A Fluid Analysis of Utility-based Wireless Scheduling Policies

Peijuan Liu, Randall Berry, Michael L. Honig

***Abstract*— We consider packet scheduling for the downlink in a wireless network, where each packet's service preferences are captured by a utility function that depends on the packet's delay. The goal is to schedule packet transmissions to maximize the total utility. We examine a simple gradient-based scheduling algorithm, the $\dot{U}R$-rule, which is a type of generalized $c\mu$-rule ($Gc\mu$) that takes into account both a user's channel condition and derived utility. We study the performance of this scheduling rule for a draining problem. We formulate a "large system" fluid model for this draining problem where the number of packets increases while the packet-size decreases to zero, and give a complete characterization of the behavior of the $\dot{U}R$ scheduling rule in this limiting regime. We then give an optimal control formulation for finding the optimal scheduling policy for the fluid draining model. Using Pontryagin's minimum principle, we show that, when the user rates are chosen from a TDM-type of capacity region, the $\dot{U}R$ rule is in fact optimal in many cases. Finally, we consider non-TDM capacity regions and show that here the $\dot{U}R$ rule is optimal only in special cases.**

I. Introduction

Efficient scheduling algorithms are a key component for providing high speed wireless data services. There has been much interest in "channel-aware" scheduling algorithms that exploit variations in channel quality across the user population to improve performance (e.g., [1]–[7]). An important consideration for such approaches is balancing the over-all system performance with each user's quality of service (QoS) requirements. For example, in a time division multiplexing (TDM) system that transmits to one user at a time, the overall throughput is maximized by always transmitting to the user with the best channel. This can result in poor performance for users with poor channel quality, especially when channel conditions vary slowly with time. To address this, various "fair" scheduling approaches have been considered, such as the *proportional fair* algorithm for the CDMA 1xEV-DO system [8].

In this paper, we consider a utility-based framework, where each packet's value as a function of delay is indicated by a utility function, which can vary across packets. The scheduling policy attempts to maximize the total system utility; in this way, the utility functions can be used to balance fairness and efficiency. We consider a simple gradient-based scheduling policy, called the $\dot{U}R$ scheduling rule [6], [7]. Here $\dot{U}$ represents a packet's marginal utility, and R is the achievable rate. This policy makes myopic decisions based only on the instantaneous parameter values, and so requires no knowledge of the fading or traffic statistics.

We analyze the performance of the $\dot{U}R$ policy for a draining model, where an initial set of packets must be sent and no new arrivals occur. For a limiting fluid version of this model, we first characterize the performance of the $\dot{U}R$ policy. We then formulate a continuous-time optimal control problem for finding the scheduling policy that maximizes the total utility. We show that in certain cases the $\dot{U}R$ scheduler is optimal; the optimality depends in part on the underlying physical layer capacity region. For a TDM type of capacity region, the $\dot{U}R$ rule is optimal for a broad class of utility functions; for a general capacity region, the $\dot{U}R$ rule is optimal only in some special cases. This generalizes our earlier work in [7], which showed the optimality for a TDM region in the special case of quadratic utility functions and uniform initial delays.

The $\dot{U}R$ policy is equivalent to the *generalized* $c\mu$ ($Gc\mu$) rule introduced in [10] for a single-server multi-class queueing system with convex delay costs.[1] In [10] the $Gc\mu$ rule is shown to be optimal in the heavy traffic regime. In [11], this rule is also shown to be heavy traffic optimal for a system with multiple flexible servers under "complete resource sharing". Here we do not consider the heavy traffic regime, but instead analyze the performance and optimality of this rule for the fluid model previously discussed. The optimality of the $Gc\mu$ rule for a different fluid "rush hour" model has been studied in [12].

We allow the utility to be an arbitrary concave decreasing function of delay. In the special case of quadratic utilities, the $\dot{U}R$ rule is equivalent to the "MaxWeight" policies studied in [1], [2], [13]. These policies are stabilizing policies in a variety of settings [1], [2]. Several other fair scheduling approaches, such as the proportional fair rule, can be viewed in terms gradient-based scheduling algorithms with utilities that depend on each user's average throughput [4], [5].

II. System Model

We consider a basic model for downlink scheduling from a single transmitter, e.g. a base station or access point. There are K classes of packets, where each class corresponds to packets intended for a specific user with the same utility. For simplicity, we assume that all packets have fixed length of L bits. We study a draining problem, where a group of packets are present at time $t = 0$ and no new arrivals occur. Each packet has a randomly chosen initial delay that reflects the length of time the packet was waiting before $t = 0$. Let

This work was supported by the Motorola-Northwestern Center for Telecommunications and by NSF under grants CCR-9903055 and CAREER award CCR-0238382. This work was performed while P. Liu was with the Dept. of ECE, Northwestern University.

P. Liu is with Rosetta Wireless Corp., Oakbrook Terrace, IL 60181 USA, Peijuan.Liu@RosettaWireless.com. R. Berry and M. Honig are with the Dept. of ECE, Northwestern University, Evanston, IL 60208 USA, {rberry,mh}@ece.northwestern.edu.

[1] A utility U that is a function of delay is equivalent to a delay cost of $-U$.

N_k denote the number of class k packets present at time $t = 0$. For $n = 1, \ldots, N_k$, let $W_{n,k}(0)$ denote the delay of the nth class k packet at $t = 0$.

