Flow-level Stability of Data Networks with Non-convex and Time-varying Rate Regions*

J. Liu†, A. Proutière‡, Y. Yi†, M. Chiang†, H. V. Poor†

ABSTRACT

In this paper we characterize flow-level stochastic stability for networks with non-convex or time-varying rate regions under resource allocation based on utility maximization. Similar to prior works on flow-level stability, we consider exogenous data arrivals with finite workloads. However, to model many realistic situations, the rate region, which constrains the feasibility of resource allocation, may be either non-convex or time-varying. When the rate region is fixed but non-convex, we derive sufficient and necessary conditions for stability, which coincide when the set of allocated rate vectors has continuous contours. When the rate region is time-varying according to some stationary, ergodic process, we derive the precise stability region. In both cases, the size of the stability region depends on the resource allocation policy, in particular, on the fairness parameter α in α-fair utility maximization. This is in sharp contrast with the substantial existing literature on stability under fixed and convex rate regions, in which the stability region coincides with the rate region for many utility-based resource allocation schemes, independently of the value of the fairness parameter. We further investigate the tradeoff between fairness and stability when rate region is non-convex or time-varying. Numerical examples of both wired and wireless networks are provided to illustrate the new stability regions and tradeoffs proved in the paper.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network communications; G.3 [**Probability and Statistics**]: Queueing theory, Stochastic processes

*This work was in part supported by NSF CNS-0417607, CNS-0625637, CCF-0448012, CCF-0635034 and DARPA CBMANET and ARO CTA.

†Department of Electrical Engineering, Princeton University, NJ. {jiapingl, yyi, chiangm, poor}@princeton.edu

‡KTH, Radio Communication Systems Electrum 418, SE-164 40 Kista, Sweden. alexandre.proutiere@radio.kth.se

SIGMETRICS'07, June 12–16, 2007, San Diego, California, USA.

General Terms

Performance

Keywords

Fairness, Resource Allocation, Stability, Network Utility Maximization.

1. INTRODUCTION

1.1 Motivation

Flows (or equivalently, end-to-end connections) in wired or wireless networks dynamically share resources (such as link capacities) according to various resource allocation schemes. These flows can be identified through their "classes", which define the set of network resources they require for the transfer of the corresponding packets. A popular family of schemes allocate resources to competing flows by distributively solving a network utility maximization problem [20]. The optimization constraint set captures the feasibility of allocations, and is referred to as *rate region.* The optimization objective in the form of utility functions can capture important notions such as traffic elasticity, user satisfaction, or fairness.

Extensive work on deterministic models of utility maximization has been conducted since the late 1990s, where flows constitute a static population and are assumed to have infinite backlog. In practice, the numbers of flows are varying as they are randomly generated by users and cease upon completion. This system can be viewed as a queuing network where the service rates depend on the solution to an optimization problem, which in turn depends on the number of active flows in each class.

A key performance requirement in data networks is that all flows are completed within a finite time, or that the numbers of active flows do not grow unbounded. Mathematically, this corresponds to the ergodicity of the process representing the numbers of flows of various classes. This property is referred to as *flow-level stability.* As will be briefly reviewed later in this section, a series of papers over the last eight years have provided necessary and sufficient conditions for flow-level stability in various models. With a couple of recent exceptions, these models assume a fixed and convex rate region. In this paper, we investigate conditions for flow-level stability when the rate region is either non-convex or time-varying.

Indeed, as explained in more detail in Subsection 2.2 and also the survey in [11], in many applications we cannot assume convexity or time-invariance of rate region in Network Utility Maximization models when studying flow-level stability. For example, non-convexity of rate region naturally arises in wireless cellular and ad-hoc networks [3, 18], and time-variation arises over time due to mobility, link failures, route or topology changes, and priority structures in resource allocation. It turns out that new proof techniques are needed to prove stability conditions in these scenarios, and intriguing tradeoffs between fairness and stability are discovered.

1.2 Related work

The first analysis of the flow-level stability focused on wired networks supporting data traffic only [5,14]. For such networks, the rate region is a (convex) polytope formed by the intersection of a finite number of linear capacity constraints, and it has been shown that all α-fair allocations with $\alpha > 0$ provide flow-level stability if and only if the vector representing the average traffic intensities of flow classes lies in the rate region. In other words, the rate region in the utility maximization problem is also the stability region under flow-level stochastic dynamics. This result has been generalized by many papers, e.g., [22,32,33], in particular, to the case of networks with arbitrary convex rate regions [6], and very recently to the case of general flow arrival processes and general flow size distributions [10,12,16,21,26,33]. It has also been shown in [6] that if the traffic intensity vector is outside of the rate region, then there is no allocation stabilizing the network at flow-level. These results imply that for fixed, convex rate regions, α-fair allocations maximize the flow-level stability region. This is sometimes called the throughput-optimality property for the utility-maximization-based resource allocations.

The analysis of the flow-level stability in case of fixed but *non-convex* rate region is generally very difficult, with very few existing work. In case of networks with two flow classes only, the stability condition of a large class of allocations can be exactly characterized [7]. However when the number of classes is greater than two, it has been generally impossible to derive an explicit and exact stability condition. This is mainly due to the fact that the stability condition depends on detailed statistical characteristics of the flow arrival processes, and flow departure processes, which are determined by the solutions of non-convex optimization problems. Some papers provide bounds on the stability region for very specific networks under particular allocations, see, e.g. [3, 25]. These papers study the stability of networks where the rate region reduces to a single point depending on the set of classes with active flows. Some other papers aim at providing exact stability conditions: in [9, 19, 29], a recursive (with respect to the number of flow classes) stability condition is given for a particular class of networks, including those studied in [3,25]. Unfortunately this kind of recursive formula often proves difficult to exploit: the stability condition of networks with S classes of flows depends on that of the network with $S-1$ classes and also on more detailed characterizations such as the probability that a given class has no active flows. Usually, these characterizations cannot be efficiently computed.

The analysis of the flow-level stability of networks with time-varying rate region has not been extensively studied so far. To the best of our knowledge, the only existing results provide the flow-level stability of wireless networks with user mobility under certain α-fair allocations [4, 8]. In [22] the largest possible stability region is studied with time-varying channel capacities under α-fair allocations, but the time-scale assumptions are completely different from our work.

As will be shown in Section 5, there are interesting tradeoffs between fairness and flow-level stability when rate region is non-convex or time-varying. This tradeoff is different from that between fairness and efficiency investigated for a static population of flows with infinite backlogs (see e.g., [27, 28] in wired networks, or [15, 24] for wireless networks, and [30] for a discussion on the absence of tradeoff in general topology). Here, we investigate the tradeoff between fairness and stochastic stability region, which quantifies the impact of fairness on the performance as perceived by users in a dynamic population of flows.

1.3 Overview

In this paper, we provide general stability conditions of α-fair allocations in networks with non-convex or time-varying rate regions. The main contributions are the followings:

(i) In networks with an arbitrary number of classes and with fixed but non-convex rate region, we give sufficient and necessary conditions for flow-level stability of α-fair allocations, for all $\alpha > 0$. We also prove that these conditions coincide when the set of allocated rate vectors is continuous (in a sense that will be defined in Section 3), leading to an explicit stability condition for such networks (Theorems 3, 4, 5, and Corollary 1).

(ii) We extend our analysis to networks with time-varying convex rate regions, for which we also provide the stability condition of α-fair allocations, for all $\alpha > 0$ (Theorems 6, 7). The results and proof techniques in (i) and (ii) can be readily combined for the general case of non-convex *and* time-varying rate region.

(iii) When rate region is either non-convex or time-varying, the stability condition is proven to depend on the chosen fairness parameter α. The exact degree of sensitivity with respect to α depends on the considered network, which can be significant (possibly changing the shape of stability region from concave to convex) or negligible. We provide examples for both situations. In two-class networks, we also prove that, as α increases, flow-level stability region shrinks (Corollary 2). In other words, there is a tradeoff between fairness and flow-level performance. Fairness can be enhanced but at the expense of reduced network stability. This is in sharp contrast to the case of fixed and convex rate region, where fairness has no impact on stability. This new phenomenon shows that the choice of the utility function is crucial to ensure a high user-level performance under non-convex or time-varying rate regions.

The paper is organized as follows. Section 2 is devoted to describing the system model and presenting the assumptions. In Sections 3 and 4, we provide the stability conditions for non-convex and time-varying rate regions, respectively. We discuss the tradeoff between fairness and stability in Section 5. We illustrate our theoretical results with concrete examples from both wired and wireless networks in Section 6, and conclude the paper in Section 7.

Notation and definitions. We first provide major definitions and notation used throughout the paper.

- For all $\boldsymbol{A}, \boldsymbol{B}$ in $\mathbb{R}^S$, $\boldsymbol{A} \le \boldsymbol{B}$ (resp. $\boldsymbol{A} < \boldsymbol{B}$) means that $\boldsymbol{A}$ is component-wisely smaller (resp. strictly smaller) than $\boldsymbol{B}$.

- A set $\mathcal{Y} \subset \mathbb{R}_+^S$ is *coordinate-convex* when the following is true: if $\boldsymbol{B} \in \mathcal{Y}$, then for all $\boldsymbol{A}$: $0 \le \boldsymbol{A} \le \boldsymbol{B}$, $\boldsymbol{A} \in \mathcal{Y}$.

- A set $\mathcal{Y} \subset \mathbb{R}_+^S$ is a *Pareto-type* set if $\nexists \boldsymbol{A}, \boldsymbol{B} \in \mathcal{Y}$ such that $\boldsymbol{A} < \boldsymbol{B}$.

- $\mathcal{Y}_\circ$ denotes the largest open subset of $\mathcal{Y}$.

- $c(\mathcal{Y})$ denotes the smallest closed set containing $\mathcal{Y}$.

