Flow-level Modeling and Optimization of Intercell Coordination with Dynamic TDD

Prajwal Osti
Aalto University
P.O.Box 13000
00076 Aalto
Espoo, Finland
prajwal.osti@aalto.fi

Samuli Aalto
Aalto University
P.O.Box 13000
00076 Aalto
Espoo, Finland
samuli.aalto@aalto.fi

Pasi Lassila
Aalto University
P.O.Box 13000
00076 Aalto
Espoo, Finland
pasi.lassila@aalto.fi

ABSTRACT

We study the intercell coordination problem between two interfering cells combined with dynamic time-division duplexing (TDD). In dynamic TDD, each station selects in each time slot whether it is serving uplink (u) or downlink (d) traffic. Thus, the system has four possible operation modes (uu, ud, du, dd). The amount of intercell interference between the stations clearly depends on the operation mode. We consider a flow-level model where traffic consists of elastic data flows in both cells (cells 1 and 2) and in both directions (uplink and downlink). We first characterize the maximal stability region, and then determine the optimal static (i.e., state-independent) policy. Our main objective is to analyze the potential gains from applying dynamic (i.e., state-dependent) policies, where the chosen operation mode depends on the instantaneous state of the system. To this end, motivated by certain stochastic optimality results in the literature, we define several priority policies. As a reference policy, we have the well-known max-weight policy, and we also develop another dynamic policy by applying the policy iteration algorithm. Notably we prove that certain simple priority policies are, in fact, stochastically optimal in some special cases, but which policy is optimal depends on the setting. To study the exact performance gains achieved by the dynamic policies, we perform extensive simulations. While our stochastic optimality results require exponential service times, in the simulations, we also study the impact of nonexponential service times and consider a physical model where the service time distribution is determined by the joint distribution of flow sizes and the random location of the corresponding user in the cell area. The max-weight policy is, as expected, performing well but the various priority policies are sometimes better and even optimal. Jointly the results indicate that dynamic policies give significant performance gains compared with the optimal static policy.

Q2SWinet'14, September 21–26, 2014, Montreal, QC, Canada.

http://dx.doi.org/10.1145/2642687.2642698.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design — *Wireless communication*; C.4 [**Performance of Systems**]: Modeling techniques, Performance attributes; G.3 [**Probability and Statistics**]: Queueing theory, Markov Processes; I.6.0 [**Simulation and Modeling**]: General

Keywords

Dynamic TDD; Elastic traffic; Intercell coordination; Markov Decision Processes; Stability; Stochastic optimality

1. INTRODUCTION

LTE-based cellular networks support the dynamic sharing of the spectrum between uplink and downlink traffic. This mechanism is known as *dynamic TDD* (time-division duplexing), which may be useful especially if traffic load between the uplink and downlink varies considerably over time, see [1, 2]. Additionally, the continuously increasing volume of mobile data in modern cellular networks is calling for denser deployment of the base stations, which increases the interference between base stations. This creates significant challenges for the (intercell) coordination of the base stations operating under dynamic TDD in order to minimize the interference, see [2, 3].

We study the intercell coordination problem combined with dynamic TDD for elastic data flows (such as TCP file transfers). Time-slot-level resource allocation in dynamic TDD has recently been considered, e.g., in [1, 4, 5]. However, the performance of elastic traffic manifests itself at the time scale where the number of flows is varying randomly, which is not covered by these time-slot-level models but requires flow-level models. On the other hand, flow-level models for cellular systems typically focus on only downlink traffic or uplink traffic, see, e.g., [6, 7, 8, 9, 10, 11, 12].

In [13], we studied how to model and optimize the dynamic TDD scheme for a single cell at the flow level. In the present paper, we extend this flow-level model to cover an essentially more complicated setting with two interfering base stations using dynamic TDD. Traffic consists of elastic data flows in both cells (cells 1 and 2) and in both directions (uplink and downlink), resulting in four separate traffic classes. The service requirement of an arbitrary flow can include the joint effect of a random size and random location determining the channel quality, as, e.g., in [6]. Thus, the transmission rate in a cell is not identical but depends on the location of the user downloading/uploading the flow. Moreover, it depends on the TDD operation mode used. According to the dynamic TDD scheme, each station selects whether it is serving uplink (u) or downlink (d) traffic. Thus, the system has four possible operation modes (uu, ud, du, dd).

Our model corresponds to a system of interacting queues, where the service rates of the queues are coupled and depend on the chosen operating policy. For such systems even stability is not easy to characterize, and, indeed, most of the available literature focuses on the stability properties, see [12, 14, 15]. Results on optimal policies are significantly more scarce. In a related discrete-time model with arbitrarily coupled service rates between the queues, the max-weight scheduler is known to be maximally stable [16] and even has certain asymptotical optimality properties [17]. Recently, for a model with only two classes representing, e.g., the downlink flows and thus excluding the dynamic TDD aspect, some interesting structural optimality properties were given in [18].

In the present paper, we first derive an explicit expression for the maximal stability condition, and then formulate the optimization problem to determine the optimal static policy that minimizes the average flow-level delay. A static policy chooses the operation mode in each time slot independently and randomly from a given distribution (and thus independently of the state of the system). We determine the optimal static policy in various scenarios. To further optimize the performance, we consider dynamic policies, where the chosen operation mode depends on the instantaneous state of the system. In the dynamic setting, inspired by the stochastic optimality results for the both stations on priority policy in [18], we define several priority policies that take into account the dynamic TDD setting. Additionally, we derive the dynamic policy called FPI resulting from applying the policy iteration algorithm from the MDP theory to the optimal static policy [19]. As a reference policy, we have the max-weight policy. We consider the stability properties of dynamic policies and argue that some of the priority policies potentially suffer from instability. Our main theoretical results show, however, that in certain special cases a simple priority policy can even be stochastically optimal.

To analyze the performance gains from the dynamic policies, we rely on numerical methods. The dynamic policies are simulated extensively in order to obtain insight into how the policies are performing relative to each other and what the optimal performance could be. As expected, the max-weight policy is robust, performing overall well. Also, the FPI policy behaves in a robust way, yielding performance that is close to that of the max-weight policy. However, by using the priority policies, we are able to obtain even better results. Namely, for a given scenario and traffic load, some of the priority policies typically yield better results than the max-weight and FPI policies. Among these, the best result then gives a feasible estimate of the optimal performance. In addition, in certain special cases we even know the optimal dynamic policy, and the simulated results give a very accurate estimate of the optimal performance. Jointly all the results give us an indication that in our scenarios the maximal gain from the optimal dynamic coordination compared with the optimal static one is considerable (approximately 50%–60%).

The rest of the paper is organized as follows. In Section 2, we describe the basic model of the problem we consider here. Then, in Section 3, we introduce the static policies and characterize the maximal stability region in the general case. The optimal static policy is considered in Section 4. This is followed by Section 5, in which we introduce the dynamic policies that are used in this paper. Then we discuss the stability issues of different dynamic policies in Section 6. In Section 7, we give the main theoretical results of the paper concerning the stochastic optimality of certain dynamic policies. This is followed by the numerical study of various dynamic policies in Section 8 with the help of simulation, and finally conclusions of the study are drawn in Section 9.

2. MODEL

We consider the problem of optimally coordinating two neighboring base stations with identical properties which operate in a dynamic TDD configuration and are affected by each other's operations. We take two adjacent cells with base stations B1 and B2, which serve elastic flows in both uplink (u) and downlink (d) directions. The flows arrive according to an independent Poisson process to each *class* (identified by the index of the station and the direction of transmission) with intensity λ_i^δ, where $\delta \in \{\mathrm{u}, \mathrm{d}\}$ denotes the direction of transfer while $i \in \{1, 2\}$ indicates the index of the base station. The size of an arriving flow has a general distribution depending on the direction δ of the flow. The sizes are assumed to be independent of each other. The location of the flow within its cell is also random and independent of the other flows.

Each station can decide to operate on u-direction or d-direction flows in each time slot. Thus, there are four different operation modes (uu, ud, du, dd). In the 'dd' mode, both stations decide to transmit in the downlink direction. Let S_d^dd denote the service time of a downlink flow in the 'dd' mode, i.e., the time needed to transfer all bits of the flow in the 'dd' mode when there is just one downlink flow in a random position within the cell. We assume that the cells are symmetric in the sense that S_d^dd represents the generic service time of a downlink flow for both cells with the mean value denoted by $\mathrm{E}[S_\mathrm{d}^\mathrm{dd}] = 1/\mu_\mathrm{d}^\mathrm{dd}$.

In the 'du' mode, B1 decides to transmit in the downlink direction and B2 in the uplink direction. Let S_d^du denote the service time of a downlink flow in the 'du' mode. Correspondingly, let S_d^ud denote the service time of a downlink flow in the 'ud' mode, where B1 decides to transmit in the uplink direction and B2 in the downlink direction. We assume that the cells are symmetric in the sense that $\mathrm{E}[S_\mathrm{d}^\mathrm{du}] = \mathrm{E}[S_\mathrm{d}^\mathrm{ud}] = 1/\mu_\mathrm{d}^\mathrm{du}$. In addition, we assume that in the 'du' and 'ud' modes, the uplink flows are not served at all because the downlink transmission power is considerably greater than the uplink power leading to heavy interference for the uplink [2].