When transmitting each packet, a transmission rate $R_{n,k}$, must be specified depending on the user's channel quality and the underlying physical layer implementation. To begin, we consider a TDM system and assume that the channel to each user is fixed over the time-period of interest.[2] In this setting, whenever a class k packet is scheduled, it will use a fixed transmission rate R_k. If the nth packet of class k is transmitted at time t, the total delay incurred is

$$D_{n,k} = W_{n,k}(0) + t + L/R_k,$$

where L/R_k is the transmission time of the packet. The utility received by sending this packet is given by $U_k(D_{n,k})$, where for all k, $U_k(\cdot)$ is a decreasing concave function. A scheduling policy is any rule for determining which packet to transmit whenever the server becomes idle, given the current delay of every waiting packet, and the utility function and transmission rate for each class. Here, we consider the following scheduling policy from [6], [7]:

$\dot{\mathbf{U}}$R scheduling rule: *Schedule the head-of-line packet from class k^* such that*

$$k^* = \arg\max_k |\dot{U}_k(\hat{W}_k(t))| R_k, \tag{1}$$

where $\hat{W}_k(t)$ denotes the head-of-line delay of class k packets at time t. Ties may be broken arbitrarily.

This scheduling rule can be viewed as taking gradient steps to maximize the average utility per packet, i.e.,

$$U_{avg} = \frac{1}{N} \sum_{k=1}^{K} \sum_{n=1}^{N_k} U_k(D_{n,k}), \tag{2}$$

where $N = \sum_k N_k$.

A. Fluid limit

Following [7], we consider a fluid limit for the above system, where the number of packets $N \to \infty$ and the packet length, $L \to 0$, while keeping a fixed load of NL bits. For convenience, we normalize $NL = 1$.[3] We also keep a fixed proportion of packets in each class, given by $p_k = \frac{N_k}{N}$. With this scaling, the time required to drain the system with any work conserving scheduling rule is given by

$$T_f = \sum_{k=1}^{K} \frac{N_k L}{R_k} = \sum_{k=1}^{K} \frac{p_k}{R_k}. \tag{3}$$

For each class k, we assume that the initial delays $\{W_{n,k}(0)\}_n$ is a sequence of *i.i.d.* random variables chosen on the compact set $[D_k^l, D_k^u]$, where $D_i^l > 0$ and $D_i^u < \infty$ are lower and upper bounds on the initial delay, respectively. Let $G_k(w) = \Pr(W_{n,k}(0) \le w)$ be the cumulative distribution function (*c.d.f.*) for the initial delays of class k packets; we assume this is strictly increasing on $[D_k^l, D_k^u]$.

Let $\mathcal{N}_k^N(t)$ denote the number of class k packets remaining at time t in a system with N initial packets (for a given scheduling policy), and let

$$f_k^N(t) = \frac{\mathcal{N}_k^N(t)}{N_k}$$

be the fraction of class i packets remaining at time t. Following similar arguments as in [7], as $N \to \infty$ with the above scaling, for each class k, $f_k^N(t) \to f_k(t)$ almost surely. Let[4] $\alpha_k(t) = -p_i \dot{f}_i(t)/R_i$, so that,

$$f_k(t) = 1 - \int_0^t \frac{\alpha_k(\tau) R_k}{p_k} d\tau. \tag{4}$$

The quantity $\alpha_k(t)$ can be interpreted as the fraction of resources devoted to class k packets at time t. It can be shown that $\alpha_k(t)$ can take on any value in $[0,1]$ and must satisfy $\sum_{k=1}^K \alpha_k(t) \le 1$ for each time t. If $\alpha_k(t) = 1$, then only class k packets are served. For a non-idling policy, $\sum_{k=1}^K \alpha_k = 1$, for all $t \in [0, T_f)$. At each time t, the scheduling algorithm specifies $\alpha_k(t)$. Equivalently, for a TDM system, we can view the scheduler as selecting rates $r_k = \alpha_k(t) R_k$ from the capacity region given by $\mathcal{C}_{TS} = \left\{ \mathbf{r} : \sum_{k=1}^K \frac{r_k}{R_k} = 1 \right\}$. This interpretation allows us to generalize this model to other capacity regions $\mathcal{C}$, in which case the scheduling policy specifies a rate vector $\mathbf{r}(t) \in \mathcal{C}$ at each time t.

Next, we turn to the packet delays in the limiting system. Let $D_k^N(t)$ be the head-of-line delay of class k packets at time t in a system with N initial packets. Assuming the scheduling policy serves packets in a longest-delay-first order. Then in the fluid system, $D_k^N(t)$ converges (almost surely) to the deterministic function $D_k(t)$, where

$$D_k(t) = H_k(f_k(t)) + t. \tag{5}$$

Here $H_k(f) = G_k^{-1}(f)$ is the inverse of the class k initial delay c.d.f., which denotes the maximum initial delay of the remaining packets in the fluid system. Finally, the average utility per packet in the fluid model is given by

$$U_{avg} = \sum_{i=1}^{K} \int_0^{T_f} \alpha_i(t) R_i U_i [D_i(t)] \, dt.$$

III. Limiting Behavior of $\dot{U}R$ Scheduler

Next we characterize the limiting behavior of the $\dot{U}R$ scheduling rule (i.e., the behavior of $\alpha_k(t)$ for each class k.) Let $S(t) = \{k : f_k(t) > 0\}$ be the set of non-empty classes at time t. Define $M_k(t) = |\dot{U}_k(D_k(t))| R_k$ to be the decision metric used by the scheduler for each class $k \in S(t)$. Among classes $k \in S(t)$, the $\dot{U}R$ scheduler transmits the head-of-line packet of the class with the maximum value of $M_k(t)$. Therefore, the $\dot{U}R$ rule satisfies the following two properties:

[2] The time-invariant channel model will be more appropriate for low-tier mobility applications such as fixed wireless access.