- Define $\mathcal{U} = \{x \in \mathbb{R}_+^S : \sum_s x_s = 1\}$ and $\mathcal{D} : \mathbb{R}_+^S \mapsto \mathcal{U}$ the application giving the direction of vectors, i.e., $\mathcal{D}(v) = v/|v|$, where $|v| = \sum_s v_s$. We say that a Pareto-type set $\mathcal{Y}$ is continuous in direction u if the two following conditions are satisfied: (i) there exists $\epsilon > 0$ such that $\{v \in \mathcal{U} : |v - u| < \epsilon\} \subset \mathcal{D}(c(\mathcal{Y}))$; (ii) the application $\mathcal{D}^{-1} : \mathcal{U} \mapsto \mathcal{Y}$ is continuous at u. Condition (i) means that there are vectors in $\mathcal{Y}$ in all directions around u. Note that $\mathcal{D}^{-1}$ is well defined since $\mathcal{Y}$ is a Pareto-type set.

- A Pareto-type set $\mathcal{Y} \in \mathbb{R}_+^S$ is said to be continuous if $\forall u \in \mathcal{D}(c(\mathcal{Y}))$, $\mathcal{Y}$ is continuous in direction u.

2. SYSTEM MODEL

2.1 Traffic demand and network state

We consider a data network where flows are randomly generated by users and cease upon completion. Flows are classified according to the set of resources required to transfer the corresponding packets. For example, in wired networks with fixed routing, the class of a flow is defined by the set of links that the flow traverses from the source to the destination. We have a finite set $\mathcal{S}$ of S classes of flows. Flows of class s are generated according to a Poisson process of intensity λ_s flows per second. The sizes of class-s flows are i.i.d. exponentially distributed with mean size $1/\mu_s$ bits. We define the traffic intensity/offered load of flows of class s by $\rho_s = \lambda_s/\mu_s$ bit/s.

At time t, the network state is denoted by $\boldsymbol{N}(t) = (N_1(t), \cdots, N_S(t))$ where $N_s(t)$ is the number of active class-s flows. $\{\boldsymbol{N}(t)\}_{t=0}^{\infty}$ is a stochastic process governed by the random arrivals and departures of flows.

2.2 Rate region

The *rate region* $\mathcal{R}$ of the network is defined as the set of achievable rate vectors $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_S)$ where ϕ_s is the total rate allocated to class-s flows. A rate vector $\boldsymbol{\phi}$ is said to be achievable if there exist resource allocation mechanisms that can realize this vector. We assume here that the rate region does not depend on the network state $\boldsymbol{N}$. For example, consider a wired network with two links of respective capacities C_1 and C_2. Two flow classes compete for the use of these resources, class-1 flows require the use of both links whereas class-2 flows require that of the second link only. The corresponding rate region is then $\mathcal{R} = \{\boldsymbol{\phi} : \phi_1 + \phi_2 \le C_2, \phi_1 \le C_1\}$.

As illustrated in the previous example, the rate regions of wired networks are often convex and coordinate-convex sets. This is also the case for some wireless systems, mainly when a centralized resource allocation is permitted and a time-sharing argument *convexifies* the rate region. See e.g. [6] for many other examples of networks with convex rate regions. However, there are many situations where the rate region loses its convexity, for example, due to distributed resource allocation in wireless networks, or due to the fact that the achievable set of capacities is discrete. In cellular networks, the fact that the transmissions of the various base stations are not coordinated leads to non-convex rate regions [3]. In particular, when the achievable transmission power levels of a base station form a countable set, the rate region is discrete [7]. In wireless LANs, mesh or ad-hoc networks, users or nodes randomly access the radio channel in a distributed manner, which again induces non-convexity [18]. See [7] and Section 6 of the present paper for the example on distributed MAC scheduling. The first contribution of this paper is to analyze the performance of networks with non-convex rate region. In Section 3, we do not make any assumption on the rate region except that it is a compact subset of $\mathbb{R}_+^S$.

The second contribution of this paper is to study networks with time-varying capacities according to some exogenous processes (independent of the evolution of the network state). For example, in wired multi-service networks supporting low-priority data traffic and high-priority real-time traffic, the available capacity for data traffic is what is left by real-time traffic. The variations can also stem from link failures or from routing table changes. In wireless systems, fading as well as user mobility (in cellular networks) or node mobility (in ad hoc networks) also generate the capacity variations. Here we denote by $\mathcal{R}(t)$ the rate region at time t. We assume that the set $\mathcal{I}$ of indices of possible states $\{\mathcal{R}_i\}$ for the process $\{\mathcal{R}(t)\}_{t=0}^{\infty}$ is finite, and that $\{\mathcal{R}(t)\}_{t=0}^{\infty}$ is stationary and ergodic. We denote by $\boldsymbol{\pi}$ the stationary distribution of $\{\mathcal{R}(t)\}$, i.e., $\mathbb{P}\{\mathcal{R}(t) = \mathcal{R}_i\} = \pi_i$, $i \in \mathcal{I}$.

2.3 Resource allocation algorithms

Resource allocation algorithms allocate network resources to different flow classes according to the current network state $\boldsymbol{N}(t)$ and the current rate region $\mathcal{R}(t)$. Since the seminal work of Kelly et al. [20], optimization approaches have been extensively used to model and design the way these algorithms share the network resources. Most of the existing resource allocations aim at maximizing a certain notion of *utility* of the network. The realized allocation is then the solution of the following optimization problem:

$$\begin{array}{ll} \text{maximize} & \sum_s N_s(t) U_s(\phi_s/N_s(t)), \\ \text{subject to} & \boldsymbol{\phi} \in \mathcal{R}(t), \end{array} \quad (1)$$

where the utility functions U_s are usually assumed to be concave and non-decreasing. Here we also assume that all flow classes share the same utility function, i.e., $U_s = U$ for all s.

A large class of resource allocations are obtained based on the utility functions $U^\alpha(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha)$ for $\alpha > 0$, and $\log(\cdot)$, for $\alpha = 1$ [28]. The parameter α represents the degree of fairness of the allocation: when $\alpha = 0$, the total throughput of the network is maximized but the allocation may lead to user starvation and thus will not be considered in this paper; $\alpha = 1$ gives the Proportional fair allocation; when $\alpha \to \infty$, it corresponds to the Max-min fairness.

We denote the optimal solution of (1) at time t by $\boldsymbol{\phi}(\boldsymbol{N}(t))$ or $\boldsymbol{\phi}(\boldsymbol{N}(t), \mathcal{R}(t))$. For time-varying rate regions, this solution is denoted by $\boldsymbol{\phi}^{(i)}(\boldsymbol{N}(t))$ if $\mathcal{R}(t) = \mathcal{R}_i$. Since $\mathcal{R}(t)$ is compact, a solution of (1) exists. However for non-convex rate region, the solution is not necessarily unique.

Note that we could replace $\boldsymbol{N}(t)$ in (1) by any vector $\boldsymbol{N}$ in $\mathbb{R}_+^S$. We then denote by $\boldsymbol{\phi}(\boldsymbol{N}, \mathcal{R}(t))$ or $\boldsymbol{\phi}(\boldsymbol{N})$ the solution of optimization problem. When the rate region $\mathcal{R}(t)$ is convex, the solution $\boldsymbol{\phi}(\boldsymbol{N})$ corresponding to any α-fair allocation is unique and has the following properties:

Property 1 (Continuity): the mapping $\boldsymbol{N} \mapsto \boldsymbol{\phi}(\boldsymbol{N})$ is continuous on $\mathbb{R}_+^S$.

Property 2 (Homogeneity): For any $\boldsymbol{N}$ and any scalar $a > 0$, $\boldsymbol{\phi}(a\boldsymbol{N}) = \boldsymbol{\phi}(\boldsymbol{N})$.

Property 3 (Pareto Efficiency): The set $\{\boldsymbol{\phi}(\boldsymbol{N}), \forall \boldsymbol{N} \in \mathbb{R}_+^S\}$ is a Pareto-type set.

The proof of Properties 1, 2 is provided in [33], and Property 3 is due to the fact that any α-fair allocation with a compact rate region is Pareto efficient.

2.4 Time scale assumptions

The global system dynamics are induced by the flow arrivals/departures, the possible variations of the rate region, and the packet-level dynamics of the underlying resource allocation algorithms. The different time-scales of these sources of system dynamics play an important role in the performance analysis, denoted as follows:

(i) $\mathbb{T}_1$: the time-scale of the flow-level dynamics,

(ii) $\mathbb{T}_2$: the time-scale of the rate region variations,

(iii) $\mathbb{T}_3$: the time scale of resource allocation algorithm's convergence.

We assume that the time-scale of flow-level dynamics is much larger than that of resource allocation algorithms, i.e., $\mathbb{T}_1 \gg \mathbb{T}_3$. When the network state changes, the resource allocation algorithms are assumed to converge instantaneously to adapt the realized rate vector to this change. This assumption is often referred to as the *time-scale separation* assumption in the literature.

When the time-scales of rate region variations and of the resource allocation algorithms are similar, i.e., $\mathbb{T}_2 \approx \mathbb{T}_3$, these algorithms can directly take advantage of the rate region variations. Such systems are said to be *opportunistic*. A typical example of such systems is channel-aware scheduling in cellular networks [2], where fast-fading variations of the channels are exploited to get a greater throughput. When the rate region variations are not that fast, i.e. $\mathbb{T}_2 \gg \mathbb{T}_3$, being opportunistic proves more difficult and these variations can be exploited only at the expense of compromising the delay allowance of users. In this paper, the rate region variations are assumed to be relatively slow, as they can be generated by phenomena such as node mobility in wireless networks and link failures in wired networks.

To summarize, we assume that $\mathbb{T}_1, \mathbb{T}_2 \gg \mathbb{T}_3$, which means that the resource allocation algorithms instantaneously adapt the rate vector to either the numbers of active flows of various classes or the rate region variations. No assumption is made on the relative time-scales $\mathbb{T}_1$ and $\mathbb{T}_2$.