Finally, let S_u^uu denote the service time of an uplink flow in the 'uu' mode, where both stations decide to transmit in the uplink direction. Because of the low uplink power, we assume that the transmissions do not interfere. In addition, we again assume that the cells are symmetric in the sense that S_u^uu represents the service time of an uplink flow for both cells with the mean value denoted by $\mathrm{E}[S_\mathrm{u}^\mathrm{uu}] = 1/\mu_\mathrm{u}^\mathrm{uu}$.

Note that the service time of a flow depends on its random size, its random location, and the operation mode used. The latter two together determine the transmission rate for the user. Thus, for example, $S_\mathrm{d}^\mathrm{dd} = Y^\mathrm{d}/c^\mathrm{dd}(Z)$, where Y^d refers to the size of the downlink flow (in bits), Z its distance from the base station, and $c^{\mathrm{dd}}(z)$ the transmission rate (in bits/s) at distance z in the 'dd' mode.

Since the intercell interference is much higher in the 'dd' mode (compared to the 'du' and 'ud' modes), we may assume that

$$\mu_{\mathrm{d}}^{\mathrm{du}} > \mu_{\mathrm{d}}^{\mathrm{dd}}. \tag{1}$$

Throughout this paper, we focus on the *low downlink interference* case which is defined by further requiring that

$$2\mu_{\mathrm{d}}^{\mathrm{dd}} > \mu_{\mathrm{d}}^{\mathrm{du}}. \tag{2}$$

The high downlink interference case, where $2\mu_{\mathrm{d}}^{\mathrm{dd}} \le \mu_{\mathrm{d}}^{\mathrm{du}}$, is also of interest and apparently leads to qualitatively different results. However, due to lack of space, it is not considered in this paper.

In each class, the flows are assumed to be served in a round robin manner at the time-slot level. With the additional assumption of time-scale separation [6, 7, 8, 9, 10, 11, 12] between time-slot and flow levels, we conclude that at the flow level each class is served according to the processor sharing (PS) service discipline.

3. STATIC POLICIES AND THE MAXIMAL STABILITY REGION

As mentioned in Section 2, the system has four operational modes (uu, ud, du, dd). In *static* sharing, these four modes are allocated some fixed time shares without taking into account the underlying state of the system. In *dynamic* sharing, instead, the modes are allocated dynamically depending on the state of the system. In this section, we consider static (sharing) policies and derive the maximal stability region. Dynamic policies are introduced then in the following section.

A *static policy* selects the mode randomly and independently (in each time slot) from a given probability distribution $p = (p^{\mathrm{uu}}, p^{\mathrm{ud}}, p^{\mathrm{du}}, p^{\mathrm{dd}})$. Elements p^m thus refer to the fraction of time the two stations are operated in the corresponding mode m.

According to our assumptions given in Section 2, static sharing results in four independent M/G/1-PS queues with service completion rates

$$\begin{aligned} \bar\mu_1^{\mathrm{u}} = \bar\mu_2^{\mathrm{u}} = p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}}, \quad \bar\mu_1^{\mathrm{d}} &= p^{\mathrm{du}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}}, \\ \bar\mu_2^{\mathrm{d}} &= p^{\mathrm{ud}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}}. \end{aligned} \tag{3}$$

A static policy p is *stable* if

$$\begin{aligned} \max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\} < p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}}, \quad \lambda_1^{\mathrm{d}} &< p^{\mathrm{du}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}}, \\ \lambda_2^{\mathrm{d}} &< p^{\mathrm{ud}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}}. \end{aligned} \tag{4}$$

Let X_i^δ denote the number of flows in station i and direction δ. For a stable policy, the mean total number of flows in the whole system is $\bar X = \bar X_1^{\mathrm{u}} + \bar X_1^{\mathrm{d}} + \bar X_2^{\mathrm{u}} + \bar X_2^{\mathrm{d}}$, where

$$\bar X_i^\delta = \frac{\lambda_i^\delta}{\bar\mu_i^\delta - \lambda_i^\delta}, \quad \delta \in \{\mathrm{u},\mathrm{d}\},\ i \in \{1,2\}.$$

By Little's result, the mean flow-level delay (i.e., the time needed to complete the transmission) is given by $\bar T = \bar X/\lambda$, where $\lambda = \lambda_1^{\mathrm{u}} + \lambda_2^{\mathrm{u}} + \lambda_1^{\mathrm{d}} + \lambda_2^{\mathrm{d}}$ denotes the total arrival rate.

Our first result characterizes the maximal stability region.

Theorem 1. *There exists a stable static policy if and only if*

$$\frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\min\{\lambda_1^{\mathrm{d}}, \lambda_2^{\mathrm{d}}\}}{\mu_{\mathrm{d}}^{\mathrm{dd}}} + \frac{|\lambda_1^{\mathrm{d}} - \lambda_2^{\mathrm{d}}|}{\mu_{\mathrm{d}}^{\mathrm{du}}} < 1. \tag{5}$$

Proof: Without loss of generality, we assume that $\lambda_1^{\mathrm{d}} \ge \lambda_2^{\mathrm{d}}$.

1° Assume first that condition (5) is satisfied, i.e.,

$$\frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{dd}}} + \frac{\lambda_1^{\mathrm{d}} - \lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} < 1,$$

Consider the static policy defined by

$$\begin{aligned} p^{\mathrm{uu}} &= \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\epsilon}{3}, \quad p^{\mathrm{dd}} = \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{dd}}} + \frac{\epsilon}{3}, \\ p^{\mathrm{du}} &= \frac{\lambda_1^{\mathrm{d}} - \lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} + \frac{\epsilon}{3}, \quad p^{\mathrm{ud}} = 0, \end{aligned}$$

where

$$\epsilon = 1 - \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} - \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{dd}}} - \frac{\lambda_1^{\mathrm{d}} - \lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} > 0.$$

It is easy to see that all conditions (4) are satisfied so that the policy is well-defined and stable.

2° Assume now that condition (5) is *not* satisfied, i.e.,

$$\frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{dd}}} + \frac{\lambda_1^{\mathrm{d}} - \lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} \ge 1, \tag{6}$$

Consider first the static policies for which

$$p^{\mathrm{dd}} \le \lambda_2^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{dd}}. \tag{7}$$

Now

$$\begin{aligned} & p^{\mathrm{uu}} + p^{\mathrm{ud}} + p^{\mathrm{du}} + p^{\mathrm{dd}} \\ & \overset{(4)}{>} \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\lambda_1^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} + \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} - p^{\mathrm{dd}}\frac{2\mu_{\mathrm{d}}^{\mathrm{dd}} - \mu_{\mathrm{d}}^{\mathrm{du}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} \\ & \overset{(2),(7)}{\ge} \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\lambda_1^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} + \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} - \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{dd}}}\frac{2\mu_{\mathrm{d}}^{\mathrm{dd}} - \mu_{\mathrm{d}}^{\mathrm{du}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} \overset{(6)}{\ge} 1, \end{aligned}$$

which contradicts with the requirement that $p = (p^{\mathrm{uu}}, p^{\mathrm{ud}}, p^{\mathrm{du}}, p^{\mathrm{dd}})$ is a probability distribution. Thus, there is no stable static policy satisfying (7).

Consider then the static policies for which

$$p^{\mathrm{dd}} > \lambda_2^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{dd}}. \tag{8}$$

Now

$$\begin{aligned} & p^{\mathrm{uu}} + p^{\mathrm{du}} + p^{\mathrm{dd}} \\ & \overset{(4)}{>} \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\lambda_1^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} + p^{\mathrm{dd}}\frac{\mu_{\mathrm{d}}^{\mathrm{du}} - \mu_{\mathrm{d}}^{\mathrm{dd}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} \\ & \overset{(1),(8)}{>} \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}} + \frac{\lambda_1^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} + \frac{\lambda_2^{\mathrm{d}}}{\mu_{\mathrm{d}}^{\mathrm{dd}}}\frac{\mu_{\mathrm{d}}^{\mathrm{du}} - \mu_{\mathrm{d}}^{\mathrm{dd}}}{\mu_{\mathrm{d}}^{\mathrm{du}}} \overset{(6)}{\ge} 1, \end{aligned}$$

which again contradicts with the requirement that p is a probability distribution. Thus, there is neither any stable static policy that satisfies (8). □

Based on Theorem 1, it is natural to define

$$\rho^{\mathrm{u}} = \frac{\max\{\lambda_1^{\mathrm{u}}, \lambda_2^{\mathrm{u}}\}}{\mu_{\mathrm{u}}^{\mathrm{uu}}}, \quad \rho^{\mathrm{d}} = \frac{\min\{\lambda_1^{\mathrm{d}}, \lambda_2^{\mathrm{d}}\}}{\mu_{\mathrm{d}}^{\mathrm{dd}}} + \frac{|\lambda_1^{\mathrm{d}} - \lambda_2^{\mathrm{d}}|}{\mu_{\mathrm{d}}^{\mathrm{du}}} \tag{9}$$

so that the maximal stability condition reads as $\rho^{\mathrm{u}} + \rho^{\mathrm{d}} < 1$.

