RC25476 (IRE1404-016) June 18, 2014
Computer Science

IBM Research Report

Nonhomogeneous Place-Dependent Markov Chains, Unsynchronised AIMD, and Network Utility Maximization

Fabian Wirth[1], Sonja Stuedli[2], Jia Yuan Yu[1], Martin Corless[3], Robert Shorten[1]

[1]IBM Research
Smarter Cities Technology Centre
Mulhuddart
Dublin 15, Ireland

[2]Newcastle University
Newcastle, Australia

[3]Purdue University
West Lafayette, IN USA

Nonhomogeneous Place-Dependent Markov Chains, Unsynchronised AIMD, and Network Utility Maximization

Fabian Wirth* Sonja Stuedli† Jia Yuan Yu* Martin Corless‡
Robert Shorten*

May 21, 2014

Abstract

In this paper we derive a convergence result for the non-homogeneous Markov chain that arises in the study of networks employing the *additive-increase multiplicative decrease* (AIMD) algorithm. We then use this result to solve the network utility maximization (NUM) problem. Using AIMD, we show that the NUM problem is solvable in a very simple manner using only intermittent feedback, no inter-agent communication, and no common clock.

Keywords: AIMD; Nonhomogeneous Markov Chains; Invariant Measure; Iterated Function Systems; Almost Sure Convergence; NUM with intermittent feedback

1 Introduction

Recent developments in the context of Smart Grid, Smart Transportation, and the internet have given rise to a rich set of optimisation problems in which a number of agents collaborate to achieve a social optimum [10, 37, 14, 29]. For example, collaborative cruise control systems are emerging in which a group of vehicles on a stretch of road share information to determine a speed limit that minimises fuel consumption subject to some constraint (traffic flow, pollution constraints, etc.). Other examples of this problem can be found in several application domains; in the energy literature [40, 18, 4]; in the electric vehicles literature [21, 11, 17, 8], in distributed load control [29]; in the study of control strategies for thermostatically controlled loads, as refrigerators or air conditioners [1, 41, 34], and of course, in the optimization literature itself [45, 26].

Roughly speaking the optimisation problems that emerge in such applications are simple to solve. Typically, one wishes to minimise a sum of strictly convex functions of a single variable subject to a linear, or perhaps, polynomial constraint. It is well known that such problems can be readily solved by a multitude of methods in a convex optimisation framework [7]. Notwithstanding this fact, solving these problems in a smart grid or smart transportation framework is challenging. The difficulties that arise in such environments are due to several factors.

The first is due to the fact that one wishes to find solutions which can be implemented with minimal communication (or even no) between individual agents, and between the agents and infrastructure. This need arises due to the fact that many of these problems are massively large scale in nature, and continuous inter-agent communication would place an undue burden on the telecommunications infrastructure [29] and due to privacy considerations.

The second difficulty is that the number of agents participating in the optimisation problem is large and time-varying, and each agent's utility is private and not communicated to other agents. An additional

*IBM Research Ireland, Damastown Industrial Estate, Mulhuddart, Dublin 15, Ireland, ({fabwirth,jiayuanyu,robshort}@ie.ibm.com).
†Newcastle University, Newcastle, Australia (sonja.stuedli@uon.edu.au)
‡School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN, USA (corless@purdue.edu)
This work is supported in part by the EU FP7 project INSIGHT under grant 318225.

further difficulty arises because agents in such applications typically have limited actuation capabilities, i.e. limited capabilities of effecting a change in their state. For example, in *internet of things applications* agents can often only influence their behaviour by switching themselves on or off. Thus, distributed algorithms for solving large scale optimisation problems in which agents with limited actuation capabilities collaborate to achieve a common goal, is a highly topical research problem.

Our objective in this paper is to develop algorithms that can be deployed in such situations. To this end, consider a network of n agents, each with a state x_i, representing an amount of allocated resource, and a cost function $f_i(\bar{x}_i)$, where $\bar{x}_i$ denotes the average value of the ith state x_i. Given this basic set up we wish to minimize:

$$\sum_{i=1}^{n} f_i(\bar{x}_i) \tag{1}$$

subject to the capacity constraint $\sum_{i=1}^{n} \bar{x}_i = C$ and for all i, $\bar{x}_i \geq 0$. Further, we wish to do this with as little inter-agent communication, and with a minimum amount of centralized actuation. As we shall see inter-agent communication is not necessary at all to achieve convergence to the optimum, nor is it necessary to communicate a feedback signal in the form of a multiplier. Rather, we will show that it is sufficient to provide the centralized, one-bit information, that the constraint has been reached. Furthermore, the conditions for convergence are independent of network dimension, depending only on the worst agent in the system.

All this can be achieved using only the *additive-increase multiplicative decrease* (AIMD) algorithm. Recall that AIMD is an algorithm continuously claim more and more of available resource in a gentle fashion until a notification (a congestion notification) is sent to them that the aggregate amount of available resource has been exceeded. This is the additive increase (AI) phase of the algorithm. Then then reduce their demand on resource by a factor between zero and one. This is the multiplicative decrease (MD) of the algorithm. The AI phase of the algorithm then restarts immediately.

To motivate this result, we make use of the following two known observations.

Observation 1 (Consensus) : The optimization problem may be formulated in a Lagrangian framework as follows [45]. We introduce the Lagrange parameter $\mu \in \mathbb{R}$ and consider

$$H(\bar{x}_1, ..., \bar{x}_N, \mu) = \sum_{i=1}^{n} f_i(\bar{x}_i) - \mu \left(\sum_{i=1}^{n} \bar{x}_i - C \right). \tag{2}$$

From the Karush-Kuhn-Tucker (KKT) conditions [7, Section 5.5.3], the following necessary and sufficient condition for optimality can be obtained by setting all partial derivatives to zero. In the optimal point $x^* \in \mathbb{R}^n$, $\mu^* \in \mathbb{R}$ we have

$$\mu^* = \frac{\partial f_i}{\partial \bar{x}_i}(x_i^*) \qquad \forall\, i = 1, \dots, N. \tag{3}$$

(subject to certain assumptions on the utility functions and constraints). In other words, the system is at optimality when the derivatives of the utility functions are in consensus.

Observation 2 (Ergodic behaviour) : It follows from the results in [42], under the assumptions of ergodicity, that the ergodic limit of a network of AIMD flows is almost surely of the form

$$\lim_{k \to \infty} \frac{1}{k} \sum_{\ell=1}^{k} x_i(\ell) = \frac{\Theta}{\lambda_i^*}, \tag{4}$$

where Θ is a network-specific constant and λ_i^* is the steady-state probability that the ith AIMD agent responds to a congestion notification.

With these two observations in mind, we can now aim to choose place-dependent probability functions $\lambda_i(\cdot)$ so that the equation for the steady state behaviour (4) is equivalent to the KKT condition (3). Suppose that each agent responds to congestion with probability

$$\lambda_i(\bar{x}_i(k)) \;=\; \Gamma\, \frac{f_i'(\bar{x}_i(k))}{\bar{x}_i(k)}\,, \tag{5}$$

where $\bar{x}_i(k)$ is the average of the last k values of x_i. Here Γ is again a network wide constant chosen to ensure that $0 < \lambda_i(\bar{x}_i) < 1$. Suppose now that $\bar{x}_i(k) \approx x_i^*$. Then we can write $\lambda_i(\bar{x}_i(k)) \approx \lambda_i^*$. Provided that for this choice (4) holds, we obtain for large k

$$\lambda_i(\bar{x}_i(k)) \;\approx\; \Gamma\frac{f_i'(\bar{x}_i(k))}{\Theta}\lambda_i(\bar{x}_i(k)) \tag{6}$$

and so $f_i'(\bar{x}_i(k)) \approx \Theta/\Gamma \approx f_j'(\bar{x}_j(k))$ for all i, j. These are precisely the KKT conditions which in many cases are both necessary and sufficient for optimality.