2.5 Flow-level stability

Our main focus in this paper is to prove necessary and sufficient conditions for flow-level stability: when will the durations of flows remain finite (almost surely)? Mathematically, stability means that the process $\{\boldsymbol{N}(t)\}_{t=0}^{\infty}$ is ergodic. With the assumptions in Section 2.4, this process is Markovian and evolves as follows: for each class s,

$$\begin{aligned} N_s(t) &\rightarrow N_s(t) + 1, \quad \text{with rate } \lambda_s \\ N_s(t) &\rightarrow N_s(t) - 1, \quad \text{with rate } \mu_s \phi_s(\boldsymbol{N}(t), \mathcal{R}(t)). \end{aligned}$$

The flow-level stability is now equivalent to the positive recurrence of the Markov process $\boldsymbol{N}(t)$, which implies the *almost sure* finiteness of the number of active flows in the system, i.e., flows that are being served or remain in the queues. In the following, we characterize the set of traffic intensity vectors $\boldsymbol{\rho} = (\rho_1, \ldots, \rho_S)$ such that flow-level stability can be realized. This set is referred to as the *stability region*, which also depends on the considered resource allocation algorithm. We say a compact set Γ is the stability region under certain resource allocation, if $\forall \boldsymbol{\rho} \in \Gamma_o$ such that the system is stable, and if $\forall \boldsymbol{\rho} \notin \Gamma$, the system is unstable.

On a related but different notion, the *maximum stability region* is defined by the union of all possible stability regions under all possible resource allocations, i.e., for any traffic intensity vector outside this set, there exists no resource allocation algorithm that can stabilize the network at flow-level. Note that such resource allocation may not be utility-based or implementable in a distributed fashion.

3. STABILITY WITH ARBITRARY FIXED RATE REGION

In this section, we investigate the flow-level stability of α-fair allocations for networks with arbitrary, but fixed rate region. We first recall the stability result for convex, coordinate-convex rate region $\mathcal{R}$, see e.g. [6, 22].

Theorem 1 (Convex rate regions). *For any convex, coordinate-convex rate region, the maximum stability region is the rate region, and is achieved by all α-fair allocations, provided that $\alpha > 0$.*

The above theorem states that α-fair allocations are optimal w.r.t. the flow-level stability. In particular, they all have the same stability region. Hence for convex, coordinate-convex rate region, fairness is not imposed at the expense of a reduction of the stability region. We now investigate the case where the rate region is not convex, in which case the stability region may strongly depend on the fairness parameter α. We begin by recalling the result providing the maximum stability region in case of arbitrary rate region [7]. We then study the case of discrete rate regions, i.e., rate regions composed by a finite number of rate vectors, and conclude this section with an analysis on the stability for arbitrary *continuous* rate regions.

3.1 Maximum stability region

The maximum stability region is given in the following theorem [7].

Theorem 2 (Maximum stability region [7]). *For a network with an arbitrary rate region $\mathcal{R}$, the maximum stability region is the smallest convex, coordinate-convex set containing $\mathcal{R}$.*

In particular, it has been proven in [7] that the so-called MaxProjection (MP) allocation introduced in [1] achieves

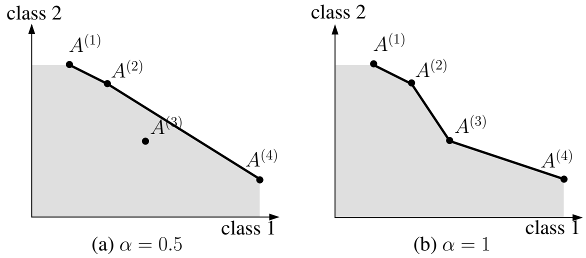

(a) $\alpha = 0.5$ (b) $\alpha = 1$

Figure 1: Different chosen rate vectors, contours and stability regions of a two-class network, for $\alpha = 0.5$ and $\alpha = 1$, with $A^{(1)} = (1, 4)$, $A^{(2)} = (2, 3.5)$, $A^{(3)} = (3, 2)$ and $A^{(4)} = (6, 1)$.

the maximum stability. For a given network state $\boldsymbol{N}$, the MP allocation allocates rates that form the solution of the following problem:

$$\max \sum_{s \in \mathcal{S}} N_s \phi_s, \text{ subject to } \boldsymbol{\phi} \in \mathcal{R} \tag{2}$$

It is worth noting that this allocation is not utility-based, thus it does not guarantee any fairness of resource allocation and there is no existing distributed implementation of this kind of allocation.

3.2 Discrete rate region

In case of arbitrary discrete rate regions, the stability condition of α-fair allocations turns out to be sensitive to detailed traffic demand characteristics, such as the flow size distribution, see e.g. [3, 25]. This explains why deriving an exact expression for the stability region proves quite challenging. However for networks with two flow classes only, the stability region is known and given by Theorem 3 [7]. In the rest of this paper, we denote by $\mathcal{R}^\alpha$ the set of rate vectors actually chosen by the α-fair allocation, i.e., the set of vectors $A \in \mathcal{R}$ such that there exists a state $\boldsymbol{N}$ with $\boldsymbol{\phi}(\boldsymbol{N}) = A$ for this allocation. Also for notational convenience, we prove the results of this section for α-fair allocations with $\alpha \neq 1$. They can be similarly proved for Proportional fair allocation given by $\alpha = 1$.

Theorem 3 (two classes [7]). *The stability region of an α-fair allocation, for $\alpha > 0$, is the smallest coordinate convex set containing the contour of $\mathcal{R}^\alpha$.*

Here for a two-class network, the contour of $\mathcal{R}^\alpha \in \mathbb{R}_+^2$ is defined as the broken line joining the allocated rate vectors from left to right. In general, $\mathcal{R}^\alpha$ depends on the allocation considered, which in turn leads to the dependence of the stability region on α. In Figure 1, we present the example of a two-class network with discrete rate region $\mathcal{R} = \{A^{(1)}, A^{(2)}, A^{(3)}, A^{(4)}\}$, and illustrate the dependence of the stability region on α. When $\alpha = 0.5$, $\mathcal{R}^\alpha = \{A^{(1)}, A^{(2)}, A^{(3)}\}$, and when $\alpha = 1$, $\mathcal{R}^\alpha = \mathcal{R}$. As a consequence, the Proportional fair allocation achieves a smaller stability region than the 0.5-fair allocation.

We generalize the result of Theorem 3 to the case of networks with an arbitrary number of flow classes. As explained above, deriving an exact expression for the stability region proves generally impossible. Hence we separately derive sufficient and necessary conditions for stability. Later, we will show that the gap between the sufficient and necessary conditions vanishes as the set of rate vectors chosen by the considered allocation gets continuous.

Consider the α-fair allocation in a network with a fixed discrete rate region $\mathcal{R} = \{A^{(1)}, \ldots, A^{(K)}\}$. We use fluid limits [13] to investigate stability, see [23] for more details. We denote by $\boldsymbol{n}$ the network state in the fluid limit with a continuous state space. Following [5, 33], the fluid limit $\boldsymbol{n}$ evolves according to

$$\frac{dn_s}{dt} = \begin{cases} \mu_s(\rho_s - A_s^{(l)}), & \text{if } n_s \neq 0, \\ \max\Big(\mu_s(\rho_s - A_s^{(l)}), 0\Big), & \text{if } n_s = 0, \end{cases} \tag{3}$$

for all $s \in \mathcal{S}$ and when the rate vector $A^{(l)}$ is allocated. The fluid limit is stable if $\boldsymbol{n}(t)$ reaches and stays at 0 within finite time. If starting from any initial point, the fluid limit reaches 0 in finite time, then the initial process $\{\boldsymbol{N}(t)\}_{t=0}^\infty$ is ergodic. The fluid limit is said to be unstable if $\|\boldsymbol{n}(t)\|$ grows at least linearly (after a finite time), and this instability implies the transience of the process $\{\boldsymbol{N}(t)\}_{t=0}^\infty$.

Define the subset $\mathcal{C}^{(j)}$ of the state space $\mathbb{R}_+^S$ (in the fluid limit) where the α-fair allocation allocates the rate vector $A^{(j)}$:

$$\mathcal{C}^{(j)} = \left\{\boldsymbol{n} : \boldsymbol{\phi}(\boldsymbol{n}) = A^{(j)}\right\}. \tag{4}$$

Note that each $\mathcal{C}^{(j)}$ is a cone due to the Property 2 (i.e., homogeneity) of the α-fair allocation. Some cones may be empty, in which case the corresponding rate vector is never allocated by the α-fair allocation. The cones defined in (4) satisfy:

(i) $\bigcup_{1 \le j \le K} \mathcal{C}^{(j)} = \mathbb{R}_+^S$,

(ii) $\mathcal{C}_o^{(j)} \bigcap \mathcal{C}_o^{(j')} = \emptyset$, for all $j \neq j'$,

(iii) the rate vector $A^{(j)}$ is allocated, if $\boldsymbol{n} \in \mathcal{C}_o^{(j)}$,

(iv) if $\boldsymbol{n} \in \mathcal{C}^{(j)} \bigcap \mathcal{C}^{(j')}$, then either $A^{(j)}$ or $A^{(j')}$ is allocated.

After defining the cone allocation, we will see in the following theorem that the stability condition depends on the comparison of traffic load and service rate (allocated rate vector).

Theorem 4 (N classes, sufficient condition). *For a discrete rate region $\mathcal{R}$, the stability region of the α-fair allocation, for $\alpha > 0$, contains Λ^α, the smallest coordinate convex set containing $\mathcal{R}^\alpha$.*

Proof. We prove the stability of the α-fair allocation for traffic vectors $\boldsymbol{\rho}$ in Λ^α whose distance to a particular rate vector $A^{(l)} \in \mathcal{R}^\alpha$ is sufficiently small so that $\boldsymbol{\rho} - A^{(l)} < 0$ and $|\rho_s - A_s^{(l)}| \ll \min_{l' \neq l} |A_s^{(l)} - A_s^{(l')}|$, $\forall s \in \mathcal{S}$. This also implies that $A^{(l)}$ is the only allocated rate vector in $\mathcal{R}$ such that $\boldsymbol{\rho} - A^{(l)} < 0$. Then stability for more general $\boldsymbol{\rho}$ is ensured using the fact that the stability region of α-fair allocations is coordinate-convex.