4. OPTIMAL STATIC POLICY

In this section, we consider the optimal static sharing problem so that the mean total number of flows (or, equivalently by Little's result, the mean flow-level delay) is minimized. Clearly, the optimal static policy exists whenever the system satisfies the maximal stability condition (5) and is obtained as the solution of the following optimization problem:

$$\begin{aligned}\bar{X}^* = \quad & \min_{p} \left\{ \frac{\lambda_1^{\mathrm{u}}}{p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda_1^{\mathrm{u}}} + \frac{\lambda_1^{\mathrm{d}}}{p^{\mathrm{du}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda_1^{\mathrm{d}}} \right. \\ & \left. + \frac{\lambda_2^{\mathrm{u}}}{p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda_2^{\mathrm{u}}} + \frac{\lambda_2^{\mathrm{d}}}{p^{\mathrm{ud}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda_2^{\mathrm{d}}} \right\}, \\ \text{s.t.} \quad & p^{\mathrm{uu}} \geq 0,\ p^{\mathrm{ud}} \geq 0,\ p^{\mathrm{du}} \geq 0,\ p^{\mathrm{dd}} \geq 0, \\ & p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda_1^{\mathrm{u}} > 0,\ p^{\mathrm{du}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda_1^{\mathrm{d}} > 0, \\ & p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda_2^{\mathrm{u}} > 0,\ p^{\mathrm{ud}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda_2^{\mathrm{d}} > 0, \\ & p^{\mathrm{uu}} + p^{\mathrm{ud}} + p^{\mathrm{du}} + p^{\mathrm{dd}} = 1. \end{aligned} \tag{10}$$

Due to the complexity of the optimization problem, we restrict ourselves to a number of scenarios that are later on illustrated by numerical examples. These scenarios are defined utilizing two free parameters, λ^{u} and λ^{d}. In *Scenario 1* (symmetric case), the uplink (and downlink) arrival rates are equal in both the stations, i.e., $\lambda_1^{\mathrm{u}} = \lambda_2^{\mathrm{u}} = \lambda^{\mathrm{u}}$ and $\lambda_1^{\mathrm{d}} = \lambda_2^{\mathrm{d}} = \lambda^{\mathrm{d}}$. In *Scenario 2* (uplink asymmetric), the downlink arrival rates are the same in both the stations while the uplink flows arrive only in B1, i.e., $\lambda_1^{\mathrm{d}} = \lambda_2^{\mathrm{d}} = \lambda^{\mathrm{d}}$, $\lambda_1^{\mathrm{u}} = \lambda^{\mathrm{u}}$, and $\lambda_2^{\mathrm{u}} = 0$. In *Scenario 3* (downlink asymmetric), the uplink arrival rates are equal in both the stations but there are no downlink flows arriving in B2, i.e., $\lambda_1^{\mathrm{u}} = \lambda_2^{\mathrm{u}} = \lambda^{\mathrm{u}}$, $\lambda_1^{\mathrm{d}} = \lambda^{\mathrm{d}}$, and $\lambda_2^{\mathrm{d}} = 0$. Finally, in *Scenario 4* (station asymmetric), the uplink (and downlink) arrival rates in station B1 are twice as large as those in station B2, i.e., $\lambda_1^{\mathrm{u}} = 2\lambda_2^{\mathrm{u}} = 2\lambda^{\mathrm{u}}$ and $\lambda_1^{\mathrm{d}} = 2\lambda_2^{\mathrm{d}} = 2\lambda^{\mathrm{d}}$. The four scenarios are summarized in Table 1.

Table 1: Summary of scenarios considered

Scenario	λ_1^{u}	λ_1^{d}	λ_2^{u}	λ_2^{d}
1	λ^{u}	λ^{d}	λ^{u}	λ^{d}
2	λ^{u}	λ^{d}	0	λ^{d}
3	λ^{u}	λ^{d}	λ^{u}	0
4	$2\lambda^{\mathrm{u}}$	$2\lambda^{\mathrm{d}}$	λ^{u}	λ^{d}

4.1 Symmetric arrivals

In Scenario 1, where $\lambda_1^{\mathrm{u}} = \lambda_2^{\mathrm{u}} = \lambda^{\mathrm{u}}$ and $\lambda_1^{\mathrm{d}} = \lambda_2^{\mathrm{d}} = \lambda^{\mathrm{d}}$, it is clearly sufficient to consider the policies for which $p^{\mathrm{du}} = p^{\mathrm{ud}}$ so that the optimization problem reads as follows:

$$\begin{aligned}\bar{X}^* \quad & = \min_{p} \left\{ \frac{2\lambda^{\mathrm{u}}}{p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda^{\mathrm{u}}} + \frac{2\lambda^{\mathrm{d}}}{p^{\mathrm{du}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda^{\mathrm{d}}} \right\}, \\ \text{s.t.} \quad & p^{\mathrm{uu}} \geq 0,\ p^{\mathrm{du}} = p^{\mathrm{ud}} \geq 0,\ p^{\mathrm{dd}} \geq 0, \\ & p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda^{\mathrm{u}} > 0,\ p^{\mathrm{du}}\mu_{\mathrm{d}}^{\mathrm{du}} + p^{\mathrm{dd}}\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda^{\mathrm{d}} > 0, \\ & p^{\mathrm{uu}} + 2p^{\mathrm{du}} + p^{\mathrm{dd}} = 1. \end{aligned}$$

To get the minimum value of $\bar{X}$, we substitute $p^{\mathrm{dd}} = 1 - 2p^{\mathrm{du}} - p^{\mathrm{uu}}$ and differentiate $\bar{X}$ with respect to p^{du} and p^{uu}, from which we observe that, for any fixed p^{uu}, $\bar{X}$ is an increasing function of p^{du} as $\frac{\partial \bar{X}}{\partial p^{\mathrm{du}}} > 0$ under the low interference condition (2). This means that the minimum is achieved at $p^{\mathrm{du}} = 0$. So the optimization problem now becomes

$$\begin{aligned}\bar{X}^* = \quad & \min_{p^{\mathrm{uu}}} \frac{2\lambda^{\mathrm{u}}}{p^{\mathrm{uu}}\mu_{\mathrm{u}}^{\mathrm{uu}} - \lambda^{\mathrm{u}}} + \frac{2\lambda^{\mathrm{d}}}{(1 - p^{\mathrm{uu}})\mu_{\mathrm{d}}^{\mathrm{dd}} - \lambda^{\mathrm{d}}}, \\ \text{s.t.} \quad & \rho^{\mathrm{u}} < p^{\mathrm{uu}} < 1 - \rho^{\mathrm{d}}, \end{aligned}$$

where $\rho^{\mathrm{u}} = \lambda^{\mathrm{u}}/\mu_{\mathrm{u}}^{\mathrm{uu}}$ and $\rho^{\mathrm{d}} = \lambda^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{dd}}$. To calculate the optimal value of p^{uu}, we set $\frac{\partial \bar{X}}{\partial p^{\mathrm{uu}}}$ to zero taking into account that $p^{\mathrm{du}} = 0$ for the optimal share, which gives

$$(p^{\mathrm{uu}})^* = \frac{(1-\rho^{\mathrm{d}})\sqrt{\rho^{\mathrm{u}}} + \rho^{\mathrm{u}}\sqrt{\rho^{\mathrm{d}}}}{\sqrt{\rho^{\mathrm{u}}} + \sqrt{\rho^{\mathrm{d}}}}, \quad (p^{\mathrm{dd}})^* = 1 - (p^{\mathrm{uu}})^*,$$

$$\bar{X}^* = \frac{2(\sqrt{\rho^{\mathrm{u}}} + \sqrt{\rho^{\mathrm{d}}})^2}{1 - \rho^{\mathrm{u}} - \rho^{\mathrm{d}}}.$$

4.2 Asymmetric arrivals

We now consider the first two cases with asymmetric arrivals. In Scenario 2 with uplink asymmetry, the uplink flows arrive only in B1 and never in B2, i.e., $\lambda_2^{\mathrm{u}} = 0$. In this case,

$$(p^{\mathrm{uu}})^* = \frac{(1-\rho^{\mathrm{d}})\sqrt{\rho^{\mathrm{u}}} + \rho^{\mathrm{u}}\sqrt{2\rho^{\mathrm{d}}}}{\sqrt{\rho^{\mathrm{u}}} + \sqrt{2\rho^{\mathrm{d}}}}, \quad (p^{\mathrm{dd}})^* = 1 - (p^{\mathrm{uu}})^*,$$

$$\bar{X}^* = \frac{(\sqrt{\rho^{\mathrm{u}}} + \sqrt{2\rho^{\mathrm{d}}})^2}{1 - \rho^{\mathrm{u}} - \rho^{\mathrm{d}}}.$$

where $\rho^{\mathrm{u}} = \lambda^{\mathrm{u}}/\mu_{\mathrm{u}}^{\mathrm{uu}}$ and $\rho^{\mathrm{d}} = \lambda^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{dd}}$.

