

On the optimal trade-off between SRPT and opportunistic scheduling

Samuli Aalto
Aalto University
P.O.Box 13000, FIN-00076
Aalto, Finland
Samuli.Aalto@tkk.fi

Aleksi Penttinen
Aalto University
P.O.Box 13000, FIN-00076
Aalto, Finland
Aleksi.Penttinen@tkk.fi

Pasi Lassila
Aalto University
P.O.Box 13000, FIN-00076
Aalto, Finland
Pasi.Lassila@tkk.fi

Prajwal Osti
Aalto University
P.O.Box 13000, FIN-00076
Aalto, Finland
Prajwal.Osti@tkk.fi

ABSTRACT

We consider service systems where new jobs not only increase the load but also improve the service ability of such a system, cf. opportunistic scheduling gain in wireless systems. We study the optimal trade-off between the SRPT (Shortest Remaining Processing Time) discipline and opportunistic scheduling in the systems characterized by compact and symmetric capacity regions. The objective is to minimize the mean delay in a transient setting where all jobs are available at time 0 and no new jobs arrive thereafter. Our main result gives conditions under which the optimal rate vector does not depend on the sizes of the jobs as long as their order (in size) remains the same. In addition, it shows that in this case the optimal policy applies the SRPT principle serving the shortest job with the highest rate of the optimal rate vector, the second shortest with the second highest rate etc. We also give a recursive algorithm to determine both the optimal rate vector and the minimum mean delay. In some special cases, the rate vector, as well as the minimum mean delay, have even explicit expressions as demonstrated in the paper. For the general case, we derive both an upper bound and a lower bound of the minimum mean delay.

1. INTRODUCTION

Modern wireless cellular systems allow highly sophisticated scheduling algorithms to be used for sharing the radio resources among the users. These systems operate in slotted time with a very short time slot duration (at millisecond time scale). The base station also has access to information about the instantaneous transmission reception conditions of every user, essentially the instantaneous transmission rates,

which are randomly varying over time due to various fading phenomena. This has given rise to *opportunistic scheduling*, where the idea is to favor those users with instantaneously high transmission rates. It is clear that the more users there are in the system, the more likely it is to have some user in a good state. Thus, the overall service rate of the system increases with the number of users, i.e., there is opportunistic scheduling gain (also sometimes referred to as multiuser diversity gain).

A well-known example of opportunistic scheduling is the PF (Proportionally Fair) scheduler [16], which combines information about the instantaneous rates with the throughput. PF belongs to a more general class of utility-based α -fair schedulers [14, 6]. Other examples of opportunistic schedulers include maxweight schedulers that combine the channel information with the queue lengths or delays [13], or rate-based schedulers that only use information about the channel-statistics [2, 3, 8]. Max-weight and utility-based schedulers have been analyzed at the time-slot level (or packet level) by assuming a fixed population of users and they have been shown to exhibit certain optimality properties, see [13, 14].

In reality the number of active users varies in the system. Models that take this into account are called flow-level models. Flows roughly represent file transfers controlled by TCP (elastic data) and the performance at flow-level characterizes, e.g., the mean delay of file transfers. An important feature in the model is the assumption of time scale separation between the flow-level dynamics and the time-slot level channel dynamics, see also [5, 4, 11]. This implies that, at the flow-level, the flows observe the time-average throughput provided by the time-slot level scheduler. For a given number of flows in the system, the set of achievable rate vectors that the time-slot level scheduler can support is characterized by the notion of the *capacity region*, see [5, 4, 11]. The general scheduling problem at the flow-level is then to determine the rate vector to be used within the capacity region given the current state at the flow-level so that, e.g., the mean flow delay is minimized.

In the *dynamic setting* the system at the flow level consists of random flow arrivals and departures. In this setting the existing provable properties of opportunistic scheduling policies are limited to results on their stability properties,

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

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

see [5, 6, 4, 1]. Utility-based policies of the throughput have been shown to achieve the maximal stability region [6], while rate-based policies may suffer from instability [5, 1].

Size-based scheduling is known to be a good choice for the systems *without* any opportunistic gain. For example, for classical single-server queues, it is well-known that the SRPT (Shortest-Remaining-Processing-Time) discipline is optimal minimizing the mean delay [12]. The idea of SRPT is to minimize the delay by getting rid of flows as soon as possible. However, with fewer flows part of the opportunistic gain is lost. As a result, combining the advantages of size-based scheduling with opportunistic scheduling gain has proven very challenging. Results on the optimal policy to minimize the mean flow-level delay in the dynamic setting for the systems with opportunistic scheduling gain are not available, owing to the difficulty of the problem. Only some heuristic algorithms have been proposed and experimented with that try to combine opportunistic scheduling gain with size-based scheduling [7, 11].

All optimality results concerning the minimization of the mean flow-level delay under opportunistic scheduling gain are related to the *transient system*, where there are initially n flows with given sizes but no new arrivals. It has been shown that in this case the optimal scheduling problem can be formulated as a dynamic program [15]. However, the dynamic program does not allow to extract any structural properties of the optimal policy. The *only* structural optimality result (that we know) appeared recently in [11], where the authors show that when the capacity regions are nested polymatroids and the opportunistic gain γ_n is increasing and concave in the number n of flows, the optimal policy is an SRPT-type discipline such that the shortest flow is served with the highest possible rate γ_1 , the second shortest with the second highest rate $\gamma_2 - \gamma_1$, etc. Importantly, the optimal rate vector does *not* depend on the sizes of the flows as long as their order (in size) remains the same. In their proof, the authors of [11] utilize the known optimality result of the SRPT-FM (SRPT-Fastest-Machine) discipline for the heterogeneous multiserver queues [9].

As in [15, 11], we consider the optimal scheduling problem under opportunistic scheduling gain in the transient setting. We focus on the situation, where the capacity regions are compact and symmetric (including, e.g., all nested polymatroids). From the wireless channel point of view, the symmetry assumption implies that the random variations at the time slot level experienced by the flows are statistically identical, see [5]. Our aim is to minimize the mean flow-level delay by determining the optimal trade-off between between SRPT and opportunistic scheduling.

Our main result gives conditions under which the optimal rate vector does not depend on the sizes of the flows as long as their order (in size) remains the same. In addition, it shows that in this case the optimal policy applies the SRPT principle serving the shortest flow with the highest rate of the optimal rate vector, the second shortest with the second highest rate etc. We also give a recursive algorithm to determine both the optimal rate vector and the minimum mean delay. In addition to the theoretic value, the presented approach provides a vast improvement over any general optimization method for numerical evaluation of the optimal scheduling problem. Since we specify the conditions for *any* family of compact and symmetric capacity regions, the result is also essentially more general than the result

given in [11] (concerning only nested polymatroids). In some special cases, the rate vector, as well as the minimum flow time, have even explicit expressions as demonstrated in the paper. In addition to the lower bound (i.e., an optimistic estimate) given already in [11], we derive an upper bound (i.e., a conservative estimate) of the minimum mean delay in the general case. Numerically, the upper bound seems to be closer to the optimum value.

The rest of the paper is organized as follows. The optimal scheduling problem is formulated and some important operating policies are introduced in Section 2. Section 3 includes the main theoretic results for compact and symmetric capacity regions. In Section 4, we demonstrate that the optimality result for nested polymatroids given in [11] is essentially a special case of our main result. Nested polymatroids are also used to determine an optimistic estimate for the minimum mean delay. Section 4 considers another family of polytopes for which the optimality result is a special case of our main result. These polytopes can be utilized to determine a conservative estimate for the minimum mean delay. In Section 6, we consider a parametric family of capacity regions (so-called α -balls) for which the optimal rate vector and the minimum flow time have explicit expressions. We also work out some numerical examples for illustrative purposes. Section 7 concludes the paper and also discusses some future research directions.