Denote by $\boldsymbol{\delta}^{(j)} = \boldsymbol{\rho} - A^{(j)}$ the drift vector in the cone $\mathcal{C}^{(j)}$. Define

$$V_{j,j'}(\boldsymbol{n}) = \sum_{s \in \mathcal{S}} n_s^\alpha \frac{(A_s^{(j)})^{1-\alpha} - (A_s^{(j')})^{1-\alpha}}{1-\alpha}. \tag{5}$$

Then at any time t, if the fluid limit is at state $\boldsymbol{n}$, $\boldsymbol{\phi}(\boldsymbol{n}) = A^{(j)}$ if and only if $V_{j,j'}(\boldsymbol{n}) \geq 0$ for all $1 \leq j' \leq K$ and the boundary of cone $\mathcal{C}^{(j)}$ and $\mathcal{C}^{(j')}$ is described by $\mathcal{C}^{(j)} \bigcap \mathcal{C}^{(j')} = \{\boldsymbol{n} : V_{j,j'}(\boldsymbol{n}) = 0\}$. Now introduce the function

$$L(\boldsymbol{n}(t)) = \sum_{1 \le l' \le K} V_{l',l}(\boldsymbol{n}(t)) \mathbf{1}_{\{\boldsymbol{n}(t) \in \mathcal{C}^{(l')}\}}. \qquad (6)$$

where we suppress the index of l in the notation of L. This function is continuous and differentiable almost everywhere (except at times where $\boldsymbol{n}(t)$ is at intersections of cones $\mathcal{C}^{(j)}$). Note that, when $\boldsymbol{n}(t) \in \mathcal{C}^{(l')}$, for $l' \ne l$, then $L(\boldsymbol{n}(t)) = V_{l',l}(\boldsymbol{n}) \ge 0$. Also note that when $L(\boldsymbol{n}(t)) = 0$, then the rate vector $A^{(l)}$ is scheduled.

Assume now that at time t, $L(\boldsymbol{n}(t)) > 0$, and that it is differentiable. This implies that there exists $l' \ne l$, such that $\boldsymbol{n}(t) \in \mathcal{C}^{(l')}$. We have:

$$\frac{dL}{dt} = \sum_{s \in \mathcal{S}} \alpha n_s^{\alpha-1} \mu_s (\rho_s - A_s^{(l')}) \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha}. \qquad (7)$$

Divide $\mathcal{S}(l')$ as $\mathcal{S}_a(l') \bigcup \mathcal{S}_b(l')$ where $\mathcal{S}_a(l') = \{s : A_s^{(l)} \ge A_s^{(l')}\}$, $\mathcal{S}_b(l') = \{s : A_s^{(l)} < A_s^{(l')}\}$.

Case (a): for $s \in \mathcal{S}_a(l')$, since we have chosen $\boldsymbol{\rho}$ to be sufficiently close to $A^{(l)}$, then either $A_s^{(l)} > \rho_s > A_s^{(l')}$ or $A_s^{(l)} = A_s^{(l')}$. Thus $\frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha} \le 0$ and $\rho_s - A_s^{(l')} \ge 0$, which gives $\alpha n_s^{\alpha-1} \mu_s (\rho_s - A_s^{(l')}) \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha} \le 0$;

Case (b): for $s \in \mathcal{S}_b(l')$, $\rho_s < A_s^{(l)} < A_s^{(l')}$, which also gives $\alpha n_s^{\alpha-1} \mu_s (\rho_s - A_s^{(l')}) \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha} < 0$.

Now we may rewrite $L(\boldsymbol{n}(t))$ as $L_a(\boldsymbol{n}(t)) + L_b(\boldsymbol{n}(t))$, where:

$$L_a(\boldsymbol{n}(t)) = \sum_{l'} \mathbf{1}_{\{\boldsymbol{n}(t) \in \mathcal{C}^{(l')}\}} \sum_{s \in \mathcal{S}_a(l')} n_s^{\alpha} \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha},$$

$$L_b(\boldsymbol{n}(t)) = \sum_{l'} \mathbf{1}_{\{\boldsymbol{n}(t) \in \mathcal{C}^{(l')}\}} \sum_{s \in \mathcal{S}_b(l')} n_s^{\alpha} \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha}.$$

We have $L_a(\boldsymbol{n}(t)) \le 0$, so that $L(\boldsymbol{n}(t)) \le L_b(\boldsymbol{n}(t))$. We now prove that $L_b(\boldsymbol{n}(t))$ reaches 0 in a finite time. Notice that

$$L_b(\boldsymbol{n}(t)) = \sum_{l'} \mathbf{1}_{\{\boldsymbol{n}(t) \in \mathcal{C}^{(l')}\}} \sum_{s \in \mathcal{S}_b(l')} b_s n_s^{\alpha},$$

where

$$b_s = \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha} > 0.$$

As in (7), for $\boldsymbol{n}(t) \in \mathcal{C}^{(l')}$, we write

$$\frac{dL_b(\boldsymbol{n}(t))}{dt} = - \sum_{s \in \mathcal{S}_b(l')} a_s n_s^{\alpha-1}(t), \qquad (8)$$

where

$$a_s = -\alpha \mu_s (\rho_s - A_s^{(l')}) \frac{(A_s^{(l')})^{1-\alpha} - (A_s^{(l)})^{1-\alpha}}{1-\alpha},$$

and it has been shown above that $a_s < 0$ for all $s \in \mathcal{S}_b(l')$. We can then easily deduce that there exists $\beta > 0$ such that for all time t (see [23]):

$$\frac{dL_b(\boldsymbol{n}(t))}{dt}(t) \le -\beta L_b(\boldsymbol{n}(t))^{\frac{\alpha-1}{\alpha}}. \qquad (9)$$

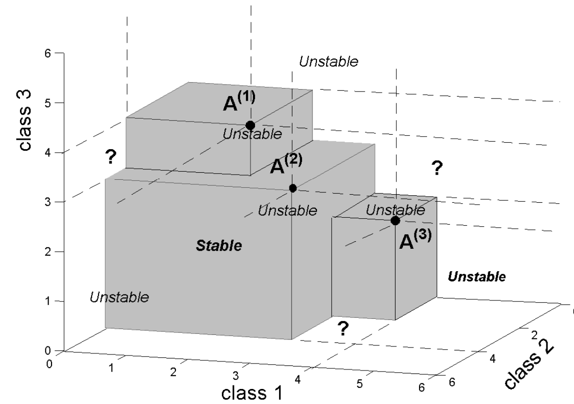

Figure 2: Lower bound on the stability region (Theorem 4) of an α-fair allocation in a 3-flow class system - $\alpha = 0.5$, $\mathcal{R}^{\alpha} = \{(4,2,2),(3,4,3),(2,3,4)\}$.

Then

$$L_b(\boldsymbol{n}(t)) \le \left(L_b(\boldsymbol{n}(0))^{1/\alpha} - \beta t/\alpha \right)^{\alpha}.$$

Thus after finite time T_1, $L_b(\boldsymbol{n}(t)) = 0$. This implies that for all $t \ge T_1$, $L_b(\boldsymbol{n}(t)) = 0$. It further implies that the rate vector $A^{(l)}$ is scheduled after time T_1 and that the system empties in a finite time. □

We now provide a necessary stability condition for α-fair allocations. Note that Theorem 2 already provides a necessary stability condition, since the traffic load cannot exceed the maximum stability region. The following result gives tighter necessary conditions.

THEOREM 5 (N CLASSES, NECESSARY CONDITIONS). *For a discrete rate region $\mathcal{R}$, the α-fair allocation, for $\alpha > 0$, is unstable if one of the following conditions holds:*

(i) *There exists $A^{(l)} \in \mathcal{R}^{\alpha}$ such that $\boldsymbol{\rho} - A^{(l)} > 0$.*

(ii) *There exists a class s such that $\rho_s > \max_{A^{(j)} \in \mathcal{R}^{\alpha}} A_s^{(j)}$.*

PROOF. Assume (i) holds. We use similar arguments as in the proof of Theorem 4 to show instability. Again we only need to consider the case when $\boldsymbol{\rho}$ is sufficiently close to $A^{(l)}$, i.e., $\boldsymbol{\rho} - A^{(l)} = \boldsymbol{\epsilon} > 0$ and $|\rho_s - A_s^{(l)}| \ll \min_{l' \ne l} |A_s^{(l)} - A_s^{(l')}|$. We use the same Lyapunov-like function $L(\boldsymbol{n})$, to prove that after a finite time the rate vector $A^{(l)}$ has to be scheduled (more precisely we prove that there exists T_3 such that $L(\boldsymbol{n}(t)) = 0$ if $t > T_3$). This implies that after a finite time, $\boldsymbol{n}(t)$ grows linearly to infinity when $t > T_3$, which further implies the the system is unstable.

Now assume (ii) holds, then at any state $\boldsymbol{n}(t)$ and with any allocated rate vector $A^{(j)}$, $\delta_s^{(j)} > 0$, i.e., the drift of class s is strictly positive and n_s is always increasing. Thus $\|\boldsymbol{n}\|$ linearly grows to infinity, and the network is unstable. □

3.3 Continuous non-convex rate region

When the number of rate vectors allocated by the α-fair allocation is small, there can be a significant gap between the sufficient and necessary conditions for stability regions derived in Theorem 4 and Theorem 5 as shown in Figure 2. When $\mathcal{R}^\alpha$ has more points, the gap reduces, and ultimately tends to 0 when $\mathcal{R}^\alpha$ becomes continuous, which is an important special case often encountered in utility maximization models. The following result formalizes this observation.