The purpose of this paper is to show that the above intuition is true. Specifically, with the place-dependent probabilities $\lambda_i(\cdot)$ chosen as in (5), we do indeed have $\bar{x}_i(k) \approx x_i^*$ for large k. Consequently, the *AIMD* algorithm can be modified to solve distributed optimisation problems in asynchronous environments in a manner that is both effective and efficient in terms of communication overhead.

Our paper is structured as follows. We begin by reviewing a recently proposed switched systems model of AIMD dynamics and known results on the stochastic stability of this model for fixed probabilities. The main result for fixed probabilities is that the long term averages converge almost surely and that this limit can be expressed analytically. We start in Section 2 by introducing notation and recalling some facts about the dynamics of stochastic AIMD algorithms. In Section 3 we present a discussion of a stochastic AIMD algorithm that solves the Network Utility Maximization problem. In Section 4 we introduce two dynamical systems, representing these algorithms. These differ in the choice of the probability laws. We then state the main convergence results. Related works are discussed in Section 5. In Section 6 we apply the results to solve the NUM problem for a network of agents. The main proofs are provided in the Appendix. There, we give intermediate results that link the fixed probability case with the place-dependent case. Specifically, we ask to what degree may the place dependent case be approximated with the fixed probability case, and over which time intervals. To this end we study the robustness properties of a deterministic system that iterates on the expectation operator. These results are then used to establish the main result of the paper.

2 Preliminaries

Our starting point is the suite of algorithms that underpin the Transmission Control Protocol (*TCP*) that is used in internet congestion control. A fundamental building block of TCP is the AIMD algorithm. To discuss AIMD in a formal setting some preliminaries are necessary.

1. Notation : The vector space of real column vectors with n entries is denoted by $\mathbb{R}^n$ with elements $x = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix}^\top$, where $x^\top$ denotes the transpose of x. For $x, y \in \mathbb{R}^n$, we write $x \gg y$ if $x_i > y_i$ for all $i = 1, \dots, n$. The convex hull of a set X is denoted by $\operatorname{conv} X$; it may be defined as the smallest convex set containing X.

We denote the canonical basis vectors in $\mathbb{R}^n$ by $e_i, i = 1, \dots, n$ and let $e := \sum_{i=1}^n e_i$. The standard 1-norm is defined by $\|x\|_1 = \sum_{i=1}^n |x_i|$, $x \in \mathbb{R}^n$. The closed ball of radius δ around 0 with respect to this norm is denoted by $\overline{B}_1(0, \delta)$. The distance of a point x to a set Z with respect to the 1-norm is denoted by

$$\operatorname{dist}_1(x, Z) := \min\{\|x - z\| \;;\; z \in Z\}\,.$$

The standard simplex Σ in $\mathbb{R}^n$ is defined by

$$\Sigma := \left\{ x \in \mathbb{R}^n_+ \;\middle|\; \sum_{i=1}^n x_i = 1 \right\} .$$

We will write Σ_n if we want to emphasize that we are working in $\mathbb{R}^n$. Note that we are only interested in dynamics on Σ. Thus when we write $\overline{B}_1(0,\delta)$ we will tacitly assume that we consider the intersection of this ball with Σ. The relative interior of Σ is defined by $\operatorname{ri}\Sigma := \{x \in \Sigma \mid x_i > 0, i = 1,\dots,n\}$. It will be sometimes useful to use the Hilbert metric $d_H(\cdot,\cdot)$ on $\operatorname{ri}\Sigma$, [22]. Recall that it is given by

$$d_H(x,y) := \max_i \log(x_i/y_i) - \min_j \log(x_j/y_j)\,, \quad x,y \in \operatorname{ri}\Sigma\,,$$

and makes $(\operatorname{ri}\Sigma, d_H)$ a complete metric space. A ball of radius δ with respect to the Hilbert metric is denoted by $B_H(x,\delta)$; again without further notice, we will understand that $B_H(x,\delta)$ is the ball contained in Σ. For the sake of analysis it is sometimes easier to work with the logarithm removed, in which case we consider

$$e^{d_H}(x,y) := \frac{\max_i x_i/y_i}{\min_j x_j/y_j}\,, \quad x,y \in \operatorname{ri}\Sigma\,.$$

Note that for $x_k, x, y, z \in \operatorname{ri}\Sigma$ we have $\|x_k - y\|_1 \to 0$ if and only if $d_H(x_k,y) \to 0$ which is again equivalent to $e^{d_H}(x_k,y) \to 1$. Furthermore, $d_H(x,y) < d_H(z,y)$ if and only if $e^{d_H}(x,y) < e^{d_H}(z,y)$.

2. AIMD algorithms and stochastic matrices : We have already mentioned that the AIMD algorithm underpins TCP. We shall not describe the TCP algorithm here. Rather, we refer the interested reader to [30, 31, 44, 43] for details of TCP and note that the dynamics of networks of AIMD flows can be described as

$$x(k+1) = A(k)x(k), \tag{7}$$

where $A(k)$ is a stochastic matrix. The matrices $A(k)$ belong to a compact set of matrices $\mathcal{A}$, which we now describe. Given two vectors $\alpha \in \operatorname{ri}\Sigma_n, \beta \in (0,1)^n$ we define a set $\mathcal{A}$ of 2^n matrices as follows. Let

$$B := \left\{ \tilde\beta \in \mathbb{R}^n \;\middle|\; \tilde\beta_i \in \{\beta_i, 1\}\,, \quad i = 1,\dots,n \right\} ,$$

which is clearly a set with 2^n elements. The *set of AIMD matrices* is then given by

$$\mathcal{A} := \left\{ \operatorname{diag}(\tilde\beta) + \alpha(e-\tilde\beta)^\top \;\middle|\; \tilde\beta \in B \right\} . \tag{8}$$

Note that $\alpha(e-\tilde\beta)^\top \in \mathbb{R}^{n\times n}$ as $\alpha \in \mathbb{R}^{n\times 1}$ and $(e-\tilde\beta)^\top \in \mathbb{R}^{1\times n}$. Such matrix sets and the dynamics of Markov chains on Σ defined by $\mathcal{A}$ have been studied in [48, 42]. In the following we assume that some enumeration of the matrices in $\mathcal{A}$ has been fixed and we will refer to the matrices $A_j \in \mathcal{A}, j = 1,\dots,2^n$. We also use the convention that the matrix A_1 is defined using β that is

$$A_1 = \operatorname{diag}(\beta) + \alpha(e-\beta)^\top . \tag{9}$$

Note that A_1 is a column stochastic, positive matrix. In particular, the eigenvalue $1 = r(A)$ is simple. It is easy to see that a corresponding positive (Perron) eigenvector is given by

$$z = \begin{bmatrix} \frac{\alpha_1}{1-\beta_1} & \cdots & \frac{\alpha_n}{1-\beta_n} \end{bmatrix}^\top . \tag{10}$$

AIMD matrices have the property that they leave the subspace

$$V := \{x \in \mathbb{R}^n \mid e^\top x = 0\}$$

invariant, as they are column stochastic. Finally, we recall the following fact about the contractive properties of A_1 from [48]. In the following statement $A_{1|V}$ denotes the restriction of A_1 to the invariant subspace V.