[3] There is no loss in generality in assuming that the product NL is normalized to 1. A system with $NL = c$ with $c \neq 1$ and rates $\{R_k\}|_{k=1}^K$, can be easily shown to behave equivalently to a system with $NL = 1$ and rates $\{R_k/c\}|_{k=1}^K$.

[4] This derivative can be shown to exist except possibly on a set of measure zero. At the values of t where $\tau_i(t)$ is not differentiable, we interpret this as the right derivative so that $r_k(t)$ is right-continuous [14].

Property 1: If $k \notin S(t)$, then $\alpha_k(t) = 0$.
Property 2: For $k \in S(t)$, $\alpha_k(t) = 0$ if there exists $j \neq k$ such that $M_k(t) < M_j(t)$.

For $k \notin S(t)$, i.e., a class which is drained at time t, we define $D_k(t) = D_k^l + t$, which is a natural extension of (5). That is, the delay for class k formally continues to increase after all class k packets have been drained. This does not affect any scheduling decisions or performance, but will be useful in optimal control formulation in Section IV.

From Properties 1 and 2 , it follows that under the $\dot{U}R$ rule, if there exists a unique class k such that $M_k(t) > M_j(t)$ for all j, then $\alpha_k(t) = 1$. This determines the behavior of the scheduling rule except at those times t when multiple classes simultaneously have the maximum value of $M_k(t)$. In these cases, the fluid scheduler splits its resources among these classes. Let $Q(t)$ be the set of non-empty classes that have the maximum value of M_k, i.e.,

$$Q(t) = \{k \in S(t) : M_k(t) \geq M_j(t) \text{ for all } j \in S(t)\}.$$

The following theorem quantifies how resources are shared among these packets when $|Q(t)| \geq 2$.

Theorem 1: For any $t < T_f$ with $|Q(t)| \geq 2$, let $\{\alpha_k(t), k \in Q(t)\}$ be the solution to:

$$\dot{M}_k(t) = -\ddot{U}_k\left[D_k(t)\right]\left(-\frac{\alpha_k(t)R_k}{p_k}\dot{H}_k\left[f_k(t)\right]+1\right)R_k = K_0(t), \tag{6}$$

where $K_0(t)$ is chosen to satisfy

$$\sum_{k \in Q(t)} \alpha_k(t) = 1. \tag{7}$$

If $0 < \alpha_k(t) < 1$ for all $k \in Q(t)$, then these are the resource allocations under the $\dot{U}R$ rule.

The proof of this follows from that fact that $\alpha_k(t)$ is right-continuous. It can be shown that a unique solution to (6) and (7) always exists. Whether or not this solution satisfies $0 < \alpha_k < 1$ for all $k \in Q(t)$ depends on the choice of $U_k(D), H_k(f)$ and R_k. Given $H_k(f)$ and R_k, we define a set of utility functions $\{U_k(D) | k = 1, \ldots, K\}$, to be *regular* if a feasible solution to (6) and (7) exists for all t where $|Q(t)| \geq 2$. For example, with $K = 2$, $R_1 > R_2$, and uniform initial delays, $U_1(D) = U_2(D) = -D^\beta$, can be shown to be regular for $\beta > 1$. In what follows, we will assume that $\{U_k(D)\}$ is regular unless stated otherwise.

As an example, consider a system with $K = 2$ and $p_1 = p_2 = 1/2$. Assume that $U_1(D) = U_2(D) = U(D)$, and that the initial delays are uniformly distributed on $[0, 1]$, i.e., for $k = 1, 2$,

$$G_k(w) = \begin{cases} w & 0 \leq w \leq 1, \\ 1 & w \geq 1. \end{cases} \tag{8}$$

In this case, (5) becomes $D_k(t) = f_k(t) + t$, and therefore

$$\dot{D}_k(t) = -2\alpha_k(t)R_k + 1,$$

with $D_k(0) = 1$ for $k = 1, 2$. From Properties 1 and 2, the fluid $\dot{U}R$ scheduler sets

$$\alpha_1(t) = \begin{cases} 1 & \text{if } |\dot{U}(D_1(t))|R_1 > |\dot{U}(D_2(t))|R_2 \text{ or } f_2(t) = 0, \\ 0 & \text{if } |\dot{U}(D_1(t))|R_1 < |\dot{U}(D_2(t))|R_2 \text{ or } f_1(t) = 0 \end{cases}$$

and $\alpha_2(t) = 1 - \alpha_1(t)$. Also, from Theorem 1, we have that for any t such that $f_1(t), f_2(t) > 0$ and $\dot{U}(D_1(t))R_1 = \dot{U}(D_2(t))R_2$, the $\dot{U}R$ rule gives

$$\alpha_1(t) = \frac{\ddot{U}(D_1(t))R_1 - \ddot{U}(D_2(t))R_2 + \ddot{U}(D_2(t))R_2^2}{\ddot{U}(D_1(t))R_1^2 + \ddot{U}(D_2(t))R_2^2}, \tag{9}$$

and $\alpha_2(t) = 1 - \alpha_1(t)$.

Define $t_k^{in} = \inf\{t | \alpha_k(t) > 0\}$ and $t_k^{out} = \inf\{t | f_k(t) = 0\}$ to be the times the server starts to serve and drains all the class k packets, respectively.