On the other hand, in Scenario 3 with downlink asymmetry, only B1 gets the downlink requests while B2 does not get any downlink service request, i.e., $\lambda_2^{\mathrm{d}} = 0$. It can be shown that

$$(p^{\mathrm{uu}})^* = \frac{(1-\rho^{\mathrm{d}})\sqrt{2\rho^{\mathrm{u}}} + \rho^{\mathrm{u}}\sqrt{\rho^{\mathrm{d}}}}{\sqrt{2\rho^{\mathrm{u}}} + \sqrt{\rho^{\mathrm{d}}}}, \quad (p^{\mathrm{du}})^* = 1 - (p^{\mathrm{uu}})^*,$$

$$\bar{X}^* = \frac{(\sqrt{2\rho^{\mathrm{u}}} + \sqrt{\rho^{\mathrm{d}}})^2}{1 - \rho^{\mathrm{u}} - \rho^{\mathrm{d}}}.$$

where $\rho^{\mathrm{u}} = \lambda^{\mathrm{u}}/\mu_{\mathrm{u}}^{\mathrm{uu}}$ and $\rho^{\mathrm{d}} = \lambda^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{du}}$.

For Scenario 4, where flows arrive in all four queues, there are, however, no explicit expressions for the optimal time shares nor the average total number of flows. Therefore we have to solve the optimization problem numerically in this case.

5. DYNAMIC POLICIES

In a dynamic setting, the stations operate in one of the four modes (uu, ud, du, dd) depending on the state of the system. Here, the state of the system refers to the vector of the number of flows in each queue, $x = (x_1^{\mathrm{u}}, x_1^{\mathrm{d}}, x_2^{\mathrm{u}}, x_2^{\mathrm{d}}) \in \mathcal{X}$. Then an action $a^{\pi}(x) \in \{\mathrm{uu}, \mathrm{ud}, \mathrm{du}, \mathrm{dd}\}$ is chosen for each state x defined by the policy π. Such dynamic policies are discussed below.

5.1 Uplink priority heuristic (HU)

In this policy, we serve both the uplink queues with the highest priority and serve the downlink only if both the uplink queues are empty. When both the uplink queues are empty, the 'dd' mode is used if there are flows in both the downlink queues, and either the 'du' or 'ud' mode is used when one of the downlink queues is empty. So, we have

$$a^{\mathrm{HU}}(x) = \begin{cases} \mathrm{uu}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} > 0, \\ \mathrm{dd}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} = 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} > 0, \\ \mathrm{du}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} = 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, x_1^{\mathrm{d}} > 0, \\ \mathrm{ud}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} = 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, x_2^{\mathrm{d}} > 0. \end{cases} \tag{11}$$

Clearly, it is a good heuristic to use when the uplink flow completion rate, $\mu_{\mathrm{u}}^{\mathrm{uu}}$, is high compared to the downlink flow completion rates, $\mu_{\mathrm{d}}^{\mathrm{du}}$ and $\mu_{\mathrm{d}}^{\mathrm{dd}}$, and the traffic load is sufficiently small. Indeed we will later observe (in Theorem 2) that HU is even stochastically optimal in certain cases with large enough $\mu_{\mathrm{u}}^{\mathrm{uu}}$. However, due to the high priority given to the uplink, it may run into stability problems at high load as we will witness in Sections 6 and 8.

5.2 Downlink priority heuristic (HD)

In this policy, the downlink is given the highest priority, and the uplink ('uu') mode is chosen only when both the downlink queues are empty. If there are flows in both the downlink queues, then the 'dd' mode is used. When the downlink queue for one station is empty, the downlink queue for the other is served at a higher rate (the 'du' or 'ud' mode). Thus,

$$a^{\mathrm{HD}}(x) = \begin{cases} \mathrm{dd}, & \mathrm{if} x_1^{\mathrm{d}} x_2^{\mathrm{d}} > 0, \\ \mathrm{du}, & \mathrm{if} x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, x_1^{\mathrm{d}} > 0, \\ \mathrm{ud}, & \mathrm{if} x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, x_2^{\mathrm{d}} > 0, \\ \mathrm{uu}, & \mathrm{if} x_1^{\mathrm{d}} + x_2^{\mathrm{d}} = 0. \end{cases} \tag{12}$$

It is reasonable to assume that this policy will give good performance when the downlink flow completion rates, $\mu_{\mathrm{d}}^{\mathrm{du}}$ and $\mu_{\mathrm{d}}^{\mathrm{dd}}$, are high compared to the uplink flow completion rate, $\mu_{\mathrm{u}}^{\mathrm{uu}}$. Under some conditions, HD even turns out to be stochastically optimal (see Theorem 3).

5.3 Modified uplink priority heuristics (H3, H4)

The uplink priority policy, HU, may become unstable at high loads (see Sections 6 and 8). To mitigate this problem, we attempt to exploit the high aggregate downlink transmission rate as much as possible (by using the 'dd' mode) and modify the HU policy to get the H3 policy. In H3, both the uplink queues are served (the 'uu' mode) when there are flows in both the queues. However, unlike the HU policy, it serves both the downlink queues (the 'dd' mode) if there are flows in both of them while at least one of the uplink queues is empty. Otherwise, it is similar to HU. Thus,

$$a^{\mathrm{H3}}(x) = \begin{cases} \mathrm{uu}, & \mathrm{if} x_1^{\mathrm{u}} x_2^{\mathrm{u}} > 0, \\ \mathrm{dd}, & \mathrm{if} x_1^{\mathrm{u}} x_2^{\mathrm{u}} = 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} > 0, \\ \mathrm{uu}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} > 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, \\ \mathrm{du}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} = 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, x_1^{\mathrm{d}} > 0, \\ \mathrm{ud}, & \mathrm{if} x_1^{\mathrm{u}} + x_2^{\mathrm{u}} = 0, x_1^{\mathrm{d}} x_2^{\mathrm{d}} = 0, x_2^{\mathrm{d}} > 0. \end{cases} \tag{13}$$

In Scenario 3 (asymmetric downlink), where $\lambda_2^{\mathrm{d}} = 0$, we should apparently not apply the 'dd' mode at all. We still serve the uplink queues (the 'uu' mode) whenever *both* of them are non-empty. But if at least one of the uplink queues is empty, then we choose the 'du' mode (instead of 'dd'). Otherwise, this modified policy, H4, is similar to H3 and HU. So, we have

$$a^{\mathrm{H4}}(x) = \begin{cases} \mathrm{uu}, & \mathrm{if} x_1^{\mathrm{u}} x_2^{\mathrm{u}} > 0, \\ \mathrm{du}, & \mathrm{if} x_1^{\mathrm{u}} x_2^{\mathrm{u}} = 0,\ x_1^{\mathrm{d}} > 0, \\ \mathrm{uu}, & \mathrm{if} x_1^{\mathrm{u}} x_2^{\mathrm{u}} = 0,\ x_1^{\mathrm{d}} = 0. \end{cases} \tag{14}$$

5.4 Max-weight policy (MW)

In this policy, the mode that has the *maximum weight* for a state is chosen. The weight is defined as the sum of the product of the service completion rate of the queues in the chosen mode and the queue lengths of the same queues. Thus, in the MW policy, the action chosen in state x is defined by

$$a^{\mathrm{MW}}(x) = \arg\max \Big\{ \mu_{\mathrm{u}}^{\mathrm{uu}} (x_1^{\mathrm{u}} + x_2^{\mathrm{u}}), \mu_{\mathrm{d}}^{\mathrm{du}} x_2^{\mathrm{d}}, \mu_{\mathrm{d}}^{\mathrm{du}} x_1^{\mathrm{d}}, \mu_{\mathrm{d}}^{\mathrm{dd}} (x_1^{\mathrm{d}} + x_2^{\mathrm{d}}) \Big\}. \tag{15}$$

This policy is maximally stable [16] and even asymptotically optimal [17] for discrete-time queues.