Lemma 2.1 *Let $\alpha \in \operatorname{ri}\Sigma_n, \beta \in (0,1)^n$ and let $\mathcal{A}$ be the corresponding set of AIMD matrices. Then for all $A_j \in \mathcal{A}$ we have $\|A_{j|V}\| \leq 1$. Also there exists a constant $c \in (0,1)$ such that for the matrix A_1 be defined by* (9) *we have*

$$\|A_{1|V}\|_1 =: c < 1.$$

3. Elementary results on AIMD : The AIMD algorithm is often studied under the assumption that probabilities are not place-dependent; namely, the probability that $A(k) = A_i$ is independent of k and $x(k)$. It shall be useful to refer to this case in the remainder of the paper and we briefly recall relevant known results here. Consider a probability distribution $p : \mathcal{A} \to [0,1]$ on the set $\mathcal{A}$ of AIMD matrices. This induces a Markov chain on Σ by setting

$$x(k+1) = A(k)x(k)\,, \quad k \in \mathbb{N}\,, \quad x(0) = x_0 \in \Sigma\,, \tag{11}$$

where $\mathbb{P}(A(k) = A_j) = p_j$ for $A_j \in \mathcal{A}$. In particular, the sequence of transition matrices $\{A(k)\}$ is iid. In the sequel, we will consider the case in which the probabilities p_j are derived from individual drop probabilities λ_i of the agents. In this case, λ_i is the probability that in (8) we have $\tilde{\beta}_i = \beta_i$. In this case, for every $\tilde{\beta} \in B$ we have

$$\mathbb{P}\Big(A(k) = \operatorname{diag}(\tilde{\beta}) + \alpha(e - \tilde{\beta})^\top\Big) = \prod_{\tilde{\beta}_i = \beta_i} \lambda_i \prod_{\tilde{\beta}_i = 1} (1 - \lambda_i)\,. \tag{12}$$

For the Markov process defined by (11) it is known from the results in [48, 42] that if $\lambda = \begin{bmatrix}\lambda_1 & \dots & \lambda_n\end{bmatrix} \gg 0$, then there is a unique, invariant, probability measure π on Σ for the Markov chain, and for every initial condition $x_0 \in \Sigma$ we have almost surely that

$$\lim_{k\to\infty} \frac{1}{k+1} \sum_{t=0}^{k} x(t; x_0) = \frac{1}{\sum_{\ell=1}^{n} \frac{\alpha_\ell}{\lambda_\ell(1-\beta_\ell)}} \begin{bmatrix} \frac{\alpha_1}{\lambda_1(1-\beta_1)} \\ \vdots \\ \frac{\alpha_n}{\lambda_n(1-\beta_n)} \end{bmatrix} =: P_\lambda\,, \tag{13}$$

where $x(\cdot; x_0)$ denotes samples paths satisfying the initial condition $x(0; x_0) = x_0$. As almost sure convergence implies convergence in probability this shows that for every $\varepsilon, \delta > 0$ there exists a k_0 such that for all $k \geq k_0$ we have

$$\mathbb{P}\left(\left\|\frac{1}{k+1} \sum_{t=0}^{k} x(t; x_0) - P_\lambda\right\|_1 > \delta\right) < \varepsilon\,. \tag{14}$$

It will also be useful to have a uniform version of (14). To this end we define the linear averaging operator

$$\overline{S}(k) := \frac{1}{k+1} \sum_{t=0}^{k} A(t-1) \cdots A(0)\,,$$

with the interpretation that the summand corresponding to $t = 0$ is the identity I_n. The interest in this expression lies on the observation that for $x_0 \in \Sigma$ we have

$$\frac{1}{k+1} \sum_{t=0}^{k} x(t; x_0) - P_\lambda = (\overline{S}(k) - P_\lambda e^\top) x_0\,. \tag{15}$$

Lemma 2.2 *Consider the Markov chain* (11) *and let $\lambda \gg 0$ as defined by* (12) *be fixed. Then for every $\varepsilon, \delta > 0$ there exists a $k_0 \in \mathbb{N}$ such that for all $k \geq k_0$*

$$\mathbb{P}\left(\left\|\overline{S}(k) - P(\lambda)e^\top\right\|_1 > \delta\right) < \varepsilon\,. \tag{16}$$

In particular, there exists a $k_0 \in \mathbb{N}$ such that for all $k \geq k_0$ and all $x_0 \in \Sigma$ we have

$$\mathbb{P}\left(\left\|\frac{1}{k+1} \sum_{t=0}^{k} x(t; x_0) - P_\lambda\right\|_1 > \delta\right) < \varepsilon\,. \tag{17}$$

Proof As Σ contains the canonical basis vectors e_i, and the norm on $\mathbb{R}^{n\times n}$ induced by $\|\cdot\|_1$ is the column sum norm, it follows that for all $S \in \mathbb{R}^{n\times n}$

$$\|S\|_1 = \max_{i=1,\dots,n} \|Se_i\|_1 = \max\{\|Sx\|_1 \mid x \in \Sigma\}\,. \tag{18}$$

Fix $\varepsilon, \delta > 0$. By (14) and in view of (15), we may choose $k_0(e_i)$ such that for all $k \geq k_0(e_i)$ we have

$$\mathbb{P}(\|(\overline{S}(k) - P_\lambda e^\top)e_i\| > \delta) < \frac{\varepsilon}{n}\,, \quad i = 1,\dots,n\,.$$

By choosing $k_0 := \max_{i=1,\dots,n} k_0(e_i)$ we thus obtain for all $k \geq k_0$ that

$$\mathbb{P}\left(\ \left\|\overline{S}(k) - P_\lambda e^\top\right\|_1 > \delta\right) = \mathbb{P}\left(\max_{i=1,\dots,n}\ \left\|\left(\overline{S}(k) - P_\lambda e^\top\right)e_i\right\|_1 > \delta\right) < \varepsilon\,, \tag{19}$$

where we have used the standard estimate $\mathbb{P}\left(\cup_{i=1}^n W_i\right) \leq \sum_{i=1}^n \mathbb{P}\left(W_i\right)$ for events W_i. The claim (17) is now an immediate consequence of (18). $\square$

4. Some comments on stochastic convergence: The main result of this paper yields conditions for almost sure convergence of the sample paths of a Markov chain. For the benefit of the reader we briefly point to relevant parts of the literature, where this notion is discussed in detail. Readers familiar with notions of stochastic convergence may skip this section.

Given a Markov chain and an initial condition, we can consider the set of all possible sample paths $\{x(k; x_0, \omega) \mid \omega \in \Omega\}$, where Ω is an index set for the set of different sample paths or trajectories of the Markov chain. In our case, the set Ω can be identified with the set of all sequences with values in $\{1,\dots,m\}$, which can be interpreted as the set of sequences $\{A(0), A(1), \dots\}$ that lead to a particular sample path. Kolmogorov's existence theorem now states that the marginal probabilities that are induced by the Markov chain on finite-time intervals define a probability measure $\mathbb{P}_\Omega$ on Ω, the set of sample paths, see [5, Sections 2, 24, 36].

The statement that convergence to a limit happens almost surely, thus means that the measure of the set of sample paths which are converging is 1; with respect to the probability measure $\mathbb{P}_\Omega$ on the sample space. Thus almost sure convergence means that the convergence happens with probability one, with the right interpretation of the probability measure.

It is furthermore known, that almost sure convergence implies convergence in probability, [5]. The latter concept is implicitly defined in (17): for every $\varepsilon > 0$ and $\delta > 0$ there exists a k_0 such that for all $k \geq k_0$ the probability of begin further away from the limit than δ is smaller than ε.

3 AIMD Based Optimisation Algorithm

In this section, we formally define the class of distributed optimization problems that can be addressed using the algorithm presented in this paper. Recall: let $n \in \mathbb{N}$, $C > 0$ and $f_i : [0, C] \to \mathbb{R}$ be strictly convex and continuously differentiable, $i = 1,\dots,n$ and consider the optimization problem

$$\begin{aligned} \text{minimize} \quad & \sum_{i=1}^n f_i(\bar{x}_i) \\ \text{subject to} \quad & \sum_{i=1}^n \bar{x}_i = C\,, \quad \bar{x}_i \geq 0\,. \end{aligned} \tag{20}$$

We are interested in finding the optimal point x^* in which the minimum is achieved. It is well-known that by compactness of the feasible space an optimal solution exists; it is unique by the assumption of strict convexity. The unique, optimal point x^* is characterized by the KKT conditions which in our case are characterized by

$$\sum_{i=1}^n x_i^* = C, \quad x_i^* \geq 0, \quad f_i'(x_i^*) = \mu^*, \quad i = 1,\dots,n\,.$$

The actual implementation of the algorithm uses the current state of the users at a given time instance t, which we denote by $x_i(t)$ and the long-term average of the states of the users denoted by $\overline{x}_i(t)$. It is the aim of the algorithm to obtain convergence of the long-term averages to the KKT point x^*.