2. PROBLEM FORMULATION

Consider a service system where the service capacity is adjustably depending on the current number of jobs.¹ More precisely said, when there are k jobs in the system (indexed with $i = 1, \dots, k$), the operator (of the system) chooses a rate vector $\mathbf{c} = (c_1, \dots, c_k)$ from the capacity region $\mathcal{C}_k \subset \mathbb{R}_+^k$. From that on, each job i is served with rate c_i until the number of jobs again changes, and a new rate vector is chosen. We assume that when choosing the rate vector the operator is aware of the (remaining) sizes of the jobs.

Assume now that, at time 0, there are n jobs in the system with sizes

$$s_1 \geq \dots \geq s_n.$$

We consider the transient system so that we do not allow any further arrivals.

An *operating policy* π is defined by a sequence of vectors $\mathbf{c}_k = (c_{k1}, \dots, c_{kk}) \in \mathcal{C}_k$ for all $k = 1, \dots, n$, where \mathbf{c}_k refers to the rate vector that the operator applies when there are k jobs in the system (called hereafter as *phase* k). It is assumed that when a job completes, the remaining $k - 1$ jobs are *re-indexed* in such a way that the remaining sizes $s_{k-1,i}$ again satisfy

$$s_{k-1,1} \geq \dots \geq s_{k-1,k-1}.$$

Thus, in the next step, the longest job is served with rate $c_{k-1,1}$, the second longest with rate $c_{k-1,2}$, etc. Let Π_n denote the family of all operating policies,

$$\Pi_n = \{\pi = (\mathbf{c}_1, \dots, \mathbf{c}_n) : \mathbf{c}_k \in \mathcal{C}_k \text{ for all } k\}.$$

Let then t_i^π denote the time when the job with original index i completes under policy π . As usual in scheduling

¹Since the concept job is often used in scheduling literature rather than flow, we will subsequently use the words job/flow interchangeably.

literature, the *flow time* T^π (a.k.a. total completion time) is defined as

$$T^\pi = \sum_{i=1}^n t_i^\pi.$$

Note that the mean delay of a job is now given by T^π/n .

In this paper, we consider the scheduling problem in which the optimal operating policy minimizes the flow time (or the mean delay, as well). Let π^* denote such an optimal policy. Thus,

$$T^{\pi^*} = \min_{\pi \in \Pi_n} T^\pi,$$

where the minimization is taken over all operating policies defined by the n fixed capacity regions $\mathcal{C}_1, \dots, \mathcal{C}_n$.

Finally we define an important category of operating policies. Policy π belongs to the class of *SRPT-HPR policies* if the corresponding rate vectors $(\mathbf{c}_1, \dots, \mathbf{c}_n)$ satisfy the following condition, for all k and j ,

$$c_{k,j} \leq c_{k,j+1}.$$

An example is given by the ordinary *SRPT* discipline that serves only the shortest job so that the rate vectors \mathbf{c}_k take the form

$$\mathbf{c}_k = (0, \dots, 0, \gamma_1),$$

where $\gamma_1 > 0$. Another example is given by the *OPS* (Opportunistic Processor Sharing) discipline, which takes the most out of the opportunistic gain in a fair way and which has been used to model the behaviour of the PF scheduler at the flow level under certain circumstances, see [5, 11]. For the OPS discipline the rate vectors \mathbf{c}_k are of form

$$\mathbf{c}_k = \left(\frac{\gamma_k}{k}, \dots, \frac{\gamma_k}{k} \right),$$

where γ_k is an increasing positive sequence referring to the opportunistic gain. An intermediate version of the two (extreme) policies, called *SRPT-OPS*, was introduced in [11]. For the SRPT-OPS discipline the rate vectors \mathbf{c}_k read as

$$\mathbf{c}_k = \left(0, \dots, 0, \frac{\gamma_{j_k}}{j_k}, \dots, \frac{\gamma_{j_k}}{j_k} \right)$$

with j_k non-zero elements, where γ_k is an increasing positive sequence referring to the opportunistic gain and j_k indicates the number of jobs (out of k) that share the service capacity available.

3. SYMMETRIC CAPACITY REGIONS

In this section we assume that the capacity regions \mathcal{C}_k have the following two properties for all $k = 1, \dots, n$:

- (i) \mathcal{C}_k is compact region of \mathbb{R}_+^k , i.e., \mathcal{C}_k is closed and bounded;
- (ii) \mathcal{C}_k is *symmetric*, i.e., if $\mathbf{c} \in \mathcal{C}_k$, then any permutation $\tilde{\mathbf{c}}$ of its components also lies in \mathcal{C}_k .

If there is only one job, $n = 1$, then the optimal policy π^* is clearly defined by the maximal service capacity

$$c_1^* = \max\{c \in \mathcal{C}_1\},$$

the existence of which is guaranteed by the compactness property (i) above. We note that the optimal policy π^* is independent of the size s_1 of the job.

Now we consider the general case where there is any number of jobs, $n \geq 1$. An operating policy π is defined by a

sequence of vectors $\mathbf{c}_k = (c_{k1}, \dots, c_{kk}) \in \mathcal{C}_k$ for $k = 1, \dots, n$. The flow time of π reads as

$$T^\pi = \sum_{k=1}^n k T_k^\pi,$$

where T_k^π refers to the length of phase k for policy π .

Let g_1, \dots, g_n be a sequence of functions with $g_k(\mathbf{c}_k)$ defined on \mathcal{C}_k for all k , G_1^*, \dots, G_n^* a sequence of positive real numbers, and $\mathbf{c}_1^*, \dots, \mathbf{c}_n^*$ a sequence of vectors with $\mathbf{c}_k^* \in \mathcal{C}_k$ for all k . These sequences are defined recursively as follows:

$$\begin{aligned} g_1(c_1) &= \frac{1}{c_{11}}, \\ G_1^* &= g_1(c_1^*) = \min_{c_1 \in \mathcal{C}_1} g_1(c_1), \\ g_k(\mathbf{c}_k) &= \frac{1}{c_{kk}} \left(k - \sum_{j=1}^{k-1} c_{kj} G_j^* \right), \\ G_k^* &= g_k(\mathbf{c}_k^*) = \min_{\mathbf{c}_k \in \mathcal{C}_k} g_k(\mathbf{c}_k), \quad k = 2, \dots, n. \end{aligned} \tag{1}$$

Note that the existence of the minimum values G_k^* is guaranteed by the compactness of capacity regions \mathcal{C}_k (Property (i)). Note also that function $g_k(\mathbf{c}_k)$ do not depend on the sizes s_1, \dots, s_n of the jobs.

PROPOSITION 1. *If the capacity regions $\mathcal{C}_1, \dots, \mathcal{C}_n$ are such that*

$$G_1^* < \dots < G_n^*,$$

then $c_{k,j+1}^ \geq c_{k,j}^*$ for all $k = 2, \dots, n$ and $j = 1, \dots, k-1$.*

PROOF. 1° Let $k \in \{2, \dots, n\}$ and $j \in \{1, \dots, k-2\}$. Suppose (contrary to our claim) that $c_{k,j+1}^* < c_{k,j}^*$. We will show below that this results in a conflict with our assumptions.

Let $\tilde{\mathbf{c}}_k^*$ denote the modification of \mathbf{c}_k^* where the service rates $c_{k,j}^*$ and $c_{k,j+1}^*$ have changed their places,

$$\tilde{\mathbf{c}}_k^* = (c_{k,1}^*, \dots, c_{k,j-1}^*, c_{k,j+1}^*, c_{k,j}^*, c_{k,j+2}^*, \dots, c_{k,k}^*).$$

Note that $\tilde{\mathbf{c}}_k^* \in \mathcal{C}_k$ since \mathcal{C}_k is symmetric (Property (ii)). Now

$$\begin{aligned} g_k(\mathbf{c}_k^*) - g_k(\tilde{\mathbf{c}}_k^*) &= G_j^* c_{k,j}^* + G_{j+1}^* c_{k,j+1}^* - G_j^* c_{k,j+1}^* - G_{j+1}^* c_{k,j}^* \\ &= (G_j^* - G_{j+1}^*) c_{k,j}^* - (G_j^* - G_{j+1}^*) c_{k,j+1}^* \\ &= (G_j^* - G_{j+1}^*) (c_{k,j}^* - c_{k,j+1}^*) \\ &< 0, \end{aligned}$$

which contradicts the definition of \mathbf{c}_k^* as the minimum point of function $g_k(\mathbf{c}_k)$.