5.5 First Policy Iteration (FPI)

In this approach, a baseline policy is improved by the method of *policy iteration* which follows from the theory of Markov decision processes [19]. In the policy iteration algorithm, the initial policy π is used to construct a new policy π', for which the action, $a^{\pi'}(x)$, in state x is given by

$$a^{\pi'}(x) = \arg\min_{a \in \mathcal{A}} \Big\{ r(x,a) - \bar{r}^{\pi} + \sum_{y \neq x} q_{xy}(a) (v^{\pi}(y) - v^{\pi}(x)) \Big\}, \tag{16}$$

where $r(x,a)$ is the instantaneous cost rate when action a is chosen, $\bar{r}^{\pi}$ is the average cost rate when the initial policy π is used, $q_{xy}(a)$ is the transition rate from state x to state y when the action a is chosen, and $v^{\pi}(x)$ is the relative value of state under the initial policy. To apply this approach, we have to assume that the service times in each mode are exponential.

We take the optimal static policy as our initial policy π and derive an improved dynamic policy π' (FPI) by applying the policy iteration algorithm once. The state is given by $x = (x_1^{\mathrm{u}}, x_1^{\mathrm{d}}, x_2^{\mathrm{u}}, x_2^{\mathrm{d}})$, and the action space consists of the four operation modes, $\mathcal{A} = \{\mathrm{uu}, \mathrm{ud}, \mathrm{du}, \mathrm{dd}\}$. The instantaneous cost rate $r(x,a)$ is equal to the total number of flows in the whole system, $x_1^{\mathrm{u}} + x_1^{\mathrm{d}} + x_2^{\mathrm{u}} + x_2^{\mathrm{d}}$, and is thus independent of the chosen action a. As we have parallel PS queues, the relative value of state of the system is the sum of the relative values of each PS queue. The relative value of state of an M/M/1-PS queue with service rate μ and arrival rate λ when there are n customers present is (see, e.g., [20])

$$v(n) = v(0) + \frac{n(n+1)}{2(\mu - \lambda)}. \tag{17}$$

Now, the quantity inside the brackets in (16) becomes

$$\begin{aligned} & x_1^{\mathrm{u}} + x_1^{\mathrm{d}} + x_2^{\mathrm{u}} + x_2^{\mathrm{d}} - \bar{r}^{\pi} + \mu^{\pi}(x,a) \\ & + \lambda_1^{\mathrm{u}} [v^{\pi}(x + e_1^{\mathrm{u}}) - v^{\pi}(x)] + \lambda_1^{\mathrm{d}} [v^{\pi}(x + e_1^{\mathrm{d}}) - v^{\pi}(x)] \\ & + \lambda_2^{\mathrm{u}} [v^{\pi}(x + e_2^{\mathrm{u}}) - v^{\pi}(x)] + \lambda_2^{\mathrm{d}} [v^{\pi}(x + e_2^{\mathrm{d}}) - v^{\pi}(x)], \end{aligned}$$

where the unit vectors e_i^δ are defined as

$$e_1^{\mathrm{u}} = (1,0,0,0), \quad e_1^{\mathrm{d}} = (0,1,0,0), \quad e_2^{\mathrm{u}} = (0,0,1,0),$$
$$e_2^{\mathrm{d}} = (0,0,0,1),$$

and $\mu^\pi(x,a)$ is the action-specific part of future cost associated with action a. By utilizing (17), we get

$$\begin{aligned}
\mu^\pi(x,\mathrm{uu}) &= \mu_{\mathrm{u}}^{\mathrm{uu}}(v^\pi(x-e_1^{\mathrm{u}}) + v^\pi(x-e_2^{\mathrm{u}}) - 2v^\pi(x)) \\
&= -\frac{\mu_{\mathrm{u}}^{\mathrm{uu}} x_1^{\mathrm{u}}}{\bar{\mu}_1^{\mathrm{u}} - \lambda_1^{\mathrm{u}}} - \frac{\mu_{\mathrm{u}}^{\mathrm{uu}} x_2^{\mathrm{u}}}{\bar{\mu}_2^{\mathrm{u}} - \lambda_2^{\mathrm{u}}}, \\
\mu^\pi(x,\mathrm{ud}) &= \mu_{\mathrm{d}}^{\mathrm{du}}(v^\pi(x-e_2^{\mathrm{d}}) - v^\pi(x)) = -\frac{\mu_{\mathrm{d}}^{\mathrm{du}} x_2^{\mathrm{d}}}{\bar{\mu}_2^{\mathrm{d}} - \lambda_2^{\mathrm{d}}}, \\
\mu^\pi(x,\mathrm{du}) &= \mu_{\mathrm{d}}^{\mathrm{du}}(v^\pi(x-e_1^{\mathrm{d}}) - v^\pi(x)) = -\frac{\mu_{\mathrm{d}}^{\mathrm{du}} x_1^{\mathrm{d}}}{\bar{\mu}_1^{\mathrm{d}} - \lambda_1^{\mathrm{d}}}, \\
\mu^\pi(x,\mathrm{dd}) &= \mu_{\mathrm{d}}^{\mathrm{dd}}(v^\pi(x-e_1^{\mathrm{d}}) + v^\pi(x-e_2^{\mathrm{d}}) - 2v^\pi(x)) \\
&= -\frac{\mu_{\mathrm{d}}^{\mathrm{dd}} x_1^{\mathrm{d}}}{\bar{\mu}_1^{\mathrm{d}} - \lambda_1^{\mathrm{d}}} - \frac{\mu_{\mathrm{d}}^{\mathrm{dd}} x_2^{\mathrm{d}}}{\bar{\mu}_2^{\mathrm{d}} - \lambda_2^{\mathrm{d}}},
\end{aligned}$$

where the completion rates $\bar{\mu}_i^\delta$ are given by (3) when applied with shares p related to the optimal static policy π. Thus, in this FPI policy, the action chosen in state x is defined by

$$a^{\mathrm{FPI}}(x) = \arg\max \left\{ \frac{\mu_{\mathrm{u}}^{\mathrm{uu}} x_1^{\mathrm{u}}}{\bar{\mu}_1^{\mathrm{u}} - \lambda_1^{\mathrm{u}}} + \frac{\mu_{\mathrm{u}}^{\mathrm{uu}} x_2^{\mathrm{u}}}{\bar{\mu}_2^{\mathrm{u}} - \lambda_2^{\mathrm{u}}}, \frac{\mu_{\mathrm{d}}^{\mathrm{du}} x_2^{\mathrm{d}}}{\bar{\mu}_2^{\mathrm{d}} - \lambda_2^{\mathrm{d}}}, \frac{\mu_{\mathrm{d}}^{\mathrm{du}} x_1^{\mathrm{d}}}{\bar{\mu}_1^{\mathrm{d}} - \lambda_1^{\mathrm{d}}}, \frac{\mu_{\mathrm{d}}^{\mathrm{dd}} x_1^{\mathrm{d}}}{\bar{\mu}_1^{\mathrm{d}} - \lambda_1^{\mathrm{d}}} + \frac{\mu_{\mathrm{d}}^{\mathrm{dd}} x_2^{\mathrm{d}}}{\bar{\mu}_2^{\mathrm{d}} - \lambda_2^{\mathrm{d}}} \right\}. \tag{18}$$

6. STABILITY OF DYNAMIC POLICIES

The stability region for a dynamic policy may be different from the maximal stability region given in Theorem 1. In this section we discuss the stability regions of the various dynamic policies we have considered.

We know that, under the assumption of exponential service times, the FPI policy improves the optimal static policy and is thus maximally stable. In other words, it is stable whenever $\rho^{\mathrm{u}} + \rho^{\mathrm{d}} < 1$, where ρ^{u} and ρ^{d} are defined in (9). Moreover, the MW policy has been proven to be maximally stable [16].

From the numerical results (given in Section 8), we observe, however, that some of the priority policies are not maximally stable. In particular we take the HU policy as an example. Since it gives the absolute priority to the uplink flows, it is obvious that the stability problems first arise in the downlink queues. As the uplink queues are always served whenever non-empty, the downlink queues will be served only when both of them are empty, i.e., with probability

$$\mathrm{P}\{X_1^{\mathrm{u}} = 0, X_2^{\mathrm{u}} = 0\} = (1-\rho_1^{\mathrm{u}})(1-\rho_2^{\mathrm{u}}), \tag{19}$$

where $\rho_1^{\mathrm{u}} = \lambda_1^{\mathrm{u}}/\mu_{\mathrm{u}}^{\mathrm{uu}}$, $\rho_2^{\mathrm{u}} = \lambda_2^{\mathrm{u}}/\mu_{\mathrm{u}}^{\mathrm{uu}}$. This means that the downlink flows do not accumulate (and the queue length remains finite) if the fraction of time the downlink queue is occupied is less than the fraction of time both the uplink queues are empty. Thus, for stability, we require that $\lambda_i^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{dd}} < (1-\rho_1^{\mathrm{u}})(1-\rho_2^{\mathrm{u}})$ for $i = 1,2$. Here we have assumed that near the stability limit only the 'dd' mode is operational due to the low DL interference assumption (2). For the symmetric setting (Scenario 1), the stability condition becomes

$$\rho^{\mathrm{u}} + \rho^{\mathrm{d}} < 1 - \rho^{\mathrm{u}}(1-\rho^{\mathrm{u}}), \tag{20}$$

where $\rho^{\mathrm{u}} = \lambda^{\mathrm{u}}/\mu_{\mathrm{u}}^{\mathrm{uu}}$, $\rho^{\mathrm{d}} = \lambda^{\mathrm{d}}/\mu_{\mathrm{d}}^{\mathrm{dd}}$. Clearly the right hand side of (20) is strictly less than 1, which shrinks the stability region of the HU policy compared to that of the maximally stable policies.