The actual algorithm implemented on each agent is based on the following assumptions. We assume that agents can infer when $\sum_{i=1}^n x_i \geq C$. If this is not the case we assume that each agent is informed of a constraint violation at the instant of its occurrence using binary feedback.[2] This is a congestion notification. Upon receipt of a congestion notification agent i updates the state $x_i(t)$:

$$x_i(t^+) = \beta x_i(t^-)$$

with probability $\lambda_i(\bar{x}_i(t^-)) = \Gamma \frac{f_i'(\bar{x}_i(t^-))}{\bar{x}_i(t^-)}$. Note that in this definition we implicitly assume that there exists a constant $\Gamma > 0$ such that $\lambda_i(\bar{x}_i) \in [0,1]$ for all values $\bar{x}_i \in [0, C]$. This restricts the admissible choices for the cost functions f_i. We will discuss different ways of treating more general functions in Remark 3.3.
At all other time instants the rate of change of $\dot{x}_i(t)$ is chosen to be a positive quantity. Note that the superscripts $+$ and $-$ denote the instants immediately prior and after a congestion notification respectively. This leads to the following discrete time algorithm that is implemented on each of the agents. We assume a common time step h is fixed and each agent i has an internal offset T_i. For the sake of abbreviation, we denote $k_i := T_i + kh, k \in \mathbb{N}$ and $x_i(k_i) := x_i(T_i + kh)$.

```
Initialization: Each agent sets its state x_i(0_i) to an arbitrary value;
                The parameter Γ is broadcast;
while agent i is active do
    if ∑_{ℓ=1}^n x_ℓ(k_i) < C then
        x_i((k+1)_i) = x_i(k_i) + α ;
    else
        x_i((k+1)_i) = βx_i(k_i) with probability λ_i(x̄_i(k_i)) = Γ f_i'(x̄_i(k_i)) / x̄_i(k_i) and
        x_i((k+1)_i) = x_i(k_i) + α otherwise;
    end
end
```

Algorithm 1: AIMD algorithm run by each agent

It is clear that the performance of the algorithm depends crucially on a number of assumptions. For example, we have assumed that the optimal point for our problem is in the interior of the simplex and that the time between sample points is the same for all agents (note that a common clock is not necessary). However, some other comments are required.

Remark 3.1 *The constant Γ is chosen to ensure that each $\lambda_i(\bar{x}_i) \in (0,1)$. Thus Γ depends on the worst utility function and must be communicated to all agents prior to the algorithms use. It is a network dependent quantity that is* **independent** *of network dimension.*

Remark 3.2 *In discrete time the algorithm could be implemented as follows: At each instant t_k all agents i perform the update*

$$x_i(t_{k+1}) = x_i(t_k) + \alpha h$$

except upon notification of congestion; at which point each agent i with probability $\lambda_i(\bar{x}_i(k))$ scales

$$x_i(t_{k+1}) = \beta x_i(t_k)\,.$$

We note that there is a small approximation error when formulating the above discrete-time model as an AIMD system.

We now briefly discuss how to reformulate NUM problems so that they satisfy the assumptions of our set-up. Note that the following list is not exhaustive.

[2]This is the intermittent feedback and the limited communication referred to throughout the paper.

Remark 3.3 *While the assumption that $\lambda_i(r)$ is well defined and in $[0,1]$ for all $r \in [0,1]$ might sound restrictive for the problem at hand, we note that the following modifications of the problem yield an feasible solution.*

(i) In case the objective functions are not increasing, we can define the constant

$$q := \max_{i=1,\dots,N} |f_i'(0)|$$

and consider the objective functions $\tilde{f}_i$, given by $\tilde{f}_i(r) = f_i(r) + qr$, $r \in [0,1]$, which are now strictly increasing. Note that this does not change the KKT point, as $\tilde{f}_i' \equiv f_i' + q$, so that the condition that all derivatives are equal is met at the same point x.

(ii) A second concern is that even if $f_i' \geq 0$ on $[0,1]$, the expression $f_i(r)/r$ might tend to ∞ as $r \to 0$, depending on the nature of the derivative of f_i at 0. In this case we may consider the definition

$$\lambda_i(r) := \min\left\{1, \Gamma \frac{f'(r)}{r}\right\}, \quad r \in [0,1]. \tag{21}$$

This amounts to a regularization of the optimization problem which we briefly outline in a simple situation. Assume that there is a unique point $r_\Gamma \in [0,1]$ such that $r_\Gamma = \Gamma f'(r_\Gamma)$. If λ_i in (21) were properly defined as in the corresponding objective function would be

$$\tilde{f}_i(r) := \begin{cases} \frac{1}{2} r^2 & r \in [0, r_\Gamma] \\ f_i(r) - f_i(r_\Gamma) + \frac{1}{2} r_\Gamma^2 & r \in [r_\Gamma, 1] \end{cases} \tag{22}$$

More generally, there could be several interlacing intervals, in which the condition $r_\Gamma < \Gamma f'(r_\Gamma)$ is satisfied or not. The important point here is that a decrease of Γ leads to a decrease of r_Γ, so that by choosing Γ small enough, the KKT point of the original problem will be found by the algorithm.

We note that although our analysis will be performed for a fixed number of agents, this is not necessary in the implementation of the algorithm. Indeed, as no information is required on the number of agents, these may join or drop out of the network at any time and the network will automatically readjust the KKT point given the new set of agents.

4 Convergence Analysis

In this section we discuss two versions of the stochastic algorithm for the approximation of the KKT point x^*. The common feature of these algorithms is that the probabilities for backing off depend on an average of past states. In the first version, we assume that there is a fixed window over which the average is taken, while in the second case the average is taken over the complete history starting at time $t_0 = 0$.

The two approaches are amenable to different methods of analysis. In the first case the problem may be recast in terms of a homogeneous Markov chain with state-dependent probabilities, sometimes also called an iterated function system (IFS). In this setting classical results ensure the existence of an attractive invariant measure and ergodicity results follow, [2, 16, 42]. However, the real convergence result of interest can be proved for the second algorithm which only gives rise to a nonhomogeneous Markov chain and for which the powerful methods that exist for the first case are not available. The method of proof relies here on a detailed analysis of the system dynamics using appropriate Lyapunov functions.

We consider a set of AIMD matrices for fixed additive increase parameter $\alpha > 0$ and multiplicative decrease parameter $\beta \in (0,1)$. We will assume there are probability functions $\lambda_i : [0,1] \to [0,1]$, $i = 1,\dots,n$ that are used by each agent to determine the probability of responding to the intermittent feedback signal, based on an average of past values of x. We assume that these functions

$$\lambda_i : [0,1] \to [0,1], \quad i = 1,\dots,n, \tag{23}$$

satisfy the following assumptions

(A1) λ_i is continuous, $i = 1, \dots, n$;

(A2) $r \mapsto r\lambda_i(r)$ is strictly increasing on $[0, 1]$, $i = 1, \dots, n$;

(A3) There exists a constant $\lambda_{\min}$ such that $\lambda_i(r) \geq \lambda_{\min} > 0$ for all $r \in [0, 1], i = 1, \dots, n$;

Note that these assumptions are satisfied for the choice of probability functions described in Section 3. In particular, (A2) is a consequence of convexity. [3] We will show in Lemma B.1 that under the above conditions there is a unique KKT point $x^* \in \Sigma$.

It is discussed in [43, 48] that the dynamics of an algorithm of the type of Algorithm 1 can be well approximated by a Markov chain of AIMD matrices. In fact, if we let $k = 0, 1, 2, \dots$ be the consecutive labels of the time instances at which the constraint is met, then the evolution form one constraint event to the next is given by (24) below, where $A(k)$ is one of the AIMD matrices describing the problem. Note that the probabilities $p_j(\cdot)$ for the matrices $A_j \in \mathcal{A}$ are now determined by the assumption that the agents act in a stochastically independent manner, so that the probability of a particular drop pattern encoded in A_j is given by the product of the probabilities of the individual agents responding or not.

The system of interest is given by the iteration of AIMD matrices in the form

$$x(k+1) = A(k)x(k)\,, \tag{24}$$

where the matrices $A(k)$ are chosen from the set of AIMD matrices $\mathcal{A}$ using a probability distribution that depends on the history of the sample path. Specifically, we consider the following two cases:

(i) **finite averaging:** We consider a fixed time window of length T. For $k \geq T$ consider the average

$$\bar{x}(k;T) := \frac{1}{T}\sum_{j=0}^{T-1} x(k-j)\,, \quad k = T, T+1, \dots\,, \tag{25}$$

and assume that

$$\mathbb{P}(A(k) = A_j) = p_j(\bar{x}(k;T))\,. \tag{26}$$

(i) **long-term averaging:** In this situation, we consider the average

$$\bar{x}(k) := \frac{1}{k+1}\sum_{j=0}^{k} x(j)\,, \quad k = 0, 1, \dots\,, \tag{27}$$

and assume that there are probability functions $p_j : \Sigma \to [0, 1]$, $j = 1, \dots, 2^n$ such that

$$\mathbb{P}(A(k) = A_j) = p_j(\bar{x}(k))\,. \tag{28}$$

4.1 Finite Averaging

The condition (26) needs to be interpreted along sample paths: for each specific realization of the Markov chain, the probabilities at time k are a function of the average along over the time interval $[k - T + 1, \dots, k]$ for the given realization. We will model this Markov chain as a Markov chain with state-dependent probabilities on the space Σ^T. In view of the evolution (24) with (25), (26), define the new variable

$$z(k) := \begin{bmatrix} x(k) & \frac{1}{2}\big(x(k) + x(k-1)\big) & \dots & \frac{1}{T}\big(x(k) + \dots + x(k-T+1)\big) \end{bmatrix}^\top . \tag{29}$$

[3]The formulation can be extended using the same arguments to the assumption that the AIMD parameters of the agents are chosen such that $\alpha \in \operatorname{ri}\Sigma, \beta \in (0, 1)^n$ satisfy the added assumption that the quotient $\alpha_i/(1 - \beta_i)$ is a constant independent of i.

It is then easy to see that the evolution of $z(k)$ is described by the Markov chain

$$z(k+1) = \begin{bmatrix} A(k) & 0 & & \dots & & 0 \\ \frac{1}{2}\left(A(k)+I\right) & 0 & & & & \\ \frac{1}{3}A(k) & \frac{2}{3}I & 0 & & & \vdots \\ \vdots & 0 & \ddots & \ddots & & \\ \vdots & & & \ddots & 0 & 0 \\ \frac{1}{T}A(k) & 0 & \dots & 0 & \frac{T-1}{T}I & 0 \end{bmatrix} z(k) =: A_T(k)z(k)\,. \tag{30}$$

Note that each matrix $A_j \in \mathcal{A}, j = 1, \dots, 2^n$ defines a matrix $A_{T,j}$ and the set of possible matrices $\mathcal{A}_T$ occurring in the Markov chain (30) is defined in this way. The Markov chain is thus defined with the place-dependent probabilities

$$\mathbb{P}\left(A_T(k) = A_{T,j}\right) = p_j\left(z_T(k)\right)\,, \tag{31}$$

where $z_T(k) \in \Sigma$ denotes the Tth component vector of $z(k)$. The following norm on $\mathbb{R}^{Tn}$ simplifies the analysis of the Markov chain considerably, as it reveals the contractive properties of the Markov chain. We define

$$\|z\|_T := \max_{i=1,\dots,T} \|z_i\|_1\,, \quad \text{where} \quad z = \begin{bmatrix} z_1^\top & \dots & z_T^\top \end{bmatrix}^\top, z_i \in \mathbb{R}^n, i = 1, \dots, T\,.$$

Lemma 4.1 *(i) For all $A_T \in \mathcal{A}_T$ the matrix norm induced by $\|\cdot\|_T$ satisfies*

$$\|A_T\|_T \leq 1\,. \tag{32}$$

(ii) The subspace

$$W := \{z \in \mathbb{R}^{Tn} \mid e^\top z_i = 0\,, \forall\, i = 1, \dots, T\} \tag{33}$$

is invariant under all $A_T \in \mathcal{A}_T$.

(iii) For all $z \in W$, $A_{T,j} \in \mathcal{A}_T$ it holds that

$$\|A_{T,j}z\|_T = \|z\|_T \quad \Rightarrow \quad A_j z_1 = z_1\,.$$

In particular, we have

$$\left\|A_{T,1|W}\right\| \leq \frac{c+T-1}{T} < 1\,, \tag{34}$$

where $c < 1$ is the constant given by Lemma 2.1.

Proof (i) This is a straightforward calculation.
(ii) This is an easy consequence of $e^\top A_j = e^\top, j = 1, \dots, 2^n$.
(iii) If $\|A_{T,j}z\|_T = \|z\|_T$, then necessarily $\|(1/i)A_j z_1 + ((i-1)/i)z_i\|_1 = \|z_i\|_1$ for an index i such that $\|z_i\|_1$ is maximal. As $\|A_j\|_1 \leq 1$ it follows that $\|A_j z_1\|_1 = \|z_i\|_1$ and so $\|z_1\|_1 = \|Az_1\|_1 = \|z_i\|_1$ as $\|z_i\|_1$ was maximal. Now it is known for the matrices $A \in \mathcal{A}$, that $\|z_1\|_1 = \|Az_1\|_1, e^\top z_1 = 0$ implies that $Az_1 = z_1$, [48, Lemma 3.8]. Finally,

$$\|A_{T,1}z\|_T = \max_{i=1,\dots,T} \left\| \frac{1}{i}Az_1 + \frac{i-1}{i}z_i \right\|_1 \leq \max_{i=1,\dots,T} \left\{ \frac{c}{i}\|z_1\|_1 + \frac{i-1}{i}\|z_i\|_1 \right\} \leq \max_{i=1,\dots,T} \left\{ \frac{c}{i} + \frac{i-1}{i} \right\} \|z_T\|_T\,.$$

This show the assertion. □

The previous result shows that the iteration of random choices of the $A_{T,i}$ is contractive when studied with respect to a suitable norm. This lies the foundation for proving the existence of a unique invariant and attractive measure for the Markov chain. Before proving this we need an assumption on the probability functions λ_i that guarantees strong contractivity on average.

Theorem 4.2 (Invariant Measure) *Assume that the probability functions λ_i satisfy (A1)-(A3) and are Lipschitz continuous. Then for all $T \geq 1$, there exists a unique invariant and attractive measure π^T on Σ^T. Furthermore, for all $z_0 \in \Sigma^T$, we have that almost surely*

$$\lim_{k\to\infty} \frac{1}{k} \sum_{\ell=0}^{k} z(l; z_0) = \int_{\Sigma^T} z \ \mathrm{d}\pi^T(z) = \mathbf{E}(\pi^T) . \tag{35}$$

Remark 4.3 *Stronger ergodicity assumptions holds as detailed in [2, 16]. We skip these for the sake of brevity.*

Proof It is easy to show that the sufficient conditions provided in [2] are satisfied. In particular, these conditions can be met by requiring that

$$\sup_{z,w\in\Sigma} \sum_{j=1}^{2^n} p_j(z) \frac{\|A_{T,j}(z-w)\|_T}{\|z-w\|_T} < 1 . \tag{36}$$

Note that $z-w \in W$, so Lemma 4.1 (i) immediately implies that the sum does not exceed 1. Assumption (A3) now ensures that for each $z \in \Sigma$, the probability $p_1(z) \geq \lambda_{\min}^n > 0$. Thus the probability of the matrix A_1 is bounded away from zero. Now Lemma 4.1 (iii) states that $\|A_1(z-w)\| \leq (c+T-1)/T\|z-w\| < \|z-w\|$. As $p_1(z) > 0$ for all z, we see that the supremum in (36) is bounded away from 0.

By Theorem 2.1 in [2] the existence of an attractive invariant measure follows. Uniqueness is then a consequence of attractivity. The ergodic property (35) now follows from [16]. □

Remark 4.4 *The previous results shows that the AIMD system is indeed converging in a strong sense; in particular, long term averages converge almost surely. Simulations suggest, that this limit gets closer to the KKT point as T increases. Also for large T, with high probability along a sample path, the average of the windows of size T is close to the KKT point.*

4.2 Long-Term Averaging

We now turn to the situation in which the probabilities for choosing the matrices depend on the long-term average of the realization. Note that (24) together with (28) do not define a Markov chain on Σ, as the probabilities do not depend on the current state $x(k)$ but rather on the complete history of a sample path. In order to obtain a formulation as a Markov chain we include the average in the state space. To this end we introduce the variable $z(k) := \begin{bmatrix} x(k)^\top & \bar{x}(x)^\top \end{bmatrix}^\top$ and the matrices

$$\bar{A}_j(k) := \begin{bmatrix} A_j & 0 \\ \frac{1}{k+2} A_j & \frac{k+1}{k+2} I \end{bmatrix}, \quad j = 1, \ldots, m, k \in \mathbb{N} . \tag{37}$$

We then arrive at the Markov chain on $\Sigma \times \Sigma$ defined by

$$z(k+1) =: \begin{bmatrix} z_1(k+1) \\ \bar{z}(k+1) \end{bmatrix} = \bar{A}(k) \begin{bmatrix} z_1(k) \\ \bar{z}(k) \end{bmatrix} = \bar{A}(k) z(k) , \tag{38}$$

with the probabilities

$$\mathbb{P}\big(\bar{A}(k) = \bar{A}_j(k)\big) = p_j(\bar{z}(k)) . \tag{39}$$

This defines a nonhomogeneous Markov chain with place dependent probabilities. Note that the non-homogeneity comes from the time-varying nature of the matrices $\bar{A}_j(k)$, whereas the functions $p_j(\cdot)$ describing the place-dependent probabilities do not depend on time.

To obtain contractive properties of the Markov chain (38) it will be of interest to study the matrices $A_j(k)$ using a particular norm. We define a norm on $\mathbb{R}^{2n}$ by setting for $x, y \in \mathbb{R}^n$

$$\left\| \begin{bmatrix} x \\ y \end{bmatrix} \right\| := \max\{\|x\|_1, \|y\|_1\} . \tag{40}$$

The matrix norm induced by this vector on $\mathbb{R}^{2n\times 2n}$ is also denoted by $\|\cdot\|$.
In the following we use the notation $\bar{\mathcal{A}} := \{\bar{A}_j(k) \mid j = 1, \ldots, 2^n, k \in \mathbb{N}\}$, which represents the set of all possible matrices appearing in (38).

Lemma 4.5 *(i) For all $\bar{A} \in \bar{\mathcal{A}}$*

$$\|\bar{A}\| \leq 1\,. \tag{41}$$

(ii) The subspace

$$W := \{(x, y) \in \mathbb{R}^{2n} \mid e^\top x = e^\top y = 0\} \tag{42}$$

is invariant under all $\bar{A} \in \bar{\mathcal{A}}$.

Proof The proof follows the lines of the proof of Lemma 4.1 and is omitted.

We stress that the key point of Lemma 4.1 was item (iii), which we used to obtain a uniform contractivity on the state space Σ^T of the Markov chain. A similar result can be obtained in the present situation, but uniformity is lost due to the time-dependent nature of the Markov chain. Unfortunately, the constant of contraction converges to 1. Considerable effort has been expensed on trying to transfer the proofs of [2, 16] to the present situation, but to no avail. We thus pursue an entirely different angle of attack in the proof of our main result.

Theorem 4.6 (Convergence) *Let the functions λ_i defined in* (23) *satisfy (A1)–(A3) and let x^* denote the KKT point guaranteed by Lemma B.1. Consider the nonhomogeneous Markov chain* (38)*. For any initial condition $z_0 = (z_1(0), \bar{z}(0)) \in \Sigma^2$ we have that the second component of $z(k; z_0)$ satisfies*

$$\lim_{k\to\infty} \bar{z}(k; z_0) = x^*\,, \quad \textit{almost surely.}$$

Remark 4.7 *Our result says that by local modification of the individual probabilities the agents can ensure almost sure convergence to the optimum. Inter-agent communication is not necessary; rather the only information needed is* 1*-bit intermittent message to all agents that congestion has occurred suffices for convergence. The main results of this section are the following:*

1. *an ergodicity result for the algorithm with finite-averaging;*
2. *a result guaranteeing almost sure convergence to the network optimum for the long-term averaging case;*
3. *the description of an easily implementable algorithms that ensures the convergence to the algorithm to the optimal point, using limited uniform communication to the agents. The algorithm is particularly suited to dumb devices that do not have extensive computational capabilities.*

Mathematically speaking, it is interesting to see an almost sure convergence result, that does not make use of the existence of an invariant measure of the stochastic process. We do expect however, that when considering the invariant measures π^T that are obtained for the case of finite time-windows, then as $T \to \infty$ the measures π^T converge to the Dirac measure in x^*.

5 Related Works

Our work lies in the intersection of two subjects: resource allocation and limiting characteristics of the stochastic version of the AIMD algorithm [23].

The AIMD literature is huge and it is not straightforward to discuss the available results in any sort of compact manner here. We refer interested readers to some recent works on this topic in the context of TCP and internet congestion control [30, 33, 28, 47, 35, 47, 49, 27, 19]. Much of this work is based on fluid approximations of AIMD dynamics; the notable exceptions are [43, 39, 42]. The latter of these papers make use of tools from iterated function systems to deduce the existence of such a unique probability distribution for standard linear AIMD networks (of which TCP is an example) under an assumption on the underlying probability model (albeit under very restrictive assumptions). To the best of our knowledge, this paper, along with the companion paper [43], established for the first time, the stochastic convergence of AIMD networks. However, the window of infinite length considered in this paper goes well beyond the set-up in these papers. In particular, the result presented in this paper may be considered as the limiting case of the results presented in these papers.

The literature on resource allocation is also immense and a full review is impossible here. Here, we briefly note that the subject of resource allocation or social welfare optimization has been studied in three prominent settings: centralized, distributed concerted, and distributed competitive.

In the first setting, there is a single decision-maker, who knows the utility function f_i of every agent, solves the network optimization problem, and assign the optimal allocation to each agent.

In the second setting, every resource user is a decision-maker that determines its own allocation according to a *fixed* policy (e.g., AIMD algorithm in the case of TCP), that is prescribed by a system operator and remains unchanged over time. When all agents follow fixed policies prescribed by the system operator, a number of distributed optimization algorithms have been proposed to iteratively converge to an optimal allocation of resources for numerous settings [15, 36, 24, 38]. Some of these algorithms are based on achieving consensus [36], others are based on distributed averaging [51, 32], and on stochastic approximation [3]. All these algorithms rely on communication between agents to achieve optimality. For example, when agents are assigned to nodes in a graph and restricted to communicate only with neighbours in that graph, the distributed dual averaging algorithm has be shown to guarantee the convergence of each agent's allocation to the optimal allocation over iterations of the algorithm [15]. Our work also considers a distributed setting with fixed-policy agents, but contains an important difference: the agents do not communicate among themselves, but are limited to an intermittent feedback signal from the network. Specifically, they only observe at each iteration whether the allocation is feasible (i.e., capacity constraint is satisfied). Another difference of our work is that our results do not depend on the existing convergence results from stochastic approximation, and hence hold under different conditions. In particular, we cannot apply the standard convergence argument for stochastic approximation because there are two time-scales (cf. [6, Chapters 6.2 and 10.4]).

In a third setting, every resource user is a decision-maker that acts strategically so as to optimize its utility function with regards to the actions of all agents. When all agents act strategically, solution concepts such as Nash equilibria are more meaningful than optimality concepts. In such a setting, a market mechanism based on bids has been proposed [25] and shown to be efficient in equilibrium under some assumptions [50]. Moreover, [20] shows that the equilibrium allocation is also the solution to an optimization problem. These works do not however provide a method for the agents to arrive at an equilibrium allocation. In contrast, our work may not consider strategic agents, but does present a set of policies that guarantee the convergence to an optimal allocation.

The link between congestion control (which encompasses the AIMD algorithm) and optimization has been noted by several authors [31, 44, 9, 46, 26, 13]. That various embodiments of TCP solve a network utility maximization problem is a cornerstone of much of the TCP literature [44]. However, we are dealing with the converse problem: given an NUM problem, is there an AIMD algorithm that solves it? Perhaps, the most closely related works in this direction are given in the following references: [9, 46, 26, 13]. Roughly speaking, these references follow two lines of direction. In the first direction, fluid-like approximations of congestion control are modified to address the NUM problem. This yields a sub-gradient like algorithm for solving the NUM problem. In the second direction, synchronized AIMD like algorithms are proposed to solve certain NUM problems using nonlinear backoff rules and nonlinear increase rules [12] [46, 26, 13]. The work presented here goes far beyond these works. First, we consider the matrix model of TCP proposed in [43] as opposed to a fluid model. Fluid approximations are valid only for very large numbers of agents, and the dynamic interaction between agents is often overlooked. The matrix model is an exact representation of AIMD dynamics under certain assumptions and can readily be implemented in existing software stacks. Furthermore, these models are often analysed using linearized approximations in contrast to our approach in which global stability (ergodicity) is proved. A further difference is that each agent responds to congestion according to its own probability function, known only to that agent. In this sense our proposed algorithms go far beyond traditional AIMD and emulate RED-like congestion control [44]. Second, we assume very limited actuation; an agent only decides to respond to congestion or not in an asynchronous manner. There is no need for a common clock and the setting is completely stochastic. In this context our results prove convergence and stability of the stochastic AIMD system and establish its suitability for solving large scale NUM problems.

6 Example

We now illustrate the application of our results. To this end consider a total of $n = 150$ agents are participating in the optimisation. Each agent has a cost function $f_i(\cdot)$ assigned which maps the share of the resource C to an associated cost. The cost-functions are chosen from the set of polynomials taking the following forms

$$\begin{aligned} f_{i1}(\bar{x}) &= a_{i1}\bar{x}^2 \\ f_{i2}(\bar{x}) &= a_{i2}\bar{x}^3 + b_{i2}\bar{x}^2 \\ f_{i3}(\bar{x}) &= a_{i3}\bar{x}^4 + b_{i3}\bar{x}^3 + c_{i3}\bar{x}^2 \\ f_{i4}(\bar{x}) &= a_{i4}\bar{x}^6 + b_{i4}\bar{x}^5 + c_{i4}\bar{x}^2. \end{aligned}$$

The parameters a_{ij},b_{ij}, and c_{ij} are the cost-factors of each function and are positive. Note that each cost is convex and strictly increasing on the interval $[0, 1]$. The objective is then

$$\begin{aligned} &\min_{x_1,\dots,x_N} \sum_{i=1}^{n} f_i(\bar{x}_i) \\ &\text{s.t.} \sum_{i=1}^{n} \bar{x}_i = C. \end{aligned} \tag{43}$$

For the simulations we choose the available resource to be equal to one, i.e. $C = 1$. Further, the cost-function type for each agent is selected randomly according to a uniform distribution. The cost-factors parameters for each agent are also selected randomly using a uniform distribution between 0 and 100. Each agent responds to congestion with probability

$$\lambda_i(\bar{x}_i(k)) = \Gamma \frac{f_i'(\bar{x}_i(k))}{\bar{x}_i(k)} \tag{44}$$

with $\bar{x}_i(k)$ as in Equation 27 in case of long-term averaging and

$$\lambda_i(\bar{x}_i(k,T)) = \Gamma \frac{f_i'(\bar{x}_i(k,T))}{\bar{x}_i(k,T)} \tag{45}$$

with $\bar{x}_i(k, T)$ as in Equation 25 in case of finite averaging. Here Γ is a network wide constant chosen to ensure that $0 < \lambda_i(k) < 1$. In our simulations we set $\Gamma = \frac{1}{1300}$. The remaining AIMD parameters are identical for all agents with $\alpha = 0.01$ and $\beta = 0.85$. From our main result we know that for large k we have that $\bar{x}_i(k) \approx x_i^*$. Thus we can write $\lambda_i(\bar{x}_i(k)) \approx \lambda_i^*$. It follows that

$$\lambda_i(\bar{x}_i(k)) \approx \Gamma \frac{f_i'(\bar{x}_i(k))}{\Theta} \lambda_i(k) \tag{46}$$

and that $f_i'(\bar{x}_i(k)) \approx f_j'(\bar{x}_j(k))$ for all i, j and large k, and where Θ is a network constant. These are precisely the KKT conditions. [4] We first simulate the long-term averaging case, where the average at time instant k is taken over all previous time-steps. Figure 1 shows the typical evolution of the derivative of the cost of seven randomly selected agents. It illustrates that the derivatives approach consensus as k increases, which minimises asymptotically the above stated optimisation problem.

Figure 6 shows the long-term average $\bar{x}_i(k)$ of seven arbitrary agents in comparison to their respective optimal state depicted by a dashed line. Figure 6 shows the absolute error between the long-term average and the optimal point for the same seven agents. With increasing time the long-term average approaches the optimal point for those seven randomly selected agents. In Figure 3 the maximal error between the long-term average and the optimal point is plotted, which approaches 0 with increasing time steps.

[4] Note that we have made several assumptions. First, we have assumed that eventually $\bar{x}_i(k) \approx x_i^*$. This follows from our main result. We have also assumed that $p_i(k) \approx p_i^*$. This follows from continuity of the $f_i(\cdot)$.

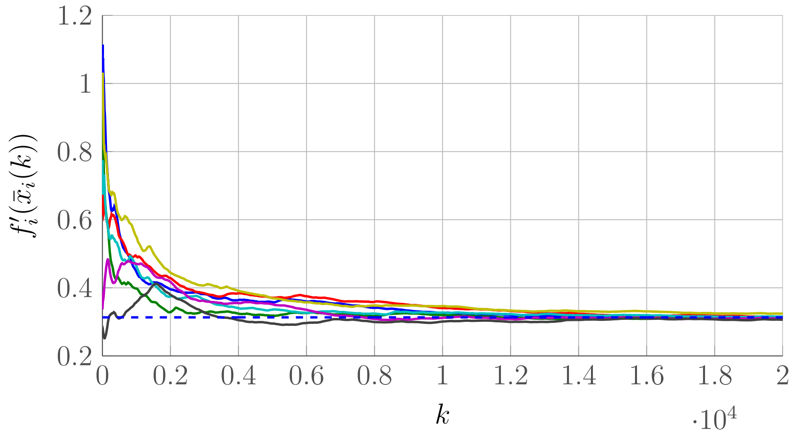

Figure 1: Evolution of the derivatives for seven randomly selected agents. The x-axis is a time index.

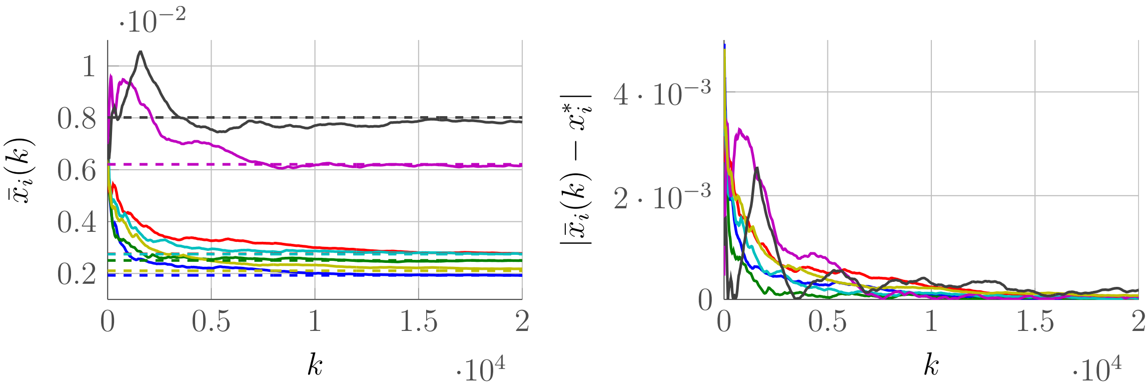

(a) $\bar{x}_i(k)$ (solid) and x_i^* (dashed)

(b) $|\bar{x}_i(k) - x_i^*|$

Figure 2: Evolution of the states in comparison to the optimal point for seven randomly selected agents. The x-axis is a time index.

We repeat our experiment for the finite averaging case, (30), with a fixed window size $T = 500$. Figure 4 shows the typical evolution of the derivative of the cost of seven randomly selected agents. It illustrates that the cost is oscillating around the optimal value.
Figure 5 shows the maximal error between the long-term average and the optimal point. Recall, that in these simulations the long term average is not used for determining the drop probabilities. Its limit exists almost surely and is given by the expectation of the underlying invariant measure, see Theorem 4.2. While the cost computed with the finite average is oscillating, the long term average of the state is still converging towards the optimal value.

7 Conclusions

In this paper we have derived a convergence result for the non-homogeneous Markov chain that arises in the study of networks employing the *additive-increase multiplicative decrease* (AIMD) algorithm. We then used this result to solve the network utility maximization problem in a very simple manner. Future work will consider the behaviour of finite window averaging systems.
We have also applied our approach to the distributed optimization of agents that have two possible states—ON and OFF—and finite memory. Initial, but extensive, simulation studies suggest that our results hold in this setting as well. Finally, it is worthwhile to note that it is possible to extend our results in the following directions: allowing updates to the increase and decrease parameters (α_i and β_i) over time to achieve faster convergence to optimality, allowing agents leaving and joining the allocation system over time, and allowing the joint allocation of multiple resources; e.g., bandwidth over distinct

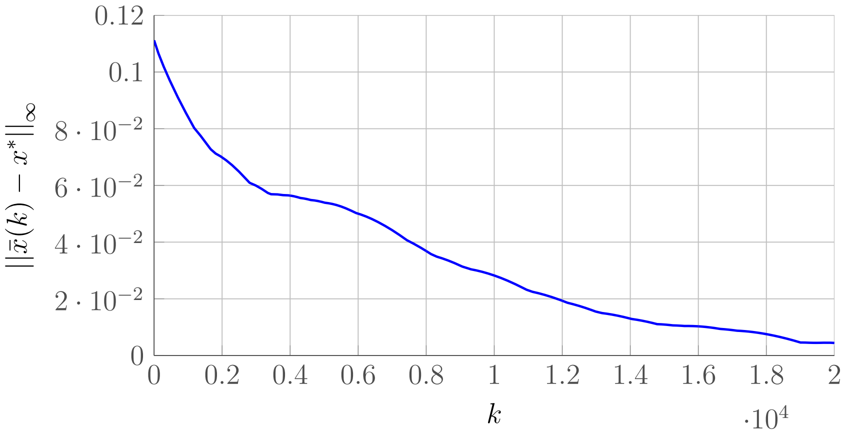

Figure 3: Evolution of the maximal absolute error between the states and the optimal, i.e. $||\bar{x}(k) - x^*||_\infty$. The x-axis is a time index.

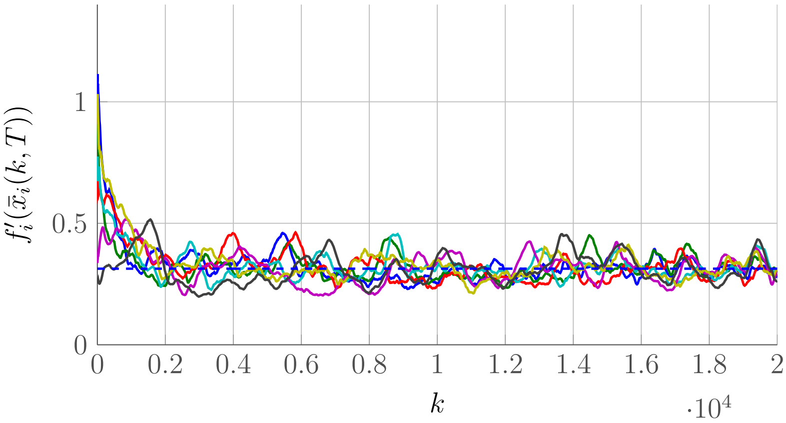

Figure 4: Evolution of the derivatives for seven randomly selected agents. The simulations are done for the finite averaging case with a fixed window size $T = 500$. The x-axis is a time index.

links of a communication network, or a combination of bandwidth and computation resources in a cloud server.

8 Acknowledgement

The authors thank Ronald Fagin for helpful comments on an earlier version of this paper.

A Preamble to the Proofs

We now briefly explain the structure of the proof of the main result, as otherwise the reader may sometimes wonder why we need certain intermediate results.

The key intuition is that in the long run the dynamics of the long term average $\bar{x}$ becomes slow. In other words, for large T and relatively short intervals of length m of the form $[T+1, T+m]$, $\bar{x}$ is almost a constant, where m is to be understood to be small when compared to T. The reason for this is the simple relation

$$\bar{x}(T+m) = \frac{T}{T+m}\bar{x}(T) + \frac{m}{T+m}\left(\frac{1}{m}\sum_{\ell=1}^{m} x(T+\ell)\right), \tag{47}$$

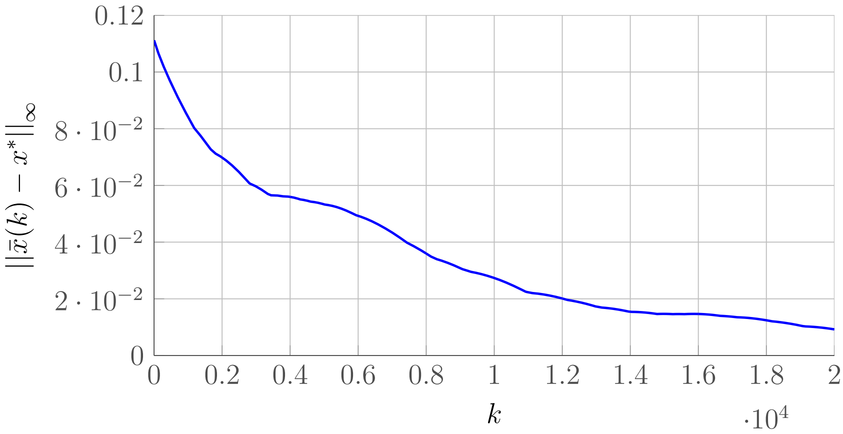

Figure 5: Evolution of the maximal absolute error between the states and the optimal, i.e. $||\bar{x}(k) - x^*||_\infty$. The simulations are done for the finite averaging case with a fixed window size $T = 500$. The x-axis is a time index.

which holds along any sample path.

If $\bar{x}$ is almost a constant on a certain interval, then the probabilities for choosing the matrices A_j are almost constant, and we can approximate the dynamics using the results on AIMD with constant probabilities; and consequently Lemma 2.2 becomes relevant. This result says that, provided that m is large enough, the average over the next m steps is close to the expectation of the AIMD Markov chain with constant probabilities. And this holds for all starting conditions $x(T)$ and with high probability.

While this basic intuition turns out to be true, we need to resolve the fact that the ergodic limit of the "fixed-probability system" depends on T. Specifically, m and T depend on each other and the precise resolution of our proof depends on understanding this relationship.

To resolve this, we use the following interpretation of (47). We denote by $P(\bar{x})$ the expectation of the invariant measure of the AIMD Markov process with the fixed probabilities $\lambda_i(\bar{x}_i), i = 1, \dots, n$. We then rewrite (47) as

$$\bar{x}(T+m) = \frac{T}{T+m}\bar{x}(T) + \frac{m}{T+m}\left(P(\bar{x}(T)) + \Delta(T)\right) , \tag{48}$$

where we interpret $\Delta(T)$ as a suitable perturbation term, that aggregates the effect that the probabilities are