2° Consider now the remaining case where $k \in \{2, \dots, n\}$ and $j = k-1$. First we note that since

$$G_{k-1}^* < G_k^* = g_k(\tilde{\mathbf{c}}_k^*) = \frac{1}{c_{kk}^*} \left(k - \sum_{j=1}^{k-1} G_j^* c_{kj}^* \right),$$

we have

$$c_{k,k}^* + c_{k,k-1}^* < 2a,$$

where constant a is defined as follows:

$$a = \frac{1}{2G_{k-1}^*} \left(k - \sum_{j=1}^{k-2} G_j^* c_{kj}^* \right).$$

Suppose (again contrary to our claim) that $c_{k,k}^* < c_{k,k-1}^*$. We will show below that also this results in a conflict with

our assumptions. Before that, we note that $c_{k,k}^* < a$, since $c_{k,k}^* + c_{k,k-1}^* < 2a$ and $c_{k,k}^* < c_{k,k-1}^*$.

Let $\tilde{\mathbf{c}}_k^*$ denote the modification of \mathbf{c}_k^* where the service rates $c_{k,k-1}^*$ and $c_{k,k}^*$ have changed their places,

$$\tilde{\mathbf{c}}_k^* = (c_{k,1}^*, \dots, c_{k,k-2}^*, c_{k,k}^*, c_{k,k-1}^*).$$

Note that $\tilde{\mathbf{c}}_k^* \in \mathcal{C}_k$ since \mathcal{C}_k is symmetric (Property (ii)). In addition, let an auxiliary function $f(x)$ be defined on \mathbb{R} as follows:

$$f(x) = 2ax - x^2.$$

2.1° Assume that $c_{k,k}^* < c_{k,k-1}^* \leq a$. Since $f(x)$ is strictly increasing for all $x \leq a$, we have

$$\begin{aligned} f(c_{k,k}^*) &< f(c_{k,k-1}^*) \\ \Leftrightarrow 2ac_{k,k}^* - (c_{k,k}^*)^2 &< 2ac_{k,k-1}^* - (c_{k,k-1}^*)^2 \\ \Leftrightarrow \frac{1}{c_{k,k-1}^*} \left(2 - \frac{c_{k,k}^*}{a}\right) &< \frac{1}{c_{k,k}^*} \left(2 - \frac{c_{k,k-1}^*}{a}\right) \\ \Leftrightarrow g_k(\tilde{\mathbf{c}}_k^*) &< g(\mathbf{c}_k^*), \end{aligned}$$

which contradicts the definition of \mathbf{c}_k^* as the minimum point of function $g(\mathbf{c}_k)$ in \mathcal{C}_k .

2.2° Assume now that $c_{k,k}^* \leq a < c_{k,k-1}^*$. We will show that also in this case $f(c_{k,k}^*) < f(c_{k,k-1}^*)$, which leads to a contradiction as shown in 2.1°. Since $c_{k,k}^* + c_{k,k-1}^* < 2a$ and $c_{k,k-1}^* > a$, we have $c_{k,k}^* < 2a - c_{k,k-1}^* \leq a$, implying that

$$\begin{aligned} f(c_{k,k-1}^*) - f(c_{k,k}^*) &> f(c_{k,k-1}^*) - f(2a - c_{k,k-1}^*) \\ &= 2ac_{k,k-1}^* - (c_{k,k-1}^*)^2 - 2a(2a - c_{k,k-1}^*) + (2a - c_{k,k-1}^*)^2 \\ &= 0, \end{aligned}$$

which completes the proof. \square

THEOREM 1. *If the capacity regions $\mathcal{C}_1, \dots, \mathcal{C}_n$ are such that*

$$G_1^* < \dots < G_n^*,$$

then the optimal operating policy is $\pi^ = (\mathbf{c}_1^*, \dots, \mathbf{c}_n^*)$ for all sizes $s_1 \geq \dots \geq s_n$, where the optimal rate vectors \mathbf{c}_k^* are defined recursively in (1). In this case, the minimum flow time T^{π^*} satisfies*

$$T^{\pi^*} = \sum_{k=1}^n s_k G_k^*.$$

In addition, $c_{k,j+1}^ \geq c_{k,j}^*$ for all $k = 2, \dots, n$ and $j = 1, \dots, k-1$ so that the optimal policy belongs to the SRPT-HPR category.*

PROOF. The result is proved by induction. For $n = 1$, the result is clearly true:

$$T^{\pi^*} = \frac{s_1}{c_1^*} = \min_{c_1 \in \mathcal{C}_1} \frac{s_1}{c_1} = \min_{\pi} T^{\pi}.$$

In addition, $G_1^* = \frac{1}{c_1^*}$ so that $T^{\pi^*} = s_1 G_1^*$ as claimed.

Assume now that $n \geq 2$ and the result is true for all values $1, \dots, n-1$. We will show that it is also true for value n .

It follows from the induction assumption that the optimal policy applies rate vectors \mathbf{c}_k^* for all $k = 1, \dots, n-1$. Thus, for any policy $\pi = (\mathbf{c}_1, \dots, \mathbf{c}_n) \in \Pi_n$, the modified policy $\tilde{\pi} = (\mathbf{c}_1^*, \dots, \mathbf{c}_{n-1}^*, \mathbf{c}_n) \in \Pi_n$ results in a smaller flow time so

that

$$\begin{aligned} T^{\pi} &\geq T^{\tilde{\pi}} \\ &= nT_n^{\tilde{\pi}} + \sum_{k=1}^{n-1} kT_k^{\tilde{\pi}} \\ &= nT_n^{\tilde{\pi}} + \sum_{k=1}^{n-1} \left(s_{i(k)} - T_n^{\tilde{\pi}} c_{n,i(k)}\right) G_k^* \\ &= n \frac{s_{i(n)}}{c_{n,i(n)}} + \sum_{k=1}^{n-1} \left(s_{i(k)} - \frac{s_{i(n)}}{c_{n,i(n)}} c_{n,i(k)}\right) G_k^* \\ &= \frac{s_{i(n)}}{c_{n,i(n)}} \left(n - \sum_{k=1}^{n-1} c_{n,i(k)} G_k^*\right) + \sum_{k=1}^{n-1} s_{i(k)} G_k^* \\ &= s_{i(n)} g((c_{n,i(1)}, \dots, c_{n,i(n)})) + \sum_{k=1}^{n-1} s_{i(k)} G_k^*, \end{aligned}$$

where $i(k)$ refers to the original index of the job that completes at the end of phase k under policy $\tilde{\pi}$. Note that $(c_{n,i(1)}, \dots, c_{n,i(n)}) \in \mathcal{C}_n$, since $\mathbf{c}_n = (c_{n1}, \dots, c_{nn}) \in \mathcal{C}_n$ and \mathcal{C}_n is symmetric (Property (ii)). Thus,

$$g((c_{n,i(1)}, \dots, c_{n,i(n)})) \geq G_n^*$$

implying that

$$T^{\pi} \geq \sum_{k=1}^n s_{i(k)} G_k^* \geq \sum_{k=1}^n s_k G_k^*,$$

where the latter inequality follows from the facts that $s_1 \geq \dots \geq s_n$ and $G_1^* < \dots < G_n^*$.

Consider then policy $\pi^* = (\mathbf{c}_1^*, \dots, \mathbf{c}_n^*)$ and let $i^*(k)$ denote the original index of the job that completes at the end of phase k under this policy π^* . It follows from Proposition 1 that $i^*(k) = k$ for all k . Thus,

$$\begin{aligned} T^{\pi^*} &= nT_n^{\pi^*} + \sum_{k=1}^{n-1} kT_k^{\pi^*} \\ &= n \frac{s_n}{c_{nn}^*} + \sum_{k=1}^{n-1} \left(s_k - \frac{s_n}{c_{nn}^*} c_{nk}^*\right) G_k^* \\ &= \frac{s_n}{c_{nn}^*} \left(n - \sum_{k=1}^{n-1} c_{nk}^* G_k^*\right) + \sum_{k=1}^{n-1} s_k G_k^* \\ &= s_n g(\mathbf{c}_n^*) + \sum_{k=1}^{n-1} s_k G_k^* \\ &= \sum_{k=1}^n s_k G_k^* \end{aligned}$$

so that $T^{\pi} \geq T^{\pi^*}$ for any $\pi \in \Pi_n$. \square

We would emphasize that all the results in this section are achieved with very general assumptions. Unlike in [11], no convexity nor coordinate-convexity is required from the capacity regions. It is only assumed that the capacity regions are compact and symmetric.