Corollary 1: If $\{U_k(D)\}$ are regular, then $t_k^{out} = \inf\{t > t_k^{in} | \alpha_k(t) = 0\}$.

In other words, once the scheduler starts serving class k packets, it continues to serve this class until all class k packets are drained. This follows from Theorem 1, which implies that once class k joins the active set $Q(t)$, it remains in $Q(t)$ until time t_k^{out}. The initiation and termination times for every class k, $\{t_k^{in}\}_{k=1}^K$ and $\{t_k^{out}\}_{k=1}^K$, mark $2K$ events[5]. Let $t^1 \leq t^2 \leq \cdots \leq t^{2K}$ denote the ordered list of these times, so that for each $j = 1, \ldots, 2K$, $t^j = t_k^{in}$ or t_k^{out} for some k. At each time t, define the *upper envelope* $\bar{M}(t)$ of $\{M_k(t)\}_{k=1}^K$ to be the value of $M_k(t)$ for all users $k \in Q(t)$. Notice that t_k^{in} and t_k^{out} satisfy $\bar{M}(t_k^{in}) = |\dot{U}_k(D_k^u + t_k^{in})|R_k$ and $\bar{M}(t_k^{out}) = |\dot{U}_k(D_k^l + t_k^{out})|R_k$.

So far, we have characterized the $\dot{U}R$ rule given the decision metrics $\{M_k(t)\}_{k=1}^K$. Next, we determine how each $M_k(t)$ evolves with t. Recall that $Q(t)$ is the set of non-empty classes receiving service at time t. Let $\bar{Q}(t) = S(t) - Q(t)$ be the set of inactive classes, which still have packets remaining to be transmitted at time t. The decision metrics and the upper envelope can be computed via the following iterative procedure:

1.) Set $j = 1$, $t^1 = 0$.
2.) While $t^j < T_f$ do:
   a.) For all k, calculate $f_k(t^j)$ and $M_k(t^j)$ and update $S(t^j), Q(t^j)$ and $\bar{Q}(t^j)$;
   b.) Set $\alpha_i(t^j) = 0$ for $i \notin S(t^j)$;
   c.) If $Q(t) = \{i\}$, set $\alpha_i(t) = 1$ and $\alpha_k(t) = 0$ for all $k \notin Q(t)$ for $t \in (t^j, t^{j+1})$.
   else if $|Q(t)| \geq 2$, calculate $\alpha_k(t)$ for $k \in Q(t)$ and $t \in (t^j, t^{j+1})$ from Theorem 1;
   d.) Evaluate $\bar{M}(t)$ from (6) for $t \in (t^j, t^{j+1})$, and compute

$$t^{j+1} = \min\left[\inf\left(t : M_k(t) = \bar{M}(t), k \in \bar{Q}(t)\right),\right. \\ \left. \inf\left(t : f_i(t) = 0, \forall i \in Q(t)\right)\right];$$

   e.) Set $j = j + 1$ and goto 2.

The quantities in step (2.a) can be computed directly from their definitions. In step (2.d), the two terms in the minimum are the smallest $t_k^{in} > t^j$ and the smallest $t_k^{out} > t^j$. Given $\bar{M}(t)$, the system behavior is completely determined. Namely, the event times $\{t^j\}$ are the intersections of $\bar{M}(t)$

[5] It is possible that some of these events coincide. In that case, we can order them arbitrarily.

with $|\dot{U}_k(D_k^u+t)|R_k$ or $|\dot{U}_k(D_k^l+t)|R_k$, for $k=1,\cdots,K$. The evolution of the decision metrics and service allocations between event times is given by Theorem 1.

Consider again the previous 2-class example with uniform initial delay distributions. In this case, let $\hat{t}$ be the solution to $\dot{U}(1+(1-2R_1)\hat{t})R_1 = \dot{U}(1+\hat{t})R_2$. Then step (2.d) of the iteration implies that if $\hat{t} \geq \frac{1}{2R_1}$, then $t_1^{out} = t_2^{in} = \frac{1}{2R_1}$, and so the scheduler sends all class 1 packets before serving any class 2 packets. Otherwise, if $\hat{t} < \frac{1}{2R_1}$, then $t_1^{out} > t_2^{in} = \hat{t}$ and the scheduler serves both class during $[t_2^{in}, t_1^{out})$, with $\alpha_1(t)$ given by (9).

IV. Optimal Fluid Scheduling Policies

We next consider finding optimal scheduling policies for the fluid system. In particular, we give sufficient conditions for the $\dot{U}R$ policy to be optimal. First, we focus on TDM capacity regions, then we consider non-TDM regions.

A. TDM capacity regions

For simplicity, we consider a two-class system with a TDM capacity region and transmission rates $R_1 > R_2$. Each class $k=1,2$ has a decreasing, concave utility, $U_k(D)$. We again assume that the initial delay for class k packets is distributed on the interval $[D_k^l, D_k^u]$ according to the c.d.f. $G_k(w)$, with a well-defined inverse $H_k(x)$. Without loss of generality, assume that $|\dot{U}_1(D_1(0))|R_1 \geq |\dot{U}_2(D_2(0))|R_2$ so that the scheduler always begins by serving class 1. Characterizing a scheduling policy is equivalent to specifying the functions $\alpha_1(t)$ and $\alpha_2(t)$ for all $t \in [0, T_f)$. We want to choose these to maximize the total utility derived; this can be formulated as:

$$\min_{\alpha_1(t),\alpha_2(t)} \int_0^{T_f} \left[-\sum_{k=1}^{2} \alpha_k(t) R_k U_k \left[H_k(f_k(t))+t\right] \right] dt \quad (10)$$

$$\text{subject to: } \dot{f}_k(t) = -\frac{\alpha_k(t)}{p_k} R_k,\ k=1,2, \quad (11)$$

$$f_k(0)=1,\ k=1,2 \text{ and } f_1(T_f)=0, \quad (12)$$

$$\alpha_1(t)+\alpha_2(t)=1, \quad (13)$$

$$\alpha_k(t) \geq 0,\ k=1,2. \quad (14)$$

This can be viewed as a continuous-time optimal control problem [15] with a fixed terminal time T_f, where the state is $\mathbf{f}(t) = (f_1(t), f_2(t))$ and $\alpha(t) = (\alpha_1(t), \alpha_2(t))$ is the control variable. Here, (11) represents the system dynamics, and (12) gives initial and final boundary conditions for the state. The final state $(f_1(T_f), f_2(T_f))$ is restricted to be on the line $f_1(T_f)=0$. Any admissible control $\alpha(t)$ also results in $f_2(T_f)=0$. However, we do not explicitly state this boundary condition. If we are given $f_1(t)$, then we can compute $f_2(t)$ and in particular, $f_1(T_f)=0$ implies $f_2(T_f)=0$. Hence the latter constraint is not independent. Furthermore, we require $\alpha(t)$ to be right-continuous.

If all the packets in class k are emptied at time $\hat{t} < T_f$, then for all $t > \hat{t}$, we have that $\alpha_k(t)=0$ and $f_k(t)=0$. To see that this must hold in the preceding formulation, note that since $f_k(\hat{t})=0$ and $f_k(T_f)=0$, $\dot{f}_k(t)=0$ for $\hat{t} < t < T_f$. Also, from (11) and (14), $\alpha_k(t)=0$ for $\hat{t} < t < T_f$.

The solution to this problem can be characterized using the Pontryagin minimum principle [15]. We first define the Hamiltonian for this problem, which is given by

$$\mathscr{H}(\mathbf{f}(t),\alpha(t),\mathbf{q}(t)) = -\sum_{k=1}^{2} \alpha_k(t) R_k \left[U(D_k(t)) + \frac{q_k(t)}{p_k} \right],$$

where $\mathbf{q}(t) = (q_1(t), q_2(t))$ is the co-state or Lagrange multiplier, and $D_k(t) = H(f_k(t))+t$. Let $\alpha^*(t)$ be an optimal control and $\mathbf{D}^*(t)$ the corresponding optimal state trajectory. According to the Pontryagin minimum principle, there exists a $\mathbf{q}^*(t)$ that satisfies the co-state equations:

$$\dot{\mathbf{q}}^*(t) = -\nabla_{\mathbf{f}} \mathscr{H}(\mathbf{f}^*(t),\alpha^*(t),\mathbf{q}^*(t)), \quad (15)$$

such that for all admissible controls $\alpha(t)$,

$$\mathscr{H}(\mathbf{f}^*(t),\alpha^*(t),\mathbf{q}^*(t)) \leq \mathscr{H}(\mathbf{f}^*(t),\alpha(t),\mathbf{q}^*(t)). \quad (16)$$

For this problem, the co-state equations (15) are:

$$\dot{q}_k(t) = \alpha_k(t) R_k \dot{U}_k(D_k(t)) \dot{H}_k(f_k(t)),\ k=1,2.$$

Furthermore, the final state conditions dictate that $q_2(T_f)=0$ [16]. Let $A_k(t) = R_k \left[U_k(D_k(t)) + \frac{q_k(t)}{p_k} \right]$ for $k=1,2$. Then the Hamiltonian can be written as

$$\mathscr{H}(\mathbf{f}(t),\alpha(t),\mathbf{q}(t)) = -A_1(t)\alpha_1(t) - A_2(t)\alpha_2(t),$$

which is linear in $\alpha_k(t)$. To satisfy (16), it follows that

$$\alpha_1^*(t) = \begin{cases} 1, & \text{if } A_1(t) > A_2(t), \\ 0, & \text{if } A_1(t) < A_2(t), \end{cases} \quad (17)$$

and $\alpha_2^*(t) = 1 - \alpha_1^*(t)$. Let $\Delta(t) = A_1(t) - A_2(t)$, so that the sign of $\Delta(t)$ determines the optimal control at time t. If $\Delta(t)=0$, then the problem is *singular* at time t. This means that (16) alone does not specify the optimal control. A *singular interval* $[t_1, t_2]$ means that the problem is singular for all t in $[t_1, t_2]$, i.e., $\Delta(t)=0$ for all $t \in [t_1, t_2]$.

Lemma 1: During any singular interval, the optimal control must satisfy (9).