7. STOCHASTIC OPTIMALITY RESULTS

In this section, we give our main theoretical results. We state that the HU policy is stochastically optimal in the special case that one of the uplink classes has no arrivals, say $\lambda_2^{\mathrm{u}} = 0$, and that the completion rate of uplink flows is sufficiently high, $\mu_{\mathrm{u}}^{\mathrm{uu}} \geq 2\mu_{\mathrm{d}}^{\mathrm{dd}}$. On the other hand, the HD policy is optimal if one of the downlink classes has no arrivals, say $\lambda_2^{\mathrm{d}} = 0$, and the maximal completion rate of downlink flows is high enough, $\mu_{\mathrm{d}}^{\mathrm{du}} \geq 2\mu_{\mathrm{u}}^{\mathrm{uu}}$. For these stochastic optimality results, we have to assume that the service times in each mode are exponential.[1]

Let $X^\pi(t)$ denote the total number of flows in the system at time t under policy π. Policy π^* is *stochastically optimal* (with respect to the total number of flows in the system) if, for all x, s, t,

$$\begin{aligned}
&\mathrm{P}\{X^{\pi^*}(t) > s \mid X^{\pi^*}(0) = x\} \\
&\quad = \min_\pi \mathrm{P}\{X^\pi(t) > s \mid X^\pi(0) = x\}.
\end{aligned}$$

Theorem 2. *The HU policy is stochastically optimal if*

$$\lambda_1^{\mathrm{u}} > 0, \quad \lambda_2^{\mathrm{u}} = 0 \quad \textit{and} \quad \mu_{\mathrm{u}}^{\mathrm{uu}} \geq 2\mu_{\mathrm{d}}^{\mathrm{dd}} > \mu_{\mathrm{d}}^{\mathrm{du}}.$$

Theorem 3. *The HD policy is stochastically optimal if*

$$\lambda_1^{\mathrm{d}} > 0, \quad \lambda_2^{\mathrm{d}} = 0 \quad \textit{and} \quad \mu_{\mathrm{d}}^{\mathrm{du}} \geq 2\mu_{\mathrm{u}}^{\mathrm{uu}}.$$

8. NUMERICAL RESULTS

We first show simulation results, where the service times are based on a spatial model taking explicitly into account the different locations of users. After that we present a systematic study with exponential service time distributions to figure out the potential performance gains of dynamic policies.

8.1 Physical model with non-exponential service times

We use a similar spatial model for the service times as in [6], where users with flows having a random size arrive in a random location in the cell where they are able to attain a certain mean transmission rate.

Consider a linear network of two sectorized base stations. The cell radius $R = 400$ [m] and hence the base stations are at a distance $2R$ from each other. Given a fixed number of flows in any class in our model, the location of the flow inside the cell is uniformly distributed in $[0, R]$, i.e., flows arrive to the base station area according to a spatial Poisson process. Let Z denote the random location of a flow with a uniform pdf $f_Z(z) = 1/R$, $z \in [0, R]$. Assume that the size of any flow is exponentially distributed with mean $\mathrm{E}[Y] = 100$ [kB]. Moreover, we denote the base station power by $P^{\mathrm{bs}} = 20$ [W] and the user equipment power by $P^{\mathrm{ue}} = 0.2$ [W]. Finally, we

[1]Due to lack of space, the proofs are not included in this paper but are given in a separate technical document [21].

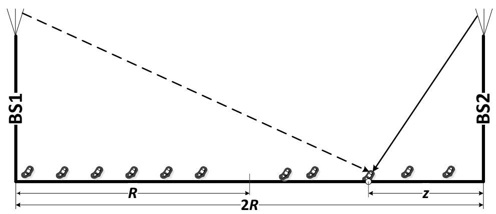

Figure 1: Linear sectorized network with two base stations.

assume that from outside the cells there is additional noise and interference equal to $I = 10^{-7}$ [W]. This is illustrated in Figure 1, which shows two base stations separated by a distance $2R$ together with a flow at distance z. The arrows reflect that in the 'dd' mode the useful signal is interfered by the signal from the other base station.

Thus, given an uplink flow at distance Z from the base station, the service time of the flow is given by $S_{\rm u}^{\rm uu} = Y/c^{\rm uu}(Z)$, where $c^{\rm uu}(z)$ represents the mean transmission rate of the flow at distance z in the 'uu' mode. Correspondingly, we get $S_{\rm d}^{\rm ud} = Y/c^{\rm ud}(Z)$, $S_{\rm d}^{\rm du} = Y/c^{\rm du}(Z)$, and $S_{\rm d}^{\rm dd} = Y/c^{\rm dd}(Z)$. Note that, due to our symmetry assumption with respect to the stations, $S_{\rm d}^{\rm ud}$ and $S_{\rm d}^{\rm du}$ are identically distributed. Now the transmission rates in the different modes at distance z are given by

$$c^{\rm uu}(z) = \min\left\{B\log_2\left(1+\frac{P^{\rm ue}(1/z)^\alpha}{I}\right), C_{\rm u}^{\max}\right\}, \quad (21)$$

$$c^{\rm du}(z) = c^{\rm ud}(z) = \min\left\{B\log_2\left(1+\frac{P^{\rm bs}(1/z)^\alpha}{I}\right), C_{\rm d}^{\max}\right\}, \quad (22)$$

$$c^{\rm dd}(z) = \min\left\{B\log_2\left(1+\frac{P^{\rm bs}(1/z)^\alpha}{P^{\rm bs}(1/(2R-z))^\alpha + I}\right), C_{\rm d}^{\max}\right\}, \quad (23)$$

where the parameter values are taken to correspond to typical values for an LTE system, such that bandwidth of the system $B = 10$ [MHz], path loss exponent $\alpha = 3$, uplink maximum rate $C_{\rm u}^{\max} = 50$ [Mbit/s], and downlink maximum rate $C_{\rm d}^{\max} = 150$ [Mbit/s]. The capacity functions reflect our interference modeling assumptions: in the 'uu', 'ud', and 'du' modes the other base station is not causing additional interference and only external interference I limits the rates, but in the 'dd' mode, the other base station is causing additional interference (dashed arrow) from distance $2R - z$.

Hence we have for the mean completion rate in the 'uu' mode

$$1/\mu_{\rm u}^{\rm uu} = {\rm E}[S_{\rm u}^{\rm uu}] = {\rm E}[Y]\cdot{\rm E}\left[1/c^{\rm uu}(Z)\right] \approx 0.48\,[1/{\rm s}].$$

Similarly, with our parameter values, we obtain $1/\mu_{\rm d}^{\rm ud} = 0.02$ [1/s] and $1/\mu_{\rm d}^{\rm dd} = 0.03$ [1/s]. Calculating also the variances, one sees that the squared coefficient of variation ranges from 1.6 ('ud' mode) to 3.3 ('uu' mode) and hence the distance dependent rates twist the original exponential size distribution to a more variable than exponential distribution. We fix the arrival rates as in Scenario 4, i.e., $\lambda_1^{\rm u} = 2\lambda_2^{\rm u} \neq 0$ and $\lambda_1^{\rm d} = 2\lambda_2^{\rm d} \neq 0$, and we vary the load within the stability limit.

We compare the performance of the physical model and the corresponding model with exponentially distributed service times in Figure 2. We clearly observe that HU and H3 become unstable quite soon. On the other hand, we see that MW and FPI are performing very robustly. Interestingly, although FPI is not guaranteed to be better than the optimal static policy in this physical model, we see that it is working very well. We also observe that HD is the best performing policy unless the load is very heavy. The simulations give some indications that HD is not necessarily maximally stable. All in all, the results for the physical model are very similar to those with exponential assumptions.

8.2 Potential performance gain of dynamic policies

We compare the performance of the proposed dynamic policies in the scenarios defined in Table 1 (in Section 4) by simulations. In all these scenarios, there are two free parameters $\lambda^{\rm u}$ and $\lambda^{\rm d}$. In our numerical studies, we consider the case where the uplink and downlink loads are equal, i.e., $\rho^{\rm u} = \rho^{\rm d} = \rho/2$. With this assumption, we still have one free parameter, the total load $\rho = \rho^{\rm u} + \rho^{\rm d}$. In addition, we assume that $\mu_{\rm d}^{\rm du} = 5$ [1/s], $\mu_{\rm d}^{\rm dd} = 3$ [1/s], and let $\mu_{\rm u}^{\rm uu}$ vary taking values in $\{1, 3, 5, 7\}$ [1/s].