The capacity regions do not even need to be nested. An easy example can be found for $n = 2$. If $\mathcal{C}_1 = [0, 1]$ and $\mathcal{C}_2 \in \mathbb{R}_2^k$ is a compact and symmetric region such that $c_{21} + c_{22} < 2$, then $G_1^* = 1$ and

$$G_2^* = \frac{1}{c_{22}^*} (2 - c_{21}^*) > 1 = G_1^*$$

so that Theorem 1 can be applied to determine the optimal policy.

Note also that (1) gives a recursive algorithm to determine both the optimal rate vector and the minimum mean delay. The proposed approach vastly facilitates the numerical evaluation of the optimal scheduling problem for any family of capacity regions that meet the presented conditions. The general optimization problem is difficult: (i) The number of possible service orders becomes quickly overwhelming when the number of flows increases, (ii) optimization on a high-dimensional capacity set may be computationally tedious, e.g., when the capacity region is a solution space of some packet level scheduling problem, and (iii) the problem needs to be solved separately for each set of flow sizes. The proposed approach avoids the combinatorial problems altogether, minimizes the need for numerical optimization on capacity sets and produces results that can be readily recycled for different flow sizes.

4. SYMMETRIC POLYMATROIDS

Capacity regions C_k , $k = 1, \dots, n$, are nested and symmetric *polymatroids* if there is an increasing sequence $\gamma_k \in \mathbb{R}_+$ (referring to the opportunistic gain) such that, for all k ,

$$C_k = \{\mathbf{c}_k \in \mathbb{R}_+^k : \sum_{i \in \mathcal{I}} c_{ki} \leq \gamma_{|\mathcal{I}|} \text{ for all } \mathcal{I} \in \{1, \dots, n\}\}.$$

Sadiq and de Veciana [11] proved that the optimal policy belongs to the SRPT-HPR category when the capacity regions are nested and symmetric polymatroids and the gain function γ_k is increasing and concave, i.e., $\gamma_{k+1} - \gamma_k$ is decreasing. In this section, we demonstrate that (with a minor additional assumption) their result is, in fact, a special case of our Theorem 1.

Given an increasing sequence γ_k , let $\theta_1, \dots, \theta_n$ denote a sequence of positive real numbers defined recursively as follows:

$$\begin{aligned} \theta_1 &= \frac{1}{\gamma_1}, \\ \theta_k &= \frac{1}{\gamma_1} \left(k - \sum_{j=1}^{k-1} (\gamma_{k+1-j} - \gamma_{k-j}) \theta_j \right), \quad k = 2, \dots, n. \end{aligned} \quad (2)$$

Sadiq [10, Proof of Theorem 5.1] has shown that the sequence θ_k is increasing when the sequence γ_k is concave, i.e., $\gamma_{k+1} - \gamma_k$ is decreasing as a function of k . Below we show that it is strictly increasing when the sequence γ_k is *strictly concave*, i.e.,

$$\gamma_1 > \gamma_2 - \gamma_1 > \dots > \gamma_n - \gamma_{n-1}.$$

PROPOSITION 2. *If the increasing sequence γ_k is strictly concave, then*

$$\theta_1 < \dots < \theta_n.$$

PROOF. The result is proved by induction. For $n = 1$, the result is trivially true.

Assume now that $n \geq 2$ and the result is true for all values $1, \dots, n-1$. We will show that it is also true for value n .

Let us denote $\gamma_0 = \theta_0 = 0$. It follows from the definition of θ_n that

$$n = \sum_{k=0}^{n-1} (\gamma_{k+1} - \gamma_k) \theta_{n-k}.$$

Correspondingly, by the definition of θ_{n-1} ,

$$n-1 = \sum_{k=0}^{n-2} (\gamma_{k+1} - \gamma_k) \theta_{n-1-k}.$$

The difference of these two equations gives thus

$$1 = \sum_{k=0}^{n-1} (\gamma_{k+1} - \gamma_k) (\theta_{n-k} - \theta_{n-1-k}).$$

By substituting n with $n-1$, we get

$$1 = \sum_{k=0}^{n-2} (\gamma_{k+1} - \gamma_k) (\theta_{n-1-k} - \theta_{n-2-k}).$$

Since $\gamma_{k+2} - \gamma_{k+1} < \gamma_{k+1} - \gamma_k$ for all k , it follows that

$$\begin{aligned} 1 &= \sum_{k=0}^{n-1} (\gamma_{k+1} - \gamma_k) (\theta_{n-k} - \theta_{n-1-k}) \\ &= \gamma_1 (\theta_n - \theta_{n-1}) + \sum_{k=0}^{n-2} (\gamma_{k+2} - \gamma_{k+1}) (\theta_{n-1-k} - \theta_{n-2-k}) \\ &< \gamma_1 (\theta_n - \theta_{n-1}) + \sum_{k=0}^{n-2} (\gamma_{k+1} - \gamma_k) (\theta_{n-1-k} - \theta_{n-2-k}) \\ &= \gamma_1 (\theta_n - \theta_{n-1}) + 1. \end{aligned}$$

Thus, $\gamma_1 (\theta_n - \theta_{n-1}) > 0$, implying that $\theta_n > \theta_{n-1}$, since $\gamma_1 > 0$. \square

To prove the main result of this section (given below in Theorem 2) we need the following auxiliary result.

PROPOSITION 3. *If the increasing sequence γ_k is strictly concave, then, for all $k = 2, \dots, n$ and $\mathbf{c}_k \in C_k$,*

$$\sum_{j=1}^{k-2} c_{kj} \theta_j + (c_{k,k-1} + c_{kk}) \theta_{k-1} \leq \sum_{j=1}^{k-2} (\gamma_{k-j+1} - \gamma_{k-j}) \theta_j + \gamma_2 \theta_{k-1}.$$

PROOF. The result follows easily from the facts that

$$0 < \theta_1 < \dots < \theta_{k-1}$$

and

$$\sum_{j=k-k'+1}^k c_{kj} \leq \gamma_{k'} = \sum_{j=3}^{k'} (\gamma_j - \gamma_{j-1}) + \gamma_2$$

for all $k' = 2, \dots, k$. The former is due to Proposition 2 and the latter follows from the properties of polymatroid C_k . \square

THEOREM 2. *If the capacity regions C_k , $k = 1, \dots, n$, are nested and symmetric polymatroids generated by an increasing and strictly concave sequence γ_k , then the optimal operating policy, for all sizes $s_1 \geq \dots \geq s_n$, is $\pi^* = (\mathbf{c}_1^*, \dots, \mathbf{c}_n^*)$, where*

$$\mathbf{c}_k^* = (\gamma_k - \gamma_{k-1}, \dots, \gamma_2 - \gamma_1, \gamma_1)$$

for all k . In this case, the minimum flow time T^{π^*} satisfies

$$T^{\pi^*} = \sum_{k=1}^n s_k \theta_k,$$

where θ_k 's are defined in (2). In addition, $c_{k,j+1}^* > c_{k,j}^*$ for all $k = 2, \dots, n$ and $j = 1, \dots, k-1$ so that the optimal policy belongs to the SRPT-HPR category.

PROOF. By Theorem 1 and Proposition 2, it is sufficient to prove that, for all k ,

$$\theta_k = \min_{\mathbf{c}_k \in \mathcal{C}_k} g_k(\mathbf{c}_k),$$

where functions g_k are defined in (1).

The result is proved by induction. For $k = 1$, the result is clearly true.