Corollary 1 (Continuous rate region). *If the set $\mathcal{R}^\alpha$ for α-fair allocation is continuous, then the stability region of this allocation is the smallest coordinate-convex set containing $c(\mathcal{R}^\alpha)$.*

Proof. When the set of allocated vectors $\mathcal{R}^\alpha$ is continuous, we can approximate it by a sequence of discrete rate regions with finite number of rate vectors. Let $\mathcal{R}^{(k)}$ be a discrete subset of $\mathcal{R}^\alpha$ such that $\mathcal{R}^{(k)} \uparrow \mathcal{R}^\alpha$ as $k \to \infty$. Now for the same system with discrete rate region $\mathcal{R}^{(k)}$, we let $\Lambda_{suf}^{(k)}$ denote the sufficient stability region defined by Theorem 4, which is the smallest coordinate-convex set containing $\mathcal{R}^{(k)}$, and $\Lambda_{nec}^{(k)}$ denote the necessary stability region defined by Theorem 5 as the complement of the unstable region. Thus if $\Lambda^{(k)}$ denotes the exact stability region of the system with rate region $\mathcal{R}^{(k)}$, we must have

$$\Lambda_{suf}^{(k)} \subseteq \Lambda^{(k)} \subseteq \Lambda_{nec}^{(k)}. \tag{10}$$

If we let Λ^α denote the smallest coordinate-convex set containing $c(\mathcal{R}^\alpha)$, by letting $k \to \infty$, $\Lambda_{suf}^{(k)} \uparrow \Lambda^\alpha$ and $\Lambda_{nec}^{(k)} \downarrow \Lambda^\alpha$, since $\mathcal{R}^{(k)} \uparrow \mathcal{R}^\alpha$ and $\mathcal{R}^\alpha$ is continuous. Hence when $k \to \infty$, the gap between the sufficient and necessary stability conditions vanishes, and the exact stability region is given as Λ^α. □

In the following section, we will consider time-varying convex rate regions, and in that case the set of allocated rate vectors is continuous by the convexity of each possible rate region. Then a similar phenomenon as described in Corollary 1 occurs, which explains why we will be able to exactly characterize the stability region.

4. STABILITY WITH TIME-VARYING RATE REGION

We now investigate the stability region of various resource allocations in networks with time-varying convex rate region. The network state is described by $(\boldsymbol{N}(t), \mathcal{R}(t))$ where we assume $\{\mathcal{R}(t)\}_{t=0}^{\infty}$ is a stationary and ergodic process as described in Subsection 2.2.

The proof techniques applied to obtain sufficient and necessary conditions for stability are similar to those used in the previous section. We will characterize the maximum stability region, and then derive the stability region of α-fair allocations. We first describe the evolution of the network in the fluid limit for any type of allocation. Consider an allocation which allocates the rate vector $\boldsymbol{\phi}^{(i)}(\boldsymbol{N})$ at state $\boldsymbol{N}$ when the rate region is $\mathcal{R}_i$, and satisfies Properties 1-3 defined in Subsection 2.3. The evolution of the system fluid limit is given by:

$$\frac{dn_s}{dt} = \lambda_s - \mu_s \sum_{i\in\mathcal{I}} \pi_i \phi_s^{(i)}(\boldsymbol{n}), \quad \forall s \in \mathcal{S}. \tag{11}$$

The proof of the above statement is presented in [23].

4.1 Maximum Stability Region

The following theorem is the analog of Theorem 2 for networks with fixed and arbitrary rate region. In that case, it turns out that the MP allocation also achieves maximum stability. Recall that the MP allocation solves (2) with rate region $\mathcal{R}(t)$ at any time t.

Theorem 6 (Maximum Stability region). *Consider a network with time-varying convex rate region $\mathcal{R}(t)$. The maximum stability region is*

$$\overline{\mathcal{R}} \ = \ \sum_{i\in\mathcal{I}} \pi_i \mathcal{R}_i \tag{12}$$

and it can be achieved by the MP allocation.

In the above theorem, the addition of sets is defined as follows: $\mathcal{R}_1 + \mathcal{R}_2 = \{x_1 + x_2 : x_1 \in \mathcal{R}_1, x_2 \in \mathcal{R}_2\}$.

Proof. *Necessary condition.* Since each $\mathcal{R}_i$ is convex, coordinate-convex, and compact, then the same properties hold for $\overline{\mathcal{R}}$. Moreover, the service rate $\sum_{i\in\mathcal{I}} \pi_i \boldsymbol{\phi}^{(i)}(\boldsymbol{n})$ in the fluid limit of any allocation belongs to $\overline{\mathcal{R}}$. If $\boldsymbol{\rho} \notin \overline{\mathcal{R}}$, we show that the fluid limit is unstable using similar arguments as those used in [6]. By convexity of $\overline{\mathcal{R}}$, there exists a half-space $\mathcal{H}$ containing $\overline{\mathcal{R}}$ such that $\boldsymbol{\rho} \notin \mathcal{H}$. Note that by compactness and coordinate convexity of $\overline{\mathcal{R}}$, we can choose this half-space such that its boundary intersects each coordinate axis. We denote by ϱ_s the intersection of the boundary of $\mathcal{H}$ with the class-s axis, which means $\mathcal{H} = \left\{\boldsymbol{\varphi} \mid \sum_{s\in\mathcal{S}} \frac{\varphi_s}{\varrho_s} \le 1\right\}$. Then $\sum_{i\in\mathcal{I}} \pi_i \phi_s^{(i)}(\boldsymbol{n}) \le \varrho_s$ for all $s \in \mathcal{S}$ and all state $\boldsymbol{n}$, and $\sum_{s\in\mathcal{S}} \frac{\rho_s}{\varrho_s} > 1$. Now consider the following Lyapunov function $L(\boldsymbol{n}) = \sum_{s\in\mathcal{S}} \frac{n_s}{\mu_s \varrho_s}$.

Necessity result follows from:

$$\frac{dL}{dt} = \sum_{s\in\mathcal{S}} \frac{\rho_s - \sum_{i\in\mathcal{I}} \pi_i \phi_s^{(i)}(\boldsymbol{n}(t))}{\varrho_s} \ge \left(\sum_{s\in\mathcal{S}} \frac{\rho_s}{\varrho_s}\right) - 1 > 0.$$

Sufficient condition. Let $\boldsymbol{\phi}^{(i),M}(\boldsymbol{n})$ denote the allocated rate vector under the MP policy when $\mathcal{R}(t) = \mathcal{R}_i$., and also let $\bar{\boldsymbol{\phi}}^M = \sum_{i\in\mathcal{I}} \pi_i \phi_s^{(i),M}$.

By [1,7], the stability region of the MP allocation is just the rate region, when the latter is fixed, convex and coordinate-convex. Then it suffices to check that the service rate $\bar{\boldsymbol{\phi}}^M$ corresponds to the MP allocation in case the rate region is $\overline{\mathcal{R}}$. We show the optimality of $\bar{\boldsymbol{\phi}}^M$ at state $\boldsymbol{n}$. Since for any $\boldsymbol{x} \in \overline{\mathcal{R}}$, $\boldsymbol{x}$ can be represented as $\sum_{i\in\mathcal{I}} \pi_i \boldsymbol{x}^{(i)}$, with $\boldsymbol{x}^{(i)} \in \mathcal{R}_i$, for all $i \in \mathcal{I}$, then we have:

$$\begin{aligned}
\sum_{s\in\mathcal{S}} n_s \sum_{i\in\mathcal{I}} \pi_i \phi_s^{(i),M}(\boldsymbol{n}) &= \sum_{i\in\mathcal{I}} \pi_i \sum_{s\in\mathcal{S}} n_s \phi_s^{(i),M}(\boldsymbol{n}) \\
&\ge \sum_{i\in\mathcal{I}} \pi_i \sum_{s\in\mathcal{S}} n_s x_s^{(i)} \\
&= \sum_{s\in\mathcal{S}} n_s \sum_{i\in\mathcal{I}} \pi_i x_s^{(i)}.
\end{aligned}$$

This completes the proof. □

Note that even with different time-scale assumptions, a similar result of Theorem 6 was provided in [22] for the specific case where $\mathcal{R}_i$ is a convex polytope, where a certain channel-aware scheduling is adopted to satisfy the maximum stability region under time-varying rate regions.

4.2 Stability Region of α-Fair Allocations

We now turn to the characterization of the stability region of α-fair allocations. Observe that by (11), the possible service rate for an α-fair allocation in the fluid limit is the average of the allocated rate vectors in the various rate regions. It is then natural to define the *average* set of rate vectors allocated by the α-fair allocation in the fluid limit as:

$$\partial\overline{\mathcal{R}_\alpha} = \{\boldsymbol{\phi} : \exists \boldsymbol{n} \in \mathbb{R}_+^S, \boldsymbol{\phi} = \sum_i \pi_i \times \boldsymbol{\phi}^{(i)}(\boldsymbol{n})\}. \tag{13}$$

This is the set of all possible service rate vectors in the fluid limit. We further define the *average rate region* in the fluid limit for the α-fair allocation as the smallest coordinate-convex set containing $\partial\overline{\mathcal{R}_\alpha}$, i.e.,

$$\overline{\mathcal{R}_\alpha} = \{\boldsymbol{y} : \exists \boldsymbol{x} \in \partial\overline{\mathcal{R}_\alpha} \quad s.t. \quad 0 \le \boldsymbol{y} \le \boldsymbol{x}\}. \tag{14}$$

In the following we assume that $\partial\overline{\mathcal{R}_\alpha}$ is a Pareto-type set, just like the sets of rate vectors allocated by the α-fair allocation when the rate region is fixed. In all examples presented in Section 6, this assumption is valid. In future work, we will further analyze the properties of $\partial\overline{\mathcal{R}_\alpha}$. Now, the stability region for α-fair allocations is given by the following result.

Theorem 7 (time-varying rate region). *For all $\alpha > 0$, the stability region of the α-fair allocation for time-varying rate region is $\overline{\mathcal{R}_\alpha}$.*

Proof. To investigate the stability of the fluid limit (11), we introduce a sequence of systems with discrete time-varying rate region which converges to the original system, and then apply similar techniques as in Theorem 4.