The results for Scenarios 2-4 are shown in Figures 3–4.[2] In these figures, the ratio of the mean flow-level delay using a dynamic policy to that of the optimal static policy, $\bar{T}^{\rm dyn}/\bar{T}^{\rm sta}$, is plotted against the total load ρ. Due to Little's result, we have the ratio of the average delays equal to the ratio of the average total number of flows, i.e., $\bar{T}^{\rm dyn}/\bar{T}^{\rm sta} = \bar{X}^{\rm dyn}/\bar{X}^{\rm sta}$. For each fixed set of parameter values, the mean delay of a dynamic policy is estimated from a simulation run consisting of 500 000 arrivals. To reduce randomness in these comparisons, we have used the same random arrival sequences and the same random flow sizes for all the policies.

To gain further insight, we applied the policy iteration algorithm from the MDP theory to numerically estimate the performance of the optimal policy. The results have been obtained by truncating the state space and solving in each iteration step the value functions for each state from the associated Howard's equations and then performing numerically in each state the optimization step (16). This iteration is continued for a fixed number of steps. In our results, the maximum number of flows in each class was limited to 7 and the number of iteration rounds was 5. The ratio of the results from the policy iteration to the optimal static policy is depicted by the lines labeled MDP. Results are shown only up to load $\rho = 0.7$ because beyond that the truncation will start affecting them significantly. In fact, due to truncation, the graphs for the iterated MDP policy are curving down already for $\rho \approx 0.6$.

In Scenario 2 (Figure 3 (a) and (b)), where we have only one uplink queue being served, at low values of $\mu_{\rm u}^{\rm uu}$, HD performs very well and HU is the worst policy. When we increase the uplink service rate, $\mu_{\rm u}^{\rm uu}$, the performance of HU and H3 both become gradually better, but as long as $\mu_{\rm u}^{\rm uu} \leq 3$, the performance of HD is still the best and very

[2] Due to lack of space, the results of Scenario 1 are only given in the related technical document [21].

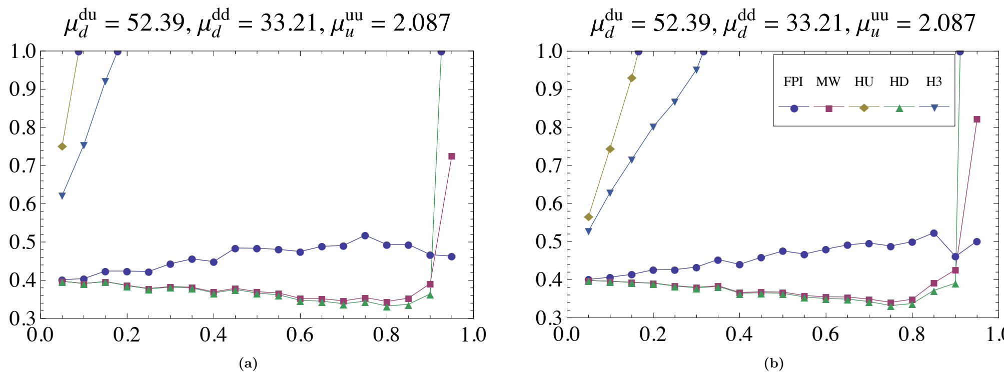

Figure 2: Performance of Scenario 4 using the physical model. The horizontal axis represents the traffic intensity ρ and the vertical axis is the ratio of the mean delay of the dynamic policy to that of the optimal static policy, $\bar{T}^{\rm dyn}/\bar{T}^{\rm sta}$. The left panel shows the performance of the different policies in the physical model while the right panel shows the performance of the same policies in an equivalent case where the service times are assumed to be exponentially distributed.

close to the optimal policy (as indicated by its closeness to the MDP curve) whenever the traffic load is not very high. When $\mu_{\rm u}^{\rm uu} = 7$ (Figure 3 (b)), we have $\mu_{\rm u}^{\rm uu} = 7 > 2\mu_{\rm d}^{\rm dd} = 6$, and thus, HU gives the best performance as it is the stochastically optimal policy according to Theorem 2. In addition, we see that MW and FPI, although not optimal, are relatively robust and maximally stable.

In Scenario 3 (Figure 3 (c) and (d)), where we have only one downlink queue being served so that HU and H3 are identical as expected and suffer from severe stability problem. When $\mu_{\rm u}^{\rm uu} = 1$ (Figure 3 (c)), then $2\mu_{\rm u}^{\rm uu} \leq \mu_{\rm d}^{\rm du} = 5$, and we have HD as the stochastically optimal policy according to Theorem 3. But when we set $\mu_{\rm u}^{\rm uu} = 7$, we have $2\mu_{\rm u}^{\rm uu} > \mu_{\rm d}^{\rm du} = 5$ and HD loses its optimality. Under this condition, the policy which gives priority to the uplink (H4) begins to perform better, although it is not optimal for all ρ. At higher values of $\mu_{\rm u}^{\rm uu}$, HD becomes worse than all of H4, FPI and MW, although it looks maximally stable like H4, FPI and MW. Moreover, based on Figures 3 (b) and 3 (c), the MDP results are numerically very close to the simulated performance of the stochastically optimal policies up to load $\rho = 0.6$.

Scenario 4 depicts the *station asymmetric* setting where the arrival rate in Station 1 is twice the arrival rate in Station 2, i.e., $\lambda_1^{\rm u} = 2\lambda_2^{\rm u} \neq 0$ and $\lambda_1^{\rm d} = 2\lambda_2^{\rm d} \neq 0$. In Figure 4, we again observe that HD is performing very nicely when $\mu_{\rm u}^{\rm uu} = 3$ or $\mu_{\rm u}^{\rm uu} = 5$ while the policies that give priority to uplink (HU and H3) perform the worst. As the value of $\mu_{\rm u}^{\rm uu}$ increases, the uplink prioritizing policies, HU and H3, start to perform better and become more stable while the performance of HD begins to deteriorate and turns into the worst performing policy for lower values of ρ when $\mu_{\rm u}^{\rm uu} = 7$. In all these cases, the performance of MW is very close to the optimal policy represented by the MDP curve. As $\mu_{\rm u}^{\rm uu}$ increases, the performance of FPI starts to get better and is practically indistinguishable from MW for $\mu_{\rm u}^{\rm uu} \geq 3$.

9. CONCLUSIONS

In this paper we have considered the intercell coordination problem between two interfering cells combined with dynamic TDD. Traffic in our model consists of elastic data flows in both cells and in both directions.

We have derived an explicit expression for the maximal stability condition. With this, it is possible to define the uplink and downlink traffic loads. The main focus is on the performance comparisons between various policies that can be used to operate the two stations. We have determined the optimal static policy for various scenarios, and considered further optimizing the performance by dynamic policies. In particular, we have experimented with various heuristic priority policies. We have also considered the dynamic policy called FPI resulting from applying the policy iteration algorithm from the MDP theory and the well-known max-weight policy in our performance comparisons.

Our main theoretical results show that under certain special conditions for the model parameters either giving absolute priority to the 'uu' mode or the 'dd' mode is stochastically optimal for exponential service times. These results generalize the earlier stochastic optimality results in a similar system with two base stations but supporting only downlink traffic. The dynamic TDD setting adds more classes and more modes to the system, which makes it significantly more challenging to establish the optimality for a given policy.

To obtain insight into the performance gains from the dynamic policies, we have carried out extensive simulations. We have observed that the priority policies, indeed, may give relatively good performance but some of them may run into stability problems when the load becomes too high. The max-weight and FPI policies are more robust and perform overall well. Even in the physical model with nonexponential service times, all the policies, excluding the ones prioritizing the uplinks, seem to be performing nicely for a wide range of load. Moreover, as we have experimented with many dynamic policies, the performance of the best among these policies appears to be close to the optimal policy. Jointly all the results give us an indication that in our scenarios the maximal gain from optimal dynamic coordination compared with the optimal static policy is approximately 50%–60%.

In this paper, we have focused on the low downlink interference case. In the future, it is worth investigating the high downlink interference case, as well.