Assume now that $k \geq 2$ and the result is true for all values $1, \dots, k-1$. We will show that it is also true for value k .

Note first that

$$\theta_k - \theta_{k-1} = \frac{1}{\gamma_1} \left(k - \sum_{j=1}^{k-2} (\gamma_{k-j+1} - \gamma_{k-j}) \theta_j - \gamma_2 \theta_{k-1} \right) > 0$$

by Proposition 2. In addition, by the induction assumption, for any $j = 1, \dots, k-1$,

$$\theta_j = \min_{\mathbf{c}_j \in \mathcal{C}_j} g_j(\mathbf{c}_j) = G_j^*.$$

Thus, for any $\mathbf{c}_k \in \mathcal{C}_k$,

$$\begin{aligned} & g_k(\mathbf{c}_k) - g_k(\mathbf{c}_k^*) \\ &= \frac{1}{c_{kk}} \left(k - \sum_{j=1}^{k-1} c_{kj} \theta_j \right) \\ &\quad - \frac{1}{\gamma_1} \left(k - \sum_{j=1}^{k-1} (\gamma_{k-j+1} - \gamma_{k-j}) \theta_j \right) \\ &= \frac{1}{c_{kk}} \left(k - \sum_{j=1}^{k-2} c_{kj} \theta_j - (c_{k,k-1} + c_{kk}) \theta_{k-1} \right) \\ &\quad - \frac{1}{\gamma_1} \left(k - \sum_{j=1}^{k-2} (\gamma_{k-j+1} - \gamma_{k-j}) \theta_j - \gamma_2 \theta_{k-1} \right) \\ &\geq \frac{1}{c_{kk}} \left(k - \sum_{j=1}^{k-2} (\gamma_{k-j+1} - \gamma_{k-j}) \theta_j - \gamma_2 \theta_{k-1} \right) \\ &\quad - \frac{1}{\gamma_1} \left(k - \sum_{j=1}^{k-2} (\gamma_{k-j+1} - \gamma_{k-j}) \theta_j - \gamma_2 \theta_{k-1} \right) \\ &= \left(\frac{1}{c_{kk}} - \frac{1}{\gamma_1} \right) \gamma_1 (\theta_k - \theta_{k-1}) \geq 0, \end{aligned}$$

where the inequality is justified by Proposition 3. \square

Note that the result given above is valid even for more general capacity regions, which include the operating points $\mathbf{c}_k^* = (\gamma_k - \gamma_{k-1}, \dots, \gamma_2 - \gamma_1, \gamma_1)$ and which are subsets of the corresponding polymatroids. Note also that such capacity regions are not required to be symmetric.

As already shown in [11], symmetric polymatroids can be utilized to determine a lower bound (i.e., an *optimistic estimate*) for the flow time whenever the original capacity regions are compact, convex and symmetric, cf. Figure 1 in Section 6.

5. SYMMETRIC OPS-LIMITED POLYTOPES

Capacity regions \mathcal{C}_k , $k = 1, \dots, n$, are nested and symmetric OPS-limited polytopes if there is an increasing sequence $\gamma_k \in \mathbb{R}_+$ (referring to the opportunistic gain) such that, \mathcal{C}_k

is the convex hull of the points $\mathcal{V}_k \subset \mathbb{R}_+^k$ for all k , where

$$\mathcal{V}_k = \bigcup_{j=0}^k \mathcal{V}_{kj}$$

and \mathcal{V}_{kj} consists of all permutations of the rate vector

$$(0, \dots, 0, \frac{\gamma_j}{j}, \dots, \frac{\gamma_j}{j})$$

with j non-zero elements. Note that these permutations correspond to rate vectors for the OPS policy when applied to j jobs (out of k). It follows that

$$\mathcal{C}_k = \left\{ \sum_{i=1}^{|\mathcal{V}_k|} \alpha_i \mathbf{v}_{ki} : \mathbf{v}_{ki} \in \mathcal{V}_k, \alpha_i \geq 0, \sum_{i=1}^{|\mathcal{V}_k|} \alpha_i = 1 \right\}.$$

It is also easy to see that $|\mathcal{V}_k| = 2^k$ and $|\mathcal{V}_{kj}| = \frac{k!}{j!(k-j)!}$.

Given an increasing sequence γ_k , let η_1, \dots, η_n denote a sequence of positive real numbers defined recursively as follows:

$$\begin{aligned} \eta_1 &= \frac{1}{\gamma_1}, \\ \eta_k &= \min_{j \in \{1, \dots, k\}} \left(\frac{j k}{\gamma_j} - \sum_{i=k+1-j}^{k-1} \eta_i \right), \quad k = 2, \dots, n. \end{aligned} \quad (3)$$

PROPOSITION 4. *If the increasing sequence γ_k is such that*

$$\eta_1 < \dots < \eta_n,$$

then, for all $k = 1, \dots, n$ and $\mathbf{v}_k = (v_{k1}, \dots, v_{kk}) \in \mathcal{V}_k$,

$$\sum_{j=1}^k v_{kj} \eta_j \leq k.$$

PROOF. For $\mathbf{v}_k = (0, \dots, 0) \in \mathcal{V}_k$ the result is trivially true. Let then $\mathbf{v}_k \in \mathcal{V}_k$ such that

$$\mathbf{v}_k = (0, \dots, 0, \frac{\gamma_j}{j}, \dots, \frac{\gamma_j}{j})$$

with $j \in \{1, \dots, k\}$ non-zero elements. It follows from (3) that

$$\eta_k \leq \left(\frac{j k}{\gamma_j} - \sum_{i=k+1-j}^{k-1} \eta_i \right),$$

which is equivalent with

$$\sum_{j=1}^k v_{kj} \eta_j \leq k.$$

Consider then any permutation $\tilde{\mathbf{v}}_k = (\tilde{v}_{k1}, \dots, \tilde{v}_{kk})$ of \mathbf{v}_k . The assumption

$$\eta_1 < \dots < \eta_n$$

implies that

$$\sum_{j=1}^k \tilde{v}_{kj} \eta_j \leq \sum_{j=1}^k v_{kj} \eta_j \leq k,$$

which completes the proof. \square

THEOREM 3. *If the capacity regions \mathcal{C}_k , $k = 1, \dots, n$, are nested and symmetric OPS-limited polytopes generated by an increasing sequence γ_k such that*

$$\eta_1 < \dots < \eta_n,$$

then the optimal operating policy, for all sizes $s_1 \geq \dots \geq s_n$, is the SRPT-OPS policy $\pi^* = (\mathbf{c}_1^*, \dots, \mathbf{c}_n^*)$, where, for all k ,

$$\mathbf{c}_k^* = (0, \dots, 0, \frac{\gamma_{j_k^*}}{j_k^*}, \dots, \frac{\gamma_{j_k^*}}{j_k^*})$$

with j_k^* non-zero elements, where j_k^* is the optimal index in (3). In this case, the minimum flow time T^{π^*} satisfies

$$T^{\pi^*} = \sum_{k=1}^n s_k \eta_k.$$

PROOF. By Theorem 1, it is sufficient to prove that, for all k ,

$$\eta_k = \min_{\mathbf{c}_k \in \mathcal{C}_k} g_k(\mathbf{c}_k),$$

where functions g_k are defined in (1).

The result is proved by induction. For $k = 1$, the result is clearly true.

Assume now that $k \geq 2$ and the result is true for all values $1, \dots, k-1$. We will show that it is also true for value k .