Proof: Notice that

$$\begin{aligned} \dot{A}_k(t) &= R_k \left[\dot{U}_k(D_k(t))\dot{D}_k(t) + \frac{\dot{q}_k(t)}{p_k} \right] \\ &= R_k \left\{ \dot{U}_k(D_k(t)) \left[-\dot{H}_k(f_k(t)) \frac{\alpha_k(t)}{p_k} R_k + 1 \right] \right. \\ &\qquad \left. + \frac{\alpha_k(t)}{p_k} R_k \dot{U}(D_k(t)) \dot{H}_k(f_k(t)) \right\} \\ &= R_k \dot{U}_k(D_k(t)), \end{aligned} \quad (18)$$

which does not depend on $\alpha_k(t)$. Furthermore, for all t in a singular interval, it must be that $\dot{\Delta}(t)=0$. Therefore, $\dot{A}_1(t) = \dot{A}_2(t)$, i.e., $R_1 \dot{U}(D_1(t)) = R_2 \dot{U}(D_2(t))$. This corresponds to the choice of $\alpha_1(t)$ in (9). ■

Notice that $\Delta(t)$ is continuous and differentiable since both $A_1(t)$ and $A_2(t)$ are continuous and differentiable. Lemma 1 implies that during any singular interval, the optimal scheduling policy behaves like the $\dot{U}R$ rule. Recall from Sect. III that the $\dot{U}R$ scheduler starts serving class 1 packets up to t_2^{in}, then serves both classes simultaneously for $t \in [t_2^{in}, t_1^{out})$ (where t_2^{in} may equal t_1^{out}), and finally serves the remaining class 2 packets until $t = T_f$. To show that the $\dot{U}R$ rule is optimal for all $t \in [0, T_f)$, we need to show that *i)* $\Delta(t)$ is unique; *ii)* $\Delta(t) > 0$ for $t \in [0, \hat{t}_2^{in})$, $\Delta(t) = 0$ for $t \in [\hat{t}_2^{in}, \min\{\hat{t}_1^{out}, \hat{t}_2^{out}\})$, and $\Delta(t) < 0$ for $t \in [\hat{t}_1^{out}, T_f)$ (if $\hat{t}_1^{out} < \hat{t}_2^{out}$) or $\Delta(t) > 0$ for $t \in [\hat{t}_2^{out}, T_f)$ (if $\hat{t}_2^{out} < \hat{t}_1^{out}$); and *iii)* $t_2^{in} = \hat{t}_2^{in}$, $t_1^{out} = \hat{t}_1^{out}$ and $t_2^{out} = \hat{t}_2^{out}$.

In the following, we assume that $U_1(D)$ and $U_2(D)$ are decreasing, strictly concave in D, and that they are regular (see Sect. III) for the given delay distributions and rates. We first show in Lemma 2 that for such utility functions, if $\Delta(t)$ is non-increasing on an interval where it is strictly positive, then it must be strictly decreasing on this interval. Next, in Lemma 3, we show that if $\Delta(t)$ is non-increasing, then the $\dot{U}R$ rule must be optimal. Finally, in Theorem 2, we give a condition on the utility functions under which the $\dot{U}R$ rule is optimal. The proofs are omitted due to space considerations and can be found in [17].

Lemma 2: Assume that $\Delta(t) > 0$ for all $t \in [a, b)$. If $\Delta(t)$ is non-increasing, i.e., $\dot{\Delta}(t) \le 0$ for all $t \in [a, b)$, then for regular utility functions, $\dot{\Delta}(t) < 0$ for all $t \in [a, b)$.

Lemma 3: For regular utility functions, if $\dot{\Delta}(t) \le 0$ for all $t \in [0, T_f]$, then the $\dot{U}R$ rule is optimal.

Theorem 2: Assume that the utility functions satisfy the following condition for all $t_0 > 0$: If $R_1 \dot{U}_1(D_1(t_0)) = R_2 \dot{U}_1(D_2(t_0))$, then for all $s > 0$,

$$\begin{aligned} &\text{i)} \quad R_1 \dot{U}_1 \left[H_1 \left(f_1(t_0) - \tfrac{R_1}{p_1} s \right) + t_0 + s \right] \\ &\qquad > R_2 \dot{U}_2 (D_2(t_0) + s); \\ \text{and} \quad &\text{ii)} \quad R_1 \dot{U}_1 (D_1(t_0) + s) \\ &\qquad < R_2 \dot{U}_2 \left[H_2 \left(f_2(t_0) - \tfrac{R_2}{p_2} s \right) + t_0 + s \right]. \end{aligned}$$

Then the $\dot{U}R$ rule is optimal.

Recall, $f_k(t_0) = G_k(D_k(t_0) - t_0)$ is the fraction of class k packets remaining at time t_0. The left-hand (right-hand) side of condition *i)* is the value of $M_1(t_0 + s)$ ($M_2(t_0 + s)$) if the scheduler serves only class 1 packets from time t_0 to $t_0 + s$. Condition *ii)* is the analogous relation if the scheduler serves only class 2 packets.

Corollary 2: With a uniform initial delay distribution, $U_1(D) = U_2(D) = U(D)$, and $R_1 > R_2$, then the $\dot{U}R$ rule is optimal in the following cases:

(1) $U(D) = -D^{\beta}$ with $\beta > 1$.
(2) $U(D) = 1 - e^{kD}$ where $k > 0$ is a constant.
(3) $U(D)$ is concave and $R_2 > 1$.

B. General capacity regions

We next consider the optimality of the $\dot{U}R$ rule for a more general 2-user capacity region $\mathscr{C}$ that is a compact, convex and coordinate convex[6] subset of $\mathbb{R}_+^2$. For an arbitrary capacity region, we define $\delta\mathscr{C}$ to be the set of Pareto dominate rates, i.e., $\mathbf{r} \in \delta\mathscr{C}$ if and only if $\mathbf{r} \in \mathscr{C}$ and there is no other $\mathbf{r}' \in \mathscr{C}$ such that $\mathbf{r}' \ge \mathbf{r}$. (All vector inequalities are component-wise.) We say that $\mathscr{C}$ has a strictly convex boundary if for any pair $\mathbf{r}, \mathbf{r}' \in \delta\mathscr{C}$, $\alpha\mathbf{r} + (1-\alpha)\mathbf{r}' \notin \delta\mathscr{C}$ for any $\alpha \in (0,1)$. One example of a capacity region $\mathscr{C}$ with a strictly convex boundary is the achievable rate region for a Gaussian broadcast channel. A rate vector $\mathbf{r} = (r_1, r_2)$ is defined to be in the *interior* of δC if $\mathbf{r} \in \delta C$ and $\mathbf{r} > \mathbf{0}$.

