Society, Providence, RI, 49–71.