Let $\mathbf{c}_k = (c_{k1}, \dots, c_{kk}) \in \mathcal{C}_k$. There are $\alpha_i \geq 0$ such that

$$\sum_{i=1}^{|\mathcal{V}_k|} \alpha_i = 1 \quad \text{and} \quad \mathbf{c}_k = \sum_{i=1}^{|\mathcal{V}_k|} \alpha_i \mathbf{v}_{ki},$$

where $\mathbf{v}_{ki} = (v_{ki1}, \dots, v_{kik}) \in \mathcal{V}_k$. By Proposition 4,

$$\sum_{j=1}^k c_{kj} \eta_j = \sum_{j=1}^k \sum_{i=1}^{|\mathcal{V}_k|} \alpha_i v_{kij} \eta_j = \sum_{i=1}^{|\mathcal{V}_k|} \alpha_i \sum_{j=1}^k v_{kij} \eta_j \leq k,$$

implying, by the induction assumption, that

$$g_k(\mathbf{c}_k) = \frac{1}{c_{kk}} \left(k - \sum_{j=1}^{k-1} c_{kj} G_j^* \right) = \frac{1}{c_{kk}} \left(k - \sum_{j=1}^{k-1} c_{kj} \eta_j \right) \geq \eta_k.$$

On the other hand, by (3),

$$\sum_{j=1}^k c_{kj}^* \eta_j = \frac{\gamma_{j^*(k)}}{j^*(k)} \sum_{j=k+1-j^*(k)}^k \eta_j = k,$$

implying that

$$g_k(\mathbf{c}_k^*) = \frac{1}{c_{kk}^*} \left(k - \sum_{j=1}^{k-1} c_{kj}^* G_j^* \right) = \frac{1}{c_{kk}^*} \left(k - \sum_{j=1}^{k-1} c_{kj}^* \eta_j \right) = \eta_k,$$

which completes the proof. \square

In fact, the result given above is valid for more general capacity regions \mathcal{C}_k that include the points in \mathcal{V}_k and that are bounded by the corresponding polytopes \mathcal{T}_k , i.e., $\mathcal{V}_k \subset \mathcal{C}_k \subset \mathcal{T}_k$. Note that such capacity regions are not required to be symmetric.

It is also important to observe that symmetric OPS-limited polytopes can easily be utilized to determine an upper bound (i.e., a *conservative estimate*) for the flow time whenever the original capacity regions are compact, convex and symmetric, cf. Figure 1 in Section 6.

An interesting open question is to determine sufficient conditions for the increasing sequence γ_k such that

$$\eta_1 < \dots < \eta_n.$$

We believe that the concavity of the sequence γ_k would be such a condition, but, until now, we have managed to prove the claim only for $n = 2$ and $n = 3$.

Next we show that it is possible to find a nested family of symmetric OPS-limited polytopes \mathcal{T}_k guaranteeing that the optimality of the ordinary SRPT discipline (serving always just the shortest job with fixed rate γ_1) whenever the (original) capacity regions \mathcal{C}_k are bounded by the corresponding polytopes \mathcal{T}_k . Note again that such capacity regions are not required to be symmetric.

THEOREM 4. Consider a sequence of compact capacity regions $\mathcal{C}_k \subset \mathbb{R}_+^k$, $k = 1, \dots, n$. Denote

$$\gamma_1 = \max\{c \in \mathcal{C}_1\}.$$

Let \mathcal{T}_k , $k = 1, \dots, n$, denote the symmetric OPS-limited polytopes generated by the increasing sequence γ_k defined by

$$\gamma_k = \frac{2n}{2n+1-k} \gamma_1.$$

If $(0, \dots, 0, \gamma_1) \in \mathcal{C}_k$ and $\mathcal{C}_k \subset \mathcal{T}_k$ for all k , then the optimal operating policy, for all sizes $s_1 \geq \dots \geq s_n$, is the SRPT policy $\pi^* = (\mathbf{c}_1^*, \dots, \mathbf{c}_n^*)$, where, for all k ,

$$\mathbf{c}_k^* = (0, \dots, 0, \gamma_1).$$

In this case, the minimum flow time T^{π^*} satisfies

$$T^{\pi^*} = \frac{1}{\gamma_1} \sum_{k=1}^n k s_k.$$

PROOF. 1° Assume first that $\mathcal{C}_k = \mathcal{T}_k$ for all k . By Theorem 3, it is sufficient to prove that, for all k ,

$$\eta_k = \frac{k}{\gamma_1} = \min_{j \in \{1, \dots, k\}} \left(\frac{jk}{\gamma_j} - \sum_{i=k+1-j}^{k-1} \eta_i \right).$$

Note that in this case, we certainly have

$$\eta_1 < \dots < \eta_n.$$

The result is proved by induction in $k = 1, \dots, n$. For $k = 1$, the result is trivially true since $\eta_1 = 1/\gamma_1$ by definition.

Assume now that $k \geq 2$ and the result is true for $j = 1, \dots, k-1$. We will show that the result is valid for $j = k$.

Since $k \leq n$ and $2x/(2x+1-j)$ is a decreasing function of x for $j > 1$, we have, for all $j \in \{2, \dots, k\}$,

$$\frac{\gamma_j}{\gamma_1} = \frac{2n}{2n+1-j} \leq \frac{2k}{2k+1-j}. \quad (4)$$

On the other hand, by (3), $\eta_k = k/\gamma_1$ if

$$\frac{k}{\gamma_1} \leq \frac{jk}{\gamma_j} - \sum_{i=k+1-j}^{k-1} \eta_i.$$

By applying the induction assumption, the condition reads as

$$\frac{k}{\gamma_1} \leq \frac{jk}{\gamma_j} - \frac{1}{\gamma_1} ((k+1-j) + \dots + (k-1)),$$

which is easily seen to be equivalent with condition

$$\frac{\gamma_j}{\gamma_1} \leq \frac{2k}{2k+1-j},$$

which is satisfied by (4). Thus, $\eta_k = k/\gamma_1$.

2° Assume now that $(0, \dots, 0, \gamma_1) \in \mathcal{C}_k \subset \mathcal{T}_k$ for all $k = 1, \dots, n$. Since $\mathbf{c}_k^* = (0, \dots, 0, \gamma_1)$ is the optimal operating point in \mathcal{T}_k and $\mathbf{c}_k^* \in \mathcal{C}_k \subset \mathcal{T}_k$ for all k , it must be the optimal operating point in \mathcal{C}_k , as well. \square

Note that if $\mathcal{C}_n \not\subset \mathcal{T}_n$, then applying the ordinary SRPT discipline is no longer optimal.

6. EXAMPLES AND NUMERICAL RESULTS

In this section we study a particular parametric family of capacity regions \mathcal{C}_k , for which the optimal rate vector and the minimum delay have explicit expressions. The minimum delay is compared against the results given by bounding the capacity region either by the tightest polymatroid (lower bound for the mean delay) or OPS-limited polytope (upper bound for the mean delay).

6.1 Optimum solutions

Let $\alpha > 1$, and consider the symmetric capacity regions \mathcal{C}_k defined as follows:

$$\mathcal{C}_k = \{\mathbf{c} \in \mathbb{R}_+^k : \sum_{j=1}^k c_j^\alpha \leq 1\}.$$

In the special case when $\alpha = 2$, the above regions represent the k -dimensional balls (sphere), and thus we refer to the above regions as α -balls, for short. The α -ball serves as a suitable example of a capacity region where the degree of scheduling gain can be easily parameterized between the extreme cases of a linear capacity region ($\alpha = 1$) and a hypercube ($\alpha \rightarrow \infty$).

To obtain the minimum delay T^{π^*} under the α -ball capacity region for n jobs, the optimizing values G_k^* , $k = 1, \dots, n$, need to be determined recursively by applying (1). Thus, at each stage k the following optimization problem is solved to determine the optimal rate vector $\mathbf{c}^* = (c_1^*, \dots, c_k^*) \in \mathcal{C}_k$ and the associated G_k^* ,

$$G_k^* = \min_{\mathbf{c}} g_k(\mathbf{c}) = \frac{1}{c_k} \left(k - \sum_{j=1}^{k-1} c_j G_j^* \right) \\ \text{s.t. } \sum_{j=1}^k c_j^\alpha \leq 1.$$

Recall that the optimizing values G_1^*, \dots, G_{k-1}^* are fixed constants that were determined already in the earlier stages. The above nonlinear optimization problem can be solved explicitly by an appropriate geometrical interpretation of the problem.

Let us denote $x = G_k^*$. The function $g(\mathbf{c}) = x$ represents a hyperplane with respect to \mathbf{c} ,

$$x c_k + \sum_{j=1}^{k-1} c_j G_j^* = k. \quad (5)$$

The solution to the optimization problem is given by determining the value of the unknown constant x such that the hyperplane $g(\mathbf{c}) = x$ touches the boundary of the capacity region \mathcal{C}_k given by

$$\sum_{j=1}^k c_j^\alpha = 1. \quad (6)$$

This means that at the optimal rate vector $\mathbf{c} = \mathbf{c}^*$, the outer normal vectors to the capacity boundary (6) and the hyperplane (5) must be equal up to a constant y , which gives us componentwise the following equations,

$$y c_j^{\alpha-1} = G_j^*, \quad j = 1, \dots, k-1, \\ y c_k^{\alpha-1} = x. \quad (7)$$

Thus, we have $k+2$ unknowns and $k+2$ equations, i.e., (5), (6) and (7). The solution for G_k^* , i.e., x , is readily obtained

in recursive form

$$G_k^* = \left(k^{\frac{\alpha}{\alpha-1}} - \sum_{j=1}^{k-1} G_j^* \right)^{\frac{\alpha-1}{\alpha}}.$$

The above recursive formula can be solved by reapplying the recursion to the $(k-1)$ th term on the right hand side which finally gives

$$G_k^* = \left(k^{\frac{\alpha}{\alpha-1}} - (k-1)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{\alpha-1}{\alpha}}. \quad (8)$$

The associated optimal rate vector $\mathbf{c}^* = (c_1^*, \dots, c_k^*)$ as given by (7) satisfies, for all $j = 1, \dots, k$,

$$c_j^* = k^{\frac{-1}{\alpha-1}} \left(j^{\frac{\alpha}{\alpha-1}} - (j-1)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{1}{\alpha}},$$

and the minimum flow time T^{π^*} for n jobs of sizes $s_1 \geq \dots \geq s_n$ is given by

$$T^{\pi^*} = \sum_{k=1}^n s_k \left(k^{\frac{\alpha}{\alpha-1}} - (k-1)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{\alpha-1}{\alpha}}.$$

For the polymatroid and the polytope bounds, the opportunistic gain γ_k is needed. The gain function corresponds in the capacity region to the point where the sum of the rates is maximized, i.e., one needs to solve

$$\gamma_k = \max_{\mathbf{c}} c_1 + \dots + c_k \\ \text{s.t. } \sum_{j=1}^k c_j^\alpha = 1. \quad (9)$$

Due to the symmetry of the capacity region, the optimal solution to (9) is clearly found to be $c_j = k^{-1/\alpha}$ for all $j = 1, \dots, k$, so that the gain function is

$$\gamma_k = k^{\frac{\alpha-1}{\alpha}}.$$

Note that the sequence γ_k , $k = 1, \dots, n$, is easily verified to be increasing and concave.

6.2 Numerical results

Next we give some numerical results on the delay performance in the α -ball capacity regions. We study the SRPT-HPR-type optimal policy in the actual α -ball capacity region and compare it against the lower bound given by the SRPT-HPR-type optimal policy in the tightest polymatroid capacity regions (covering the α -balls) and the upper bound given by the optimal SRPT-OPS policy in the tightest OPS-limited polytope capacity regions (inside the α -balls). These capacity regions are illustrated in Figure 1 for $n = 2$ and $\alpha = 2$, where the solid line represents the actual capacity region. The polymatroid capacity region is the outer bounding dashed line (i.e., capacity region is larger than original and hence it gives a lower bound for delay) and the inner bounding dashed line corresponds to the OPS-limited polytope capacity region (i.e., capacity region is smaller and hence it gives an upper bound on the delay). In the figure, we have also indicated (i) the rate vector $(c_1^*, c_2^*) = (\frac{1}{2}, \frac{\sqrt{3}}{2})$ associated with the optimal policy in the α -ball, (ii) the rate vector $(\gamma_2 - \gamma_1, \gamma_1) = (\sqrt{2} - 1, 1)$ associated with the optimal policy in the polymatroid capacity region, and (iii) the rate vector $(\frac{\gamma_2}{2}, \frac{\gamma_2}{2}) = (\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$ associated with the optimal policy in the OPS-limited polytope capacity region.

Note that the optimal policy in the OPS-limited polytope capacity region, i.e., SRPT-OPS, considered in Section 5

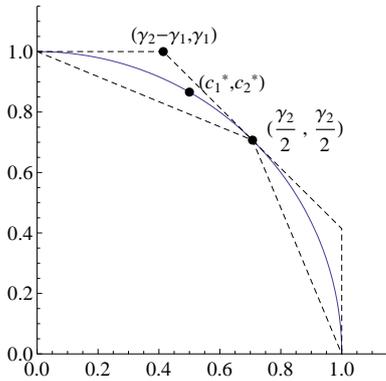


Figure 1: The capacity regions for $n = 2$ jobs with $\alpha = 2$. The actual α -ball capacity region is shown with solid line and the polymatroid and the OPS-limited polytope capacity regions are the outer and inner bounding dashed lines, respectively.

represents the size-based optimal policy when the opportunistic scheduler is only able to achieve the gain given by the PF scheduler. For this optimality, it was required that the sequence η_k is strictly increasing. This is verified in Figure 2 for the α -ball capacity region, which depicts the difference $\eta_k - \eta_{k-1}$ as a function of the number of jobs k for $\alpha = 2$ (upper panel) and $\alpha = 1.2$ (lower panel). The reason for the somewhat irregular behavior is the minimum operation required in solving (3), which makes it also difficult to analytically prove the increasing property. However, numerically we can observe that the differences remain positive (i.e., the sequence is increasing) and the figure also suggests that the differences will remain positive for any value of k .

To study the performance of the different policies, we consider the mean delay per job as a function of the number of jobs in the system. We simulated the system with random initial sizes of the jobs taken from an exponential distribution with unit mean, and the results were obtained as an average over 10^5 such realizations. Note that in the simulations there are no new arrivals and thus the only randomness comes from the random initial sizes. In the results, we additionally show the mean delay of the OPS policy representing a practical point of reference, which corresponds to the size-oblivious fair policy that serves, given n jobs in the system, all jobs in parallel at rate γ_n/n .

The results are given in Figure 3, where the upper panel corresponds to the case with $\alpha = 2$ and the lower panel to the case with $\alpha = 1.2$. The curves in each figure from bottom up correspond to (i) the polymatroid lower bound, (ii) the minimum mean delay in the α -ball capacity regions, (iii) the polytope upper bound, and (iv) the mean delay for the OPS policy. We can observe that the polymatroid lower bound becomes more loose for smaller values of α , while the polytope upper bound remains quite accurate and seems to give a good approximation to the actual optimum delay. Also, at small values of α there is less scheduling gain and OPS gives significantly poorer performance than the size-based optimum policy.

Finally, we consider the performance of the OPS policy relative to the SRPT-HPR-type optimal policy in the α -balls. Similarly as before, each simulation consisted of averaging over 10^5 realizations of random initial sizes drawn

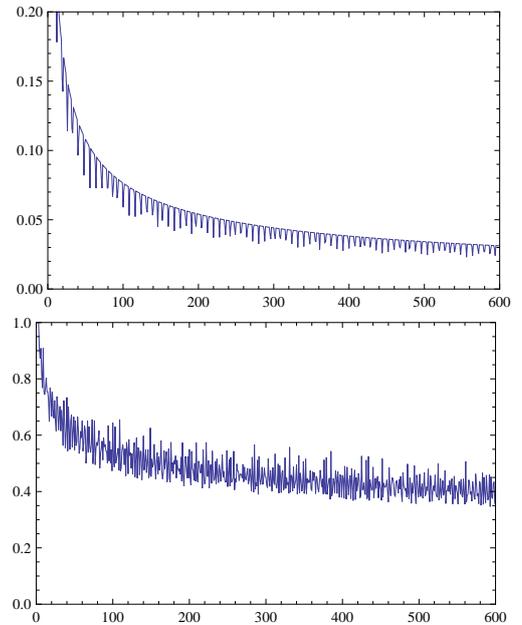


Figure 2: The differences $\eta_k - \eta_{k-1}$ in OPS-limited polytopes as a function of number of jobs for $\alpha = 2$ (upper panel) and $\alpha = 1.2$ (lower panel) remain positive, i.e., the η_k are an increasing sequence.

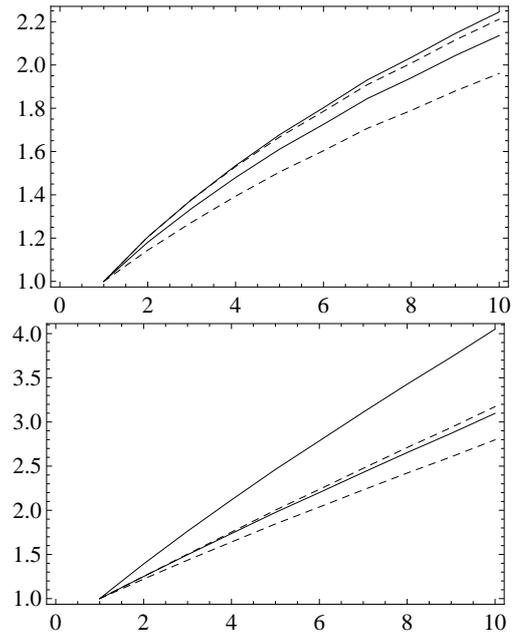


Figure 3: The mean delay as a function of number of jobs for $\alpha = 2$ (upper panel) and $\alpha = 1.2$ (lower panel). The curves from bottom up correspond to the polymatroid lower bound, the minimum mean delay for the α -balls, the polytope upper bound, and the mean delay for OPS.

from an exponential distribution with unit mean. The results are shown in Figure 4, which gives the mean delay ratio

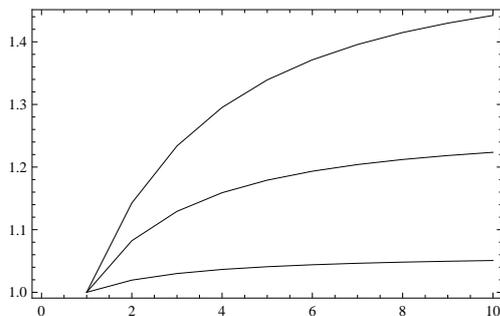


Figure 4: The mean delay ratio of the OPS policy relative to optimal policy as a function of number of jobs. The curves from bottom up correspond to $\alpha = \{2, 1.3, 1.1\}$, respectively.

of the OPS policy relative to the optimal policy in the α -balls as a function of the number of jobs. The curves from the bottom up correspond to $\alpha = \{2, 1.3, 1.1\}$, respectively. As can be seen, the gain from the optimal policy increases the smaller α is, e.g., for $\alpha = 2$ the benefit is only marginal. However, for smaller values of α the gain can be more than 40%.

7. CONCLUSIONS

We have considered the minimization of the mean flow-level delay in a transient setting for service systems where the service ability can improve as the number of jobs increases. The situation reflects the opportunistic scheduling gain observed, e.g. in modern wireless cellular networks. Our key result is that under the given conditions the SRPT principle is optimal in a sense that the shortest flow is served at the highest rate of the optimal rate vector, the second shortest at the second highest rate etc. Importantly, the optimal rate vector does not depend on the sizes of the flows. We provided a recursive algorithm to determine the optimal rate vector as well as the minimum mean delay. Also upper and lower bounds for the delay were derived by applying the main result to systems with specific polytope capacity regions. The results even allow solving the optimal scheduling problem in closed form for certain special cases and vastly facilitates numerical evaluation in the general case.

We restricted ourselves to the operating policies where the rate vector is kept unchanged until the number of jobs is decreased. However, our results give strong indication that the optimal operating policy would be the same even if we allowed *continuous* control over the rate vector. This kind of generalization might be approached by dynamic programming techniques.

While we used modern wireless cellular systems as an example of systems with opportunistic scheduling gain, where the overall service rate of the system increases with the number of users, there are many more examples as well. A similar phenomenon can be observed, e.g., in peer-to-peer file sharing systems due to the double-role of leechers as being customers and servers at the same time. Or one may consider a multi-server queue with m parallel servers and n jobs, where the total service capacity is $\gamma_n = \min\{n, m\}$. As long as there are free servers, each new job activates a new server increasing thus the whole service capacity of the system. Queueing systems with opportunistic scheduling gain

may briefly be called *scalable queues*.

In the future we plan to study the optimal scheduling problems related to various scalable queues in the dynamic setting. Even for the multi-server queue this is an open problem. It is only known that the SRPT-FM discipline minimizes the mean delay in the transient system starting with a fixed number of jobs but not allowing any (further) arrivals [9].

Another promising research direction is to consider systems where the flow level capacity regions are not explicitly defined but rather result from some adjustable packet level scheduling schemes. Combining the optimization tasks of the flow level and the packet level seems to allow solving optimal scheduling problems even in cases where determining the capacity sets would be otherwise tedious or difficult, especially at higher dimensions. An example of such a system is a weighted PF scheduler.

Acknowledgement

This research has been partially supported by the AWA (Advances in Wireless Access) project, funded by Ericsson, Nokia-Siemens Networks and TEKES.

8. REFERENCES

- [1] S. Aalto and P. Lassila. Flow-level stability and performance of channel-aware priority-based schedulers. In *Proceedings of 6th Euro-NF conference on Next Generation Internet (NGI 2010)*, pages 1–8, 2010.
- [2] F. Berggren and R. Jäntti. Asymptotically fair transmission scheduling over fading channels. *IEEE Transactions on Wireless Communications*, 3(1):326–336, Jan. 2004.
- [3] T. Bonald. A score-based opportunistic scheduler for fading radio channels. In *Proc. of European Wireless*, pages 283–292, Feb. 2004.
- [4] T. Bonald, S. Borst, N. Hegde, M. Jonckheere, and A. Proutière. Flow-level performance and capacity of wireless networks with user mobility. *Queueing Systems*, 63:131–164, December 2009.
- [5] S. Borst. User-level performance of channel-aware scheduling algorithms in wireless data networks. *IEEE/ACM Transactions on Networking*, 13(3):636–647, 2005.
- [6] S. Borst and M. Jonckheere. Flow-level stability of channel-aware scheduling algorithms. In *Proc. of WiOpt*, pages 1–6, April 2006.
- [7] M. Hu, J. Zhang, and J. Sadowsky. Traffic aided opportunistic scheduling for wireless networks: algorithms and performance bounds. *Computer Networks*, 46(4):505–518, 2004.
- [8] D. Park, H. Seo, H. Kwon, and B. Lee. Wireless packet scheduling based on the cumulative distribution function of user transmission rates. *IEEE Transactions on Communications*, 53(11):1919–1929, nov 2005.
- [9] M. Pinedo. *Scheduling: Theory, Algorithms and Systems*. Springer, third edition, 2008.
- [10] B. Sadiq. *Balancing SRPT prioritization vs opportunistic gain in wireless systems with flow dynamics*. PhD thesis, University of Texas at Austin, USA, august 2010. PhD Thesis.

- [11] B. Sadiq and G. de Veciana. Balancing srpt prioritization vs opportunistic gain in wireless systems with flow dynamics. In *Proceedings of 22nd International Teletraffic Congress (ITC-22)*, pages 1–8, 2010.
- [12] L. E. Schrage. A proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 16:687–690, 1968.
- [13] A. Stolyar. Maxweight scheduling in a generalized switch: state space collapse and equivalent workload minimization in heavy traffic. *Annals of Applied Probability*, 14(1):1–53, 2004.
- [14] A. Stolyar. On the asymptotic optimality of the gradient scheduling algorithm for multiuser throughput allocation. *Operations Research*, 53(1):12–25, 2005.
- [15] B. Tsybakov. File transmission over wireless fast fading downlink. *IEEE Transactions on Information Theory*, 48(8):2323–2337, Aug. 2002.
- [16] P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Transactions on Information Theory*, 48(6):1277–1294, jun 2002.