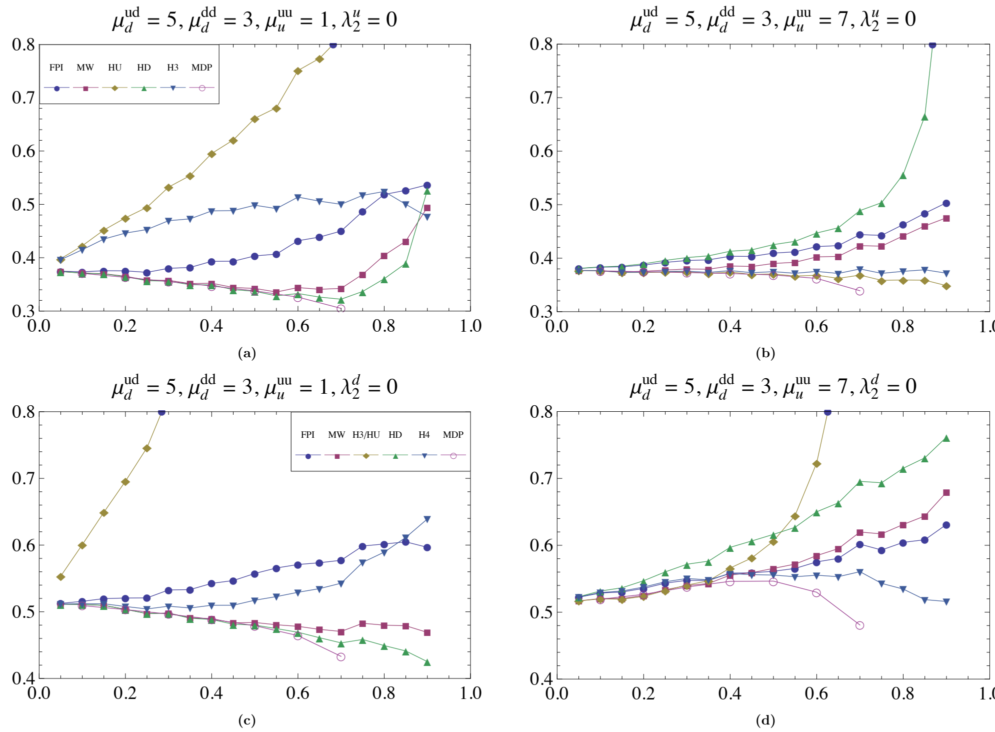

Figure 3: Performance of Scenario 2 (panels (a) and (b)) and Scenario 3 (panels (c) and (d)). The horizontal axis represents the total load ρ and the vertical axis is the ratio of the mean delay of the dynamic policy to that of the optimal static policy, $\bar{T}^{\text{dyn}}/\bar{T}^{\text{sta}}$. The legend in panel (a) shows the dynamic policies involved in Scenario 1 and the legend in panel (c) shows the different dynamic policies involved in Scenario 3.

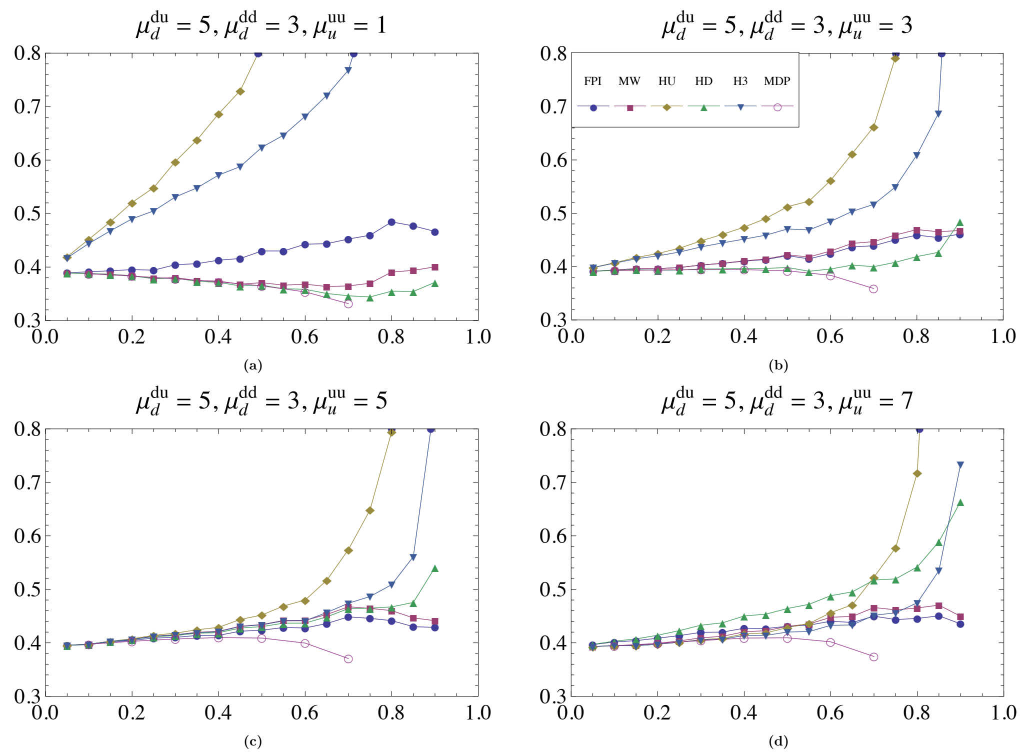

Figure 4: Performance of Scenario 4. The horizontal axis represents the traffic intensity ρ and the vertical axis is the ratio of the mean delay of the dynamic policy to that of the optimal static policy, $\bar{T}^{\text{dyn}}/\bar{T}^{\text{sta}}$. The legend in panel (b) shows the different dynamic policies involved.

10. REFERENCES

[1] R. Susitaival, H. Wiemann, J. Östergaard, and A. Larmo. Internet access performance in LTE TDD. In *Proc. of IEEE VTC 2010-Spring*, May 2010.

[2] Z. Shen, A. Khoryaev, E. Eriksson, and X. Pan. Dynamic uplink-downlink configuration and interference management in TD-LTE. *IEEE Communications Magazine*, 50(11):51–59, 2012.

[3] T. Nakamura, S. Nagata, A. Benjebbour, Y. Kishiyama, T.Hai, S. Xiaodong, Y. Ning, and L. Nan. Trends in small cell enhancements in LTE advanced. *IEEE Communications Magazine*, 51(2):98–105, 2013.

[4] C. Chiang, W. Liao, T. Liu, I. Chan, and H. Chao. Adaptive downlink and uplink channel split ratio determination for TCP-based best effort traffic in TDD-based WiMAX networks. *IEEE Journal Selected Areas in Communications*, 27(2):182–190, 2009.

[5] M. Al-Rawi and R. Jäntti. A dynamic TDD inter-cell interference coordination scheme for long term evolution networks. In *IEEE PIMRC*, April 2011.

[6] T. Bonald and A. Proutière. Wireless downlink data channels: user performance and cell dimensioning. In *Proc. of ACM MobiCom*, pages 339–352, September 2003.

[7] H. van den Berg, R. Litjens, and J. Laverman. HSDPA flow level performance: the impact of key system and traffic aspects. In *Proc. of ACM MSWiM*, pages 283–292, October 2004.

[8] S. Borst. User-level performance of channel-aware scheduling algorithms in wireless data networks. *IEEE/ACM Transactions on Networking*, 13(3):636–647, 2005.

[9] S. Aalto and P. Lassila. Impact of size-based scheduling on flow-level performance in wireless downlink data channels. In *Proc. of the 20th International Teletraffic Congress (ITC 20)*, pages 1096–1107, June 2007.

[10] D. C. Dimitrova, H. van den Berg, G. Heijenk, and R. Litjens. Flow level performance comparison of packet scheduling schemes for UMTS EUL. In *Proc. of the 6th International Conference on Wired/Wireless Internet Communications (WWIC'08)*, pages 27–40, May 2008.

[11] T. Bonald and N. Hegde. Capacity gains of some frequency reuse schemes in OFDMA networks. In *Proc. of IEEE GLOBECOM*, December 2009.

[12] S. Borst, N. Hegde, and A. Proutière. Interacting queues with server selection and coordinated scheduling — application to cellular data networks. *Annals of Operations Research*, 170:59–78, 2009.

[13] P. Lassila, A. Penttinen, and S. Aalto. Flow-level modeling and analysis of dynamic TDD in LTE. In *Proc. of the Eighth Euro-NF Conference on Next Generation Internet (NGI 2012)*, June 2012.

[14] S. Borst, M. Jonckheere, and L. Leskelä. Stability of parallel queueing systems with coupled service rates. *Discrete Event Dynamic Systems*, 18:447–472, 2008.

[15] J. Liu, A. Proutière, Y. Yi, M. Chiang, and H.V. Poor. Flow-level stability of data networks with non-convex and time-varying rate regions. *SIGMETRICS Performance Evaluation Review*, 35(1):239–250, 2007.

[16] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, 1992.

[17] A. Stolyar. Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Annals of Applied Probabability*, 14(1):1–53, 2004.

[18] IM Verloop and R. Núñez-Queija. Asymptotically optimal parallel resource assignment with interference. *Queueing Systems*, 65(1):43–92, 2010.

[19] Martin L Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, volume 414. Wiley-Interscience, 2009.

[20] Esa Hyytiä, Jorma Virtamo, Samuli Aalto, and Aleksi Penttinen. M/M/1-PS queue and size-aware task assignment. *Performance Evaluation*, 68(11):1136–1148, 2011.

[21] Prajwal Osti, Samuli Aalto, and Pasi Lassila. Flow-level modeling and optimization of intercell coordination with dynamic TDD (extended version). Technical report, 2014. http://users.comnet.aalto.fi/samuli/intercell-extended-version.pdf.