Denote by $\partial\mathcal{R}_i$ the smallest Pareto-type set $\mathcal{Y}$ such that the smallest coordinate-convex containing $\mathcal{Y}$ is $\mathcal{R}_i$. Note that this definition makes sense since $\mathcal{R}_i$ is convex. We consider a sequence of systems where the k-th system has time-varying, finite and discrete rate regions such that $\mathcal{R}_i^{(k)}$ is a subset of $\partial\mathcal{R}_i$, for all $i \in \mathcal{I}$. In particular, each discrete $\mathcal{R}_i^{(k)} = \{A^{(i,1)}, \ldots, A^{(i,k)}\}$ is a Pareto-type set. The considered sequence is such that $\mathcal{R}_i^{(k)} \uparrow \partial\mathcal{R}_i$ as $k \to \infty$. By (11), the set of allocated rate vectors in the fluid limit for the k-th system is $\partial\overline{\mathcal{R}_\alpha^{(k)}} = \{\sum_{i\in\mathcal{I}} \pi_i A^{(i,l_i)} : \exists \boldsymbol{n} \in \mathbb{R}_+^S, \boldsymbol{\phi}^{(i)}(\boldsymbol{n}) = A^{(i,l_i)}, \forall i \in \mathcal{I}\}$. If we define $\overline{\mathcal{R}_\alpha^{(k)}}$ as the smallest coordinate convex set containing $\partial\overline{\mathcal{R}_\alpha^{(k)}}$, then $\overline{\mathcal{R}_\alpha^{(k)}} \uparrow \overline{\mathcal{R}_\alpha}$ as $k \to \infty$. We now show the sufficient and necessary stability regions for the k-th system in a similar way as we proved Theorems 4 and 5, and observe that the gap between the sufficient and necessary conditions goes to zero when $k \to \infty$.

In the k-th system, when the rate region is $\mathcal{R}_i^{(k)}$, the α-allocation is a cone policy, allocating the rate vector $A^{(i,l_i)}$ if the state $\boldsymbol{n}(t)$ belongs to the cone $\mathcal{C}^{(i,l_i)}$. We introduce the cones $\mathcal{C}^{(l_1,\cdots,l_I)}$, which correspond to the states where the rate vector in the fluid limit is $\sum_{i\in\mathcal{I}} \pi_i A^{(i,l_i)}$. These cones are defined by intersections of cones:

$$\begin{aligned}\mathcal{C}^{(l_1,\cdots,l_I)} &= \bigcap_{i\in I} \mathcal{C}^{(i,l_i)} \\ &= \left\{\boldsymbol{n} : V_{l_i,l}^{(i)}(t) \ge 0, \forall 1 \le l \le K, \forall i \in \mathcal{I}\right\}.\end{aligned}$$

For each possible rate region $\mathcal{R}_i^{(k)}$, we define

$$V_{j,j'}^{(i)}(\boldsymbol{n}) = \sum_{s\in\mathcal{S}} n_s(t)^\alpha \left((A^{(i,j)})^{1-\alpha} - (A^{(i,j')})^{1-\alpha}\right)/(1-\alpha).$$

Now as in the proof of Theorem 4, to prove stability when there exists $\boldsymbol{\nu} \in \overline{\mathcal{R}_\alpha^{(k)}}$ such that $\boldsymbol{\rho} < \boldsymbol{\nu}$, we can use the following function:

$$L(\boldsymbol{n}(t)) = \sum_{i\in\mathcal{I}} \sum_{l'_i} V_{l'_i,l_i}^{(i)}(\boldsymbol{n}) \mathbf{1}_{\{\boldsymbol{n}(t)\in\mathcal{C}^{(l'_i)}\}}. \tag{15}$$

On the other hand, if $\boldsymbol{\rho} \notin \overline{\mathcal{R}_\alpha}$, there exists a sufficiently large k' such that for all $k > k'$ either one of the following conditions holds: there exists $\boldsymbol{A} \in \partial\overline{\mathcal{R}_\alpha^{(k)}}$ such that $\boldsymbol{\rho} - \boldsymbol{A} > 0$; $\rho_s > \max_{\boldsymbol{A}\in\partial\overline{\mathcal{R}_\alpha^{(k)}}} A_s$ for some class s. These two conditions correspond to cases (i) and (ii) in Theorem 5 respectively. Then the fluid limit is unstable with such $\boldsymbol{\rho}$ for any discrete rate region $\mathcal{R}_i^{(k)}$ where $k > k'$. By continuity, for all $\boldsymbol{\rho} \notin \overline{\mathcal{R}_\alpha}$ the system is unstable. □

5. FAIRNESS-STABILITY TRADEOFF

In this section, we discuss the tradeoff between fairness and flow-level stability. Namely, we study the sensitivity of the stability region of α-fair allocations to the fairness parameter α. When the rate region is fixed convex and coordinate-convex, we know from Theorem 1 that the stability region is insensitive to α. This property is lost for networks with non-convex or time-varying rate region. In this case, quantifying the sensitivity of the stability region w.r.t. α proves quite challenging, and we restrict the analysis to the case of networks with two classes only.

5.1 Sensitivity in case of non-convex rate region

A preliminary sensitivity analysis in case of non-convex rate region has been performed in [7]. It has been proved that there exist two fairness parameters γ and β with $\gamma < \beta$ such that for all $\alpha > \beta$ and $\alpha < \gamma$, the stability region of α-fair allocation is minimum and maximum, respectively. This result indicates that the stability region tends to be larger for smaller values of α. In particular Max-min fairness always leads to the smallest stability region, whereas the allocation maximizing the network throughput leads to the greatest region.

5.2 Sensitivity in case of time-varying rate region

We now investigate the sensitivity of the stability region to the choice of the resource allocation for time-varying rate regions. We provide two results indicating that for time-varying rate region, the stability region of α-fair allocations is reduced when α grows.

We consider time-varying rate regions satisfying the *scaling rule*, i.e., $\mathcal{R}_i = a^{i,j} \times \mathcal{R}_j$ for all $i,j \in \mathcal{I}$ and some $a^{i,j} \in \mathbb{R}_+^S$, where the product is defined as the component-wise scaling by factor $a_s^{i,j}$ for class-s coordinate. The scaling rule indicates that the shape of rate regions does not change. This assumption is valid in many practical systems since the type of network and resource allocation scheme (e.g., time or power sharing in wireless networks) determines the shape of the rate region. We present the following two corollaries to show the fairness-stability tradeoff in two-class networks, and refer the readers to [23] for detailed proofs.

Corollary 2. *For $\alpha \ge 1$, the stability region $\overline{\mathcal{R}_\alpha}$ is decreasing as α increases, i.e., $\overline{\mathcal{R}_{\alpha_1}} \subseteq \overline{\mathcal{R}_{\alpha_2}}$ if $\alpha_1 > \alpha_2$. In particular, if $\partial\mathcal{R}_1, \partial\mathcal{R}_2$ are hyperplanes, the monotonicity of $\overline{\mathcal{R}_\alpha}$ holds for all $\alpha > 0$.*

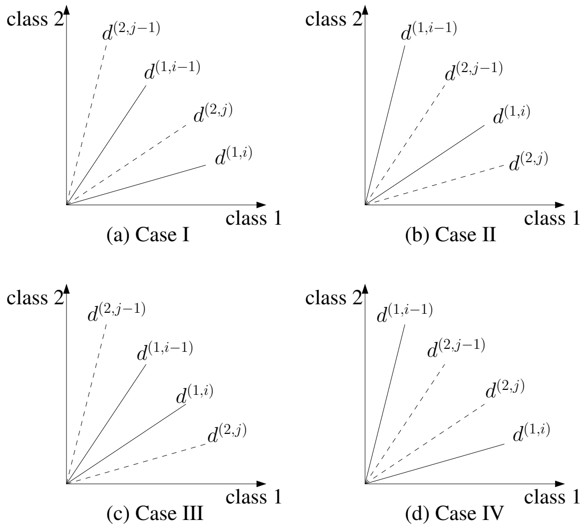

Figure 3: Cones and boundaries as $A^{(1,i)} + A^{(2,j)}$ is allocated.

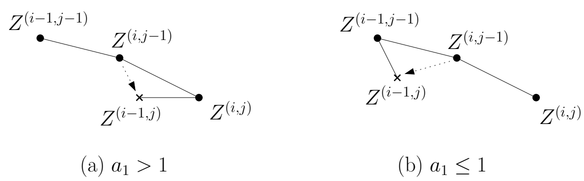

Figure 4: The transitions of allocated rate vectors when α increases.

PROOF. Here we sketch the main idea and key steps of the proof. Let $\boldsymbol{\phi}^{(i,\alpha)}(\boldsymbol{n})$ denote the α-fair allocated rate vector at state $\boldsymbol{n}$ when $\mathcal{R}(t) = \mathcal{R}_i$. First we consider $\mathcal{R}(t) = \mathcal{R}_1$ or $\mathcal{R}_2$ with $\mathcal{R}_2 = a \times \mathcal{R}_1$, $a = (a_1, a_2)$ and $a_2 > a_1$, where the result can be extended to finitely many possible rate regions by induction. Here without loss of generality we assume that $\pi_1 = \pi_2$ (just scaling the two rate regions). Then we proceed with discrete approximation of rate regions as in the proof of Theorem 7, where $\mathcal{R}_1^{(k)} = \{A^{(1,1)}, \ldots, A^{(1,k)}\}$. Let us fix the allocation, i.e., fix α. If $\mathcal{C}^{(1,i)}$ denotes the set of states $\boldsymbol{n}$ when $A^{(1,i)} \in \mathcal{R}_1^{(k)}$ is allocated, then by [7] $\mathcal{C}^{(1,i)} \bigcap \mathcal{C}^{(1,i+1)}$ is a line containing $(0,0)$, and denote by $d^{(1,i)}$ the tangent of the angle between this line and the class 2 axis. When the rate vector $Z^{(i,j)} = A^{(1,i)} + A^{(2,j)}$ is allocated, i.e., $\mathcal{C}^{(1,i)} \bigcap \mathcal{C}^{(2,j)} \neq \emptyset$, there are four possible cases for $\mathcal{C}^{(1,i)}$ and $\mathcal{C}^{(2,i)}$ as shown in Figure 3. Now we will slightly increase α and concentrate on the analysis of the four consecutive vectors in the contour of the set of allocated vectors, i.e., $Z^{(i-1,j-1)}$, $Z^{(i,j-1)}$, $Z^{(i-1,j)}$, $Z^{(i,j)}$. If α increases, the relative positions of $\mathcal{C}^{(1,i)}$ and $\mathcal{C}^{(2,j)}$ will change, which leads to transitions of allocated vectors as illustrated in Figure 4. By analyzing all the possible transitions with $\alpha \geq 1$, we can conclude that the new allocated rate vector $Z^{(i-1,j)}$ is always below contour of the set of allocated vectors before α increases. In particular, if $\partial\mathcal{R}_1, \partial\mathcal{R}_2$ are hyperplanes in $\mathbb{R}_+^2$, the above arguments hold for all $\alpha > 0$. $\square$

The following corollary analyzes the stability of allocations close to the maximum throughput allocation, where the stability region converges to the maximum stability region $\overline{\mathcal{R}}$

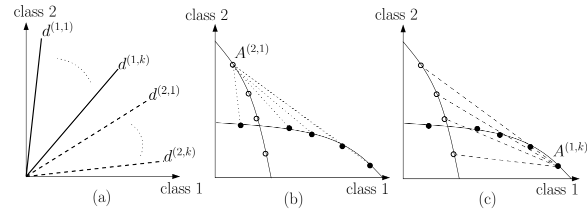

Figure 5: When $\alpha \to 0$ (a) the cones and boundaries; (b) and (c) the allocated rate vectors in $\mathcal{R}_1^{(k)}$ and $\mathcal{R}_2^{(k)}$, respectively.

defined in (12) for sufficiently small α.