With such a capacity region, the $\dot{U}R$ scheduling policy is defined to be a policy that selects a rate vector $\mathbf{r}(t) = (r_1(t), r_2(t))$ at each time t such that

$$\mathbf{r}(t) = \arg\max_{\mathbf{r} \in \mathscr{C}} \sum_{k=1}^{2} |\dot{U}_k(D_k(t))| r_k. \tag{19}$$

Note that with the preceding assumptions, this optimization problem always has a solution $\mathbf{r} \in \delta\mathscr{C}$, and if $\mathscr{C}$ has a strictly convex boundary, then the solution is unique. For a given capacity region, $\mathscr{C}$, at each time t, the solution to (19) depends only on the ratio $V(t) \equiv \frac{\dot{U}_1(D_1(t))}{\dot{U}_2(D_2(t))}$. If $\mathscr{C}$ has a strictly convex boundary, then given any point $\hat{\mathbf{r}}$ in the interior of δC there is a unique value of the ratio $V(t)$ for which $\hat{\mathbf{r}}$ is the solution to (19).

The corresponding optimal control problem is:

$$\min_{r_1(t), r_2(t)} \quad \int_0^{T_f} \left[-\sum_{k=1}^{2} r_k(t) U_k \left[H_k(f_k(t)) + t \right] \right] dt \tag{20}$$

$$\begin{aligned} \text{subject to:} \quad & \dot{f}_k(t) = -\frac{r_k(t)}{p_k},\ k = 1, 2, \\ & f_k(0) = 1, \text{and } f_k(T_f) = 0,\ \forall k = 1, 2, \\ & \mathbf{r}(t) \in \mathscr{C} \end{aligned}$$

Here, the time to drain the system, T_f is in general not the same for all non-idling scheduling policies. Therefore, this is not a fixed-terminal time problem, rather, the terminal state is specified.

The Hamiltonian is now given by

$$\mathscr{H}(\mathbf{f}(t), \mathbf{r}(t), \mathbf{q}(t)) = -A_1(t) r_1(t) - A_2(t) r_2(t),$$

where $A_k(t) = U_k(D_k(t)) + \frac{q_k(t)}{p_k}$, and the co-state satisfies $\dot{q}_k(t) = r_k \dot{U}_k(D_k(t)) \dot{H}_k(f_k(t))$. Therefore, the optimal control, $\mathbf{r}^*(t)$, satisfies

$$\mathbf{r}^*(t) = \arg\max_{\mathbf{r} \in \mathscr{C}} (-A_1(t) r_1 - A_2(t) r_2), \tag{21}$$

for each time t. As for (19), this always has a solution in $\delta\mathscr{C}$, and if $\mathscr{C}$ has a strictly convex boundary, then (21) has a unique solution; i.e. there are no singular intervals.

Proposition 1: If the capacity region $\mathscr{C}$ has a strictly convex boundary, and at time $t = 0$, the solution to (19) is in the interior of $\delta\mathscr{C}$, then a necessary condition for the $\dot{U}R$ rule to be optimal is that there exists a constant K such that the $\dot{U}R$ rule gives $\dot{U}_1(D_1(t)) = K \dot{U}_2(D_2(t))$, for all $t \in [0, T_f]$.

At $t = 0$, the solution to (19) depends only on the utilities through the ratio $V(0) = \frac{\dot{U}_1(D_1(0))}{\dot{U}_2(D_2(0))}$. The assumption that the

[6] A set $\mathscr{X} \subset \mathbb{R}_+^n$ is said to be *coordinate convex* if $\mathbf{x} \in \mathscr{X}$ implies that $\mathbf{y} \in \mathscr{X}$ for all $\mathbf{y}$ such that $\mathbf{0} \le \mathbf{y} \le \mathbf{x}$.

solution to (19) is in the interior of $\delta\mathscr{C}$ and that $\delta\mathscr{C}$ is strictly convex implies that there is only one value of $V(0)$ that will give this solution. This proposition then says that the $\dot{U}R$ rule is optimal if and only if the $\dot{U}R$ scheduler gives $V(t) = K$ for all t. This implies that the $\dot{U}R$ rate allocation is fixed for all time t.