COROLLARY 3. *The maximum stability region $\overline{\mathcal{R}}$ is achieved as $\alpha \to 0$, i.e., $\overline{\mathcal{R}_\alpha} \to \overline{\mathcal{R}}$.*

PROOF. Following the proof of Corollary 2, we assume $a_2 > a_1$. When $\alpha \to 0$, $d^{(1,i)}$ tends to the tangent of the angle between the line segment connecting $A^{(1,i)}, A^{(1,i+1)}$ and class 1 axis, and $d^{(2,i)} \to \infty$. Figure 5.(a) shows the relative positions of $\mathcal{C}^{(1,i)}$ and $\mathcal{C}^{(2,i)}$ in this case. Then the allocated rate vectors are $A^{(1,k)} + A^{(2,i)}$ and $A^{(1,i)} + A^{(2,1)}$ for all $1 \leq i \leq k$, as shown in Figure 5.(b), (c). When $a_2 > a_1$, by convexity of $\mathcal{R}_1$, $\mathcal{R}_2$, this also implies that when $h_k \to 0$, the allocated vectors when $\alpha \to 0$ are on the boundary of $\overline{R} = \mathcal{R}_1 + \mathcal{R}_2$. This finally shows that the maximum stability region $\overline{R}$ can be achieved as $\alpha \to 0$. $\square$

In fact, we will see in Section 6 that in some cases, there exists $0 < \alpha_0 < 1$ such that when $\alpha < \alpha_0$, the maximum stability is achieved, i.e., $\overline{\mathcal{R}_\alpha} = \overline{\mathcal{R}}$.

We conclude this subsection by presenting possible cases where with time-varying rate region, the stability region of α-fair allocations can still be insensitive to α. This can be the case when all flow classes experience the same capacity variations. This special case for insensitivity is true for a system with arbitrary number of classes, if for all $i \in \mathcal{I}$, there exists a constant c_i such that $\mathcal{R}_i = c_i\mathcal{R}_1$, $\overline{\mathcal{R}_\alpha} = \sum_{i \in \mathcal{I}} \pi_i c_i \mathcal{R}_i$. This is because when solving (1) with different rate regions $\mathcal{R}_i$ where $\mathcal{R}_i = c_i\mathcal{R}_1$, by scaling the decision variable $\boldsymbol{\phi}^{(j,\alpha)}$ with the same constant c_i, we have $\boldsymbol{\phi}^{(i,\alpha)} = c_i \boldsymbol{\phi}^{(1,\alpha)}$. An example of such systems is the downlink of a cell in a wireless network where the power of the base station allocated to data traffic may vary because of the presence of high-priority traffic such as voice.

As discussed in this section, when the rate region is nonconvex or time-varying, the stability region of a resource allocation scheme depends on the chosen fairness parameter α. In the cases we studied, fairness can be imposed only at the expense of reducing the stability region. Then in a number of practical networks where this fairness stability trade-off exists, it becomes crucial to choose a fairness objective that achieve the right balance between fairness and performance.

6. EXAMPLES

In this section, we present some numerical experiments to illustrate the analytical results derived in the previous sections on various types of data networks: wired networks and wireless networks with centralized or distributed resource allocation. The sensitivity of the stability region of α-fair allocations to the fairness parameter α strongly depends on the considered network. We also observe that the sensitivity is usually much higher for wireless networks than for wired networks due to the sharp variation of rate regions.

6.1 Wired networks with link failures

In this subsection, we investigate time-varying rate regions in wired networks due to link failures. The different set of time-varying broken links generate various *link failure states*, which in turn defines time-varying rate regions. We study two different cases depending on the underlying routing and flow management mechanism: (i) multi-path routing without flow splitting, and (ii) multi-path routing with flow splitting.

A wired network is represented as a set of L links and K routes where each route k is defined as a subset r_k of the set of links $\{1, \ldots, L\}$. Let $C = \{C_1, \ldots, C_L\}$ be the capacity vector with $C_l > 0$. We refer to the routing matrix as the $K \times L$-dimensional matrix R whose k, l-entry is equal to 1 if $l \in r_k$, and 0 otherwise. The routing matrix varies with link failure states, and we denote by R_i the routing matrix in link failure state i.

6.1.1 Multi-path routing without flow splitting

We now assume that each class is assigned a set of routes for each link failure state i, but at any instant of time, each class s can choose only a single route in the subset of routes $m_s(i)$. We let $\mathcal{M}_i$ be the set of $S \times K$ stochastic matrices such that on each row s, the s, k-entries are equal to 0 for all k except in the set $m_s(i)$. Each matrix $M \in \mathcal{M}$ corresponds to a particular route choice. We let $\boldsymbol{\phi}$ denote a row vector, then the rate region in the link failure state i is the convex hull of the capacity sets associated with the routing matrices $M \in \mathcal{M}_i$, given by:

$$\mathcal{R}_i \quad = \quad \text{convex hull of } \{\boldsymbol{\phi} : \exists M \in \mathcal{M}_i, \boldsymbol{\phi} M R_i \leq C\}.$$

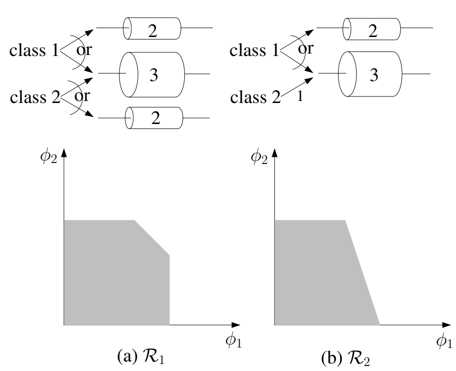

Figure 6: A wired network with link failures: multi-path routing without flow-splitting.

Consider the example of Figure 6, where $C = (2, 3, 2)$, and

$$R_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \; R_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\mathcal{M}_1 = \left\{ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \right\},$$

$$\mathcal{M}_2 = \left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} \right\}.$$

Figure 8.(a) shows the stability regions for different values of α when $\pi_1 = \pi_2 = 1/2$. We observe that the stability region decreases as α increases. However the sensitivity to α is rather limited, and when $\alpha < 0.5$, the maximum stability region is achieved.

6.1.2 Multi-path routing with flow splitting

Suppose now that for link failure state i, each class s can use all routes in the set $m_s(i)$ at the same time. Abusing the notation, we again let $\mathcal{M}$ be the set of $S \times K$ stochastic matrices such that on each row s, the s, k-entries are equal to 0 for all k except those in the set $m_s(i)$. Each matrix $M \in \mathcal{M}_i$ corresponds to a particular traffic splitting scheme at the failure state i. Then, the rate region is given by:

$$\mathcal{R}_i \quad = \quad \{\boldsymbol{\phi} : \exists M \in \mathcal{M}_i, \boldsymbol{\phi} M R_i \leq C\}.$$

Consider the example in Figure 7 with three links, where $C = (3, 2, 3)$, and

$$R_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \; R_2 = R_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\mathcal{M}_1 = \left\{ \begin{pmatrix} 2/3 & 1/3 & 0 \\ 0 & 1/3 & 2/3 \end{pmatrix} \right\}, \mathcal{M}_2 = \left\{ \begin{pmatrix} 2/3 & 1/3 & 0 \\ 0 & 1 & 0 \end{pmatrix} \right\},$$

$$\mathcal{M}_3 = \left\{ \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1/3 & 2/3 \end{pmatrix} \right\}.$$

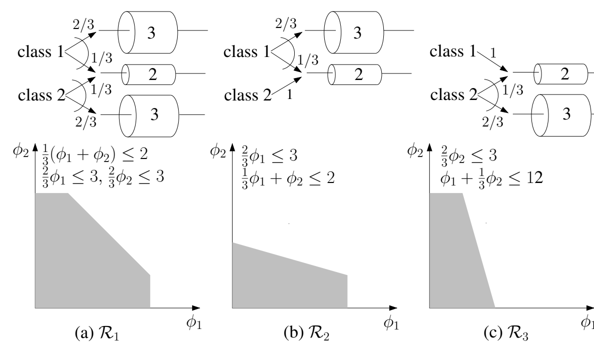

Figure 7: A wired network with link failures: multi-path routing with flow-splitting.

Figure 8.(b) shows the change of stability regions for different values of α, where we assume that $\pi_1 = \pi_2 = \pi_3 = 1/3$. As illustrated, the sensitivity to α is more substantial compared to the case without flow splitting. When $\alpha < 0.2$, the maximum stability region is achieved.

6.2 Wireless cellular networks with random interference

We consider the downlink of a cell covered by base station (BS) 1. BS 1 serves two classes of flows generated by some users with fixed positions, as shown in Figure 9. The rate regions of the system are time-varying due to variations of the interference generated by BS 2 and 3. To simplify the analysis, we assume that BS 2 and 3 cannot be active at the same time, and when active they transmit at full and fixed power. The system is symmetric, and when BS 2 (resp. 3) is on and BS 3 (resp. 2) is off, the noise plus interference

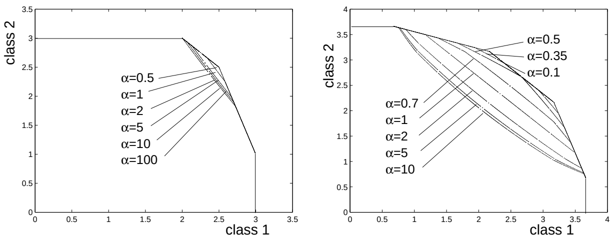

(a) No flow-splitting (b) Flow-splitting

Figure 8: Stability regions: multi-path routing.