As an example, consider a system with uniform initial delays on $[0,1]$ for each class, and $U_k(D_k) = w_k U(D_k)$, $k = 1,2$, where $U(D)$ is the same for both classes and w_k is a class dependent weight. In this case

$$\frac{\dot{U}_1(D_1(0))}{\dot{U}_2(D_2(0))} = \frac{w_1 \dot{U}(1)}{w_2 \dot{U}(1)} = \frac{w_1}{w_2},$$

so that at time $t = 0$, (19) corresponds to maximizing the weighted sum rate $(w_1 r_1 + w_2 r_2)$ for the two classes. If the maximum weighted sum rate is achieved at an interior point of δC, then according to Prop. 1, for the $\dot{U}R$ rule to be optimal, it must give $D_1(t)$ and $D_2(t)$ that satisfy $\dot{U}(D_1(t)) = \frac{w_1}{w_2}\dot{U}(D_2(t))$ for all t. Since the utilities are the same, this implies that $D_1(t) = D_2(t)$ for all t, and so $\dot{f}_1(t) = \dot{f}_2(t)$, or equivalently $\frac{r_1}{p_1} = \frac{r_2}{p_2}$, where r_1 and r_2 are the rates that maximize the weighted sum rate for the two users. In other words, the line $r_1 = \frac{p_1}{p_2} r_2$ must intersect δC at the point that maximizes the weighted sum rate. For a given capacity region and utility weights, this implies that there is only one particular ratio of p_1 and p_2 for which the $\dot{U}R$ rule might be optimal, and this ratio must be "matched" to the utility weights.

Proposition 1 provides a necessary condition for the $\dot{U}R$ rule to be optimal. We have not shown sufficiency of these conditions in general, but we can show them in some special cases. For example, when the initial delays are uniform and utilities are affine. The difficulty is that the problem is not jointly convex in the control and state variables, which precludes appealing to standard sufficiency results.

V. Conclusions

We have presented an analysis of a simple utility-based scheduling rule for a downlink wireless data service, which schedules packets with the largest product of marginal utility and achievable rate. We studied the performance of this scheduler for a fluid draining model where the utility is a function of delay. Assigned to each packet are a randomly chosen initial delay and rate. In this setting, we are able to characterize how scheduling resources, or the total service time, is split among the remaining packets as time progresses. We can also explicitly compute performance measures such as average utility and delay.

We next considered finding the optimal scheduling policy for a two user systems, where the the objective is to maximize the total utility per packet. This was formulated as a continuous time optimal control problem. First we considered the problem with a TDM capacity region. Using Pontryagin's minimum principle, we characterized the optimal scheduling policy. Comparing the optimal and $\dot{U}R$ schedulers shows that both behave exactly the same whenever the service time is split between the two classes. For a general utility function, the way in which the optimal scheduler alternates service between the two classes may differ from the $\dot{U}R$ rule. We specified conditions on the utility function, which guarantee that this order is the same, so that the $\dot{U}R$ rule is optimal. These conditions include many utility functions of interest. Finally, we considered the optimal scheduling policy for a non-TDM capacity region with a strictly convex boundary. In that case, we showed that much stronger conditions are needed for the $\dot{U}R$ rule to be optimal.

In this work, we have not considered dynamically changing channels and retransmissions, which arise in mobile wireless data systems. The $\dot{U}R$ rule can, in principle, be modified to take these additional features into account. Associated modeling and performance issues are topics for further study.

References

[1] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queue with randomly varying connectivity", in *IEEE Transactions on Information Theory*, Vol. 39, pp. 466-478, March 1993.

[2] M. Andrews, et al., "Scheduling in a Queueing System with Asynchronously Varying Service Rates," To appear in *Probability in the Engineering and Informational Sciences*.

[3] S. Shakkottai and A. L. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule", in *Analytic Methods in Applied Probability*. Series 2, Volume 207, pp. 185-202.

[4] R. Agrawal and V. Subramanian, "Optimality of Certain Channel Aware Scheduling Policies," *Proc. of 2002 Allerton Conference on Communication, Control and Computing*, Oct. 2002.

[5] H. Kushner and P. Whiting, "Asymptotic Properties of Proportional-Fair Sharing Algorithms," *Proc. of 2002 Allerton Conference on Communication, Control and Computing*, Oct. 2002.

[6] P. Liu, R. Berry, and M. Honig, "Delay-Sensitive Packet Scheduling in Wireless Networks," *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC)*, New Orleans, LA, March 2003

[7] R. Berry, P. Liu and M. Honig, "Design and Analysis of Downlink Utility-Based Schedulers," *2002 Allerton Conference on Communication, Control and Computing*, Monticello, IL, October 2002.

[8] P. Bender, et al., "CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users", *IEEE Commun. Mag.*, pp. 70-77, July 2000.

[9] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic Beam-forming using Dumb Antennas," *IEEE Trans. on Information Theory*, vol. 48, June 2002.

[10] J. A. Van Mieghem, "Dynamic Scheduling with Convex Delay Costs: the Generalized $c\mu$ Rule," *Annals of Applied Probability*, 5(3), 1995.

[11] A. Mandelbaum and A. L. Stoylar, "$GC\mu$ Scheduling of Flexible Servers: Asymptotic Optimality in Heavy Traffic," *2002 Allerton Conf. on Communication, Control and Computing*, Oct. 2002.

[12] R. Haji and G. F. Newell, "Optimal Strategies for Priority Queues with Nonlinear Costs of Delay," *SIAM Journal on Applied Mathematics*, Volume 20, Issue 2, pages 224-240, March, 1971.

[13] A. L. Stolyar, "MaxWeight scheduling in a generalized switch: state space collapse and equivalent workload minimization in Heavy Traffic", submitted, 2001.

[14] F. Riesz and B. Nagy, *Functional Analysis*, Ungar, New York, 1955.

[15] M. Athans and P. Falb, *Optimal Control, An Introduction to the Theory and its Applications,* Mc Graw-Hill, New York, 1966.

[16] D. E. Kirk, *Optimal Control Theory an Introduction,* Prentice-Hall, NJ, 1970.

[17] P. Liu, R. Berry, and M. Honig, "A fluid analysis of a wireless utility-based scheduling policy," *under submission*.