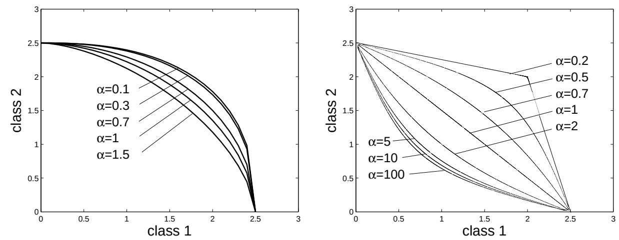

(a) CD (b) TD

Figure 10: Stability regions for wireless cellular networks with random interference.

at the position of class-1 (resp. class-2) users is 12 dB and that for class-2 (resp. class-1) users is 0 dB. The downlink resources are shared at the BS 1 either according to a successive decoding scheme for the broadcast channel (Code Division, denoted by CD) or to a time division multiple access scheme (denoted by TD). For CD channels, the rate region is given as,

$$\phi_1 \leq \log_2 (1 + P_1/\eta_1),$$
$$\phi_2 \leq \log_2 (1 + P_2/\eta_2).$$

where P_i and η_i are the downlink transmission and noise power for each class $i = 1, 2$, respectively, and $P = P_1 + P_2$ is the maximum transmission power of the base station. The corresponding rate regions are presented in Figure 9.

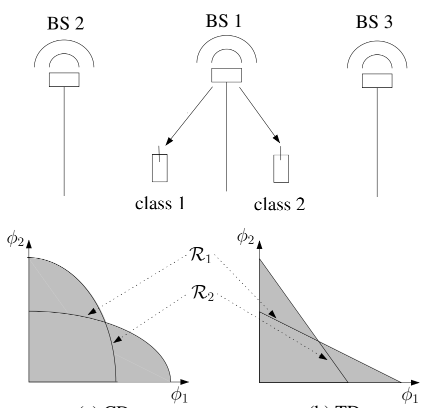

(a) CD (b) TD

Figure 9: Wireless cellular network with random interference.

When BS 1 allocates its full power to users of class 1 (or 2), the corresponding flows are served at rate 4 or 1 depending on the activities of BS 2 and 3. Now the stability region for different α is shown in Figure 10. The maximum stability region is achieved when $\alpha < 0.1$ or $\alpha < 0.2$ for CD or TD systems, respectively. We also observe that the sensitivity of α is very significant in this case due to the sharp variation of rate regions in wireless network.

6.3 Random access in wireless networks

We conclude this section by an example of a network with non-convex but fixed rate region. For this example, the results of Section 3 allow us to exactly characterize the stability region of some α-fair allocation.

The model is similar to that considered in [17, 31], and may represent typical WLANs or multi-hop wireless networks with random access Aloha-type algorithm. The network is a collection $\mathcal{L}$ of L wireless links. We consider L flow classes, flows of class l use link l only. The links interact through interference. We assume that a transmission on link l can be successful only if none of the neighboring links is transmitting, otherwise there is a collision. Denote by $\mathcal{L}_l$ the set of links interfering link l. Time is slotted and packet transmissions exactly last one slot. At the beginning of each slot, links try to access the channel in a distributed manner, each link l transmits with probability p_l. The rate of link l, and then of class-l flows, is given by:

$$\phi_l = p_l \prod_{k \in \mathcal{L}_l} (1 - p_k). \tag{16}$$

The rate region of this system is in general non-convex and given by:

$$\mathcal{R} = \{\boldsymbol{\phi} : \forall l, \phi_l \leq p_l \prod_{k \in \mathcal{L}_l} (1 - p_k), p_k \in (0, 1)\}. \tag{17}$$

It is shown that for Proportional fair allocation with non-convex rate region (17), the transmission probabilities at state $\boldsymbol{N}$ are:

$$p_l = \frac{N_l}{N_l + \sum_{k: l \in \mathcal{L}_k} N_k}, \tag{18}$$

and by varying the network state $\boldsymbol{N}$ that the conditions of Subsection 3.3 are satisfied, namely that the set of the allocated rate vectors is continuous and actually equal to the entire rate region in this case. Therefore we can conclude that the stability region of the Proportional fair allocation is exactly equal to the largest open subset of $\mathcal{R}$.

7. CONCLUSION AND FUTURE WORK

In practical networks, the rate region that constrains the resource allocation may not follow the standard assumptions of convexity and time-invariance. The characterization of stability region becomes more challenging for either non-convex or time-varying rate regions, and its size and shape become dependent on the chosen resource allocation policy. In this paper, we have studied this dependence for a large class of utility-based allocations, the α-fair allocations.

For networks with an arbitrary number of classes and with fixed and non-convex rate region, we have given sufficient and necessary conditions for flow-level stability of α-fair allocations, for all $\alpha > 0$. We then extended the analysis to networks with time-varying and convex rate region, for which we have characterized the stability condition of α-fair allocations, for all $\alpha > 0$. We have also studied the sensitivity of the stability region of α-fair allocations to the fairness parameter α, and demonstrated an intriguing tradeoff between fairness and flow-level stability. This tradeoff raises further questions on how to choose a resource allocation policy.

8. REFERENCES

[1] M. Armony and N. Bambos. Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing Systems: Theory and Applications*, 44(3):209–252, 2003.

[2] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi. CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users. *IEEE Comm. Mag.*, 38(4):70–77, 2000.

[3] T. Bonald, S. Borst, N. Hegde, and A. Proutière. Wireless data performance in multicell scenarios. In *Proceedings of ACM Sigmetrics/Performance*, 2004.

[4] T. Bonald, S. Borst, and A. Proutière. How mobility impacts the flow-level performance of wireless data networks. In *Proceedings of IEEE Infocom*, 2004.

[5] T. Bonald and L. Massoulié. Impact of fairness on internet performance. In *Proceedings of ACM Sigmetrics/Performance*, 2001.

[6] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing Systems: Theory and Applications*, 53(1-2):65–84, 2006.

[7] T. Bonald and A. Proutière. Flow-level stability of utility-based allocations for non-convex rate regions. In *Proceedings of the 40th Conference on Information Sciences and Systems*, March 2006.

[8] S. Borst, N. Hegde, and A. Proutière. Capacity of wireless networks with intra- and inter-cell mobility. In *Proceedings of IEEE Infocom*, March 2006.

[9] S. Borst, L. Leskela, and M. Jonckheere. Stability of parallel queueing systems with coupled rates, 2007. Submitted.

[10] M. Bramson. Stability of networks for max-min fair routing. In *Presentation at the 13th INFORMS Applied Probability Conference*, 2005.

[11] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE*, 95(1):255–312, January 2007.

[12] M. Chiang, D. Shah, and A. Tang. Stochastic stability of network utility maximization: General file size distribution, 2006. Submitted to *IEEE Transactions on Information Theory*. Partially appeared in *Proceedings of Allerton Conference*, Sept., 2006.

[13] J. G. Dai. On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *Annals of Applied Probability*, 5:49–77, 1995.

[14] G. de Veciana, T. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Transactions on Networking*, 1:2–14, 2001.

[15] V. Gambiroza, B. Sadeghi, and E. W. Knightly. End-to-end performance and fairness in multihop wireless backhaul networks. In *Proceedings of ACM Mobicom*, 2004.

[16] H. C. Gromoll and R. Williams. Fluid limit of a network with fair bandwidth sharing and general document size distribution, 2006. Preprint.

[17] P. Gupta and A. Stolyar. Optimal throughput allocation in general random-access networks. In *Proceedings of the 38th Conference on Information Sciences and Systems*, March 2006.

[18] N. Hegde and A. Proutière. Packet and flow level performance of wireless multihop networks. In *Proceedings of IEEE Globecom*, 2006.

[19] M. Jonckheere and S. Borst. Stability of multi-class queueing systems with state-dependent service rates. In *Proceedings of IEEE Value Tools*, October 2006.

[20] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, volume 49, 1998.

[21] A. Lakshmikantha, C. L. Beck, and R. Srikant. Connection level stability analysis of the internet using the sum of squares (SOS) techniques. In *Proceeding of CISS*, 2004.

[22] X. Lin, N. B. Shroff, and R. Srikant. On the connection-level stabiltiy of congestion-controlled communication networks, 2005. submitted to *IEEE Transaction on Information Theory*.

[23] 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, 2007. Technical report, Princeton University.

[24] H. Luo, S. Lu, and V. Bharghavan. A new model for packet scheduling in multihop wireless networks. In *Proceedings of Mobicom*, 2000.

[25] W. Luo and A. Ephremides. Stability of n interacting queues in random-access systems. *IEEE Transactions on Information Theory*, 45(5):1579–1587, 1999.

[26] L. Massoulié. Structural properties of proportional fairness: Stability and insensitivity, 2006. Submitted.

[27] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Transactions on Networking*, 10(3):320–328, 2002.

[28] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000.

[29] W. Szpankowski. Stability conditions for some multi-queue distributed systems: Buffered random access systems. *Annals of Applied Probability*, 26:498–515, 1994.

[30] A. Tang, J. Wang, and S. Low. Counter-intuitive throughput behaviors in networks under end-to-end control. *IEEE/ACM Transactions on Networking*, 14(2):355–368, 2006.

[31] X. Wang and K. Kar. Cross-layer rate control for end-to-end proportional fairness in wireless networks with random access. In *Proceedings of MobiHoc*, 2005.

[32] H. Ye. Stability of data networks under optimization-based bandwidth allocation. *IEEE Transactions on Automatic Control*, 48(7):1238–1242, 2003.

[33] H. Ye, J. Ou, and X. Yuan. Stability of data networks: Stationary and bursty models. *Operations Research*, 53:107–125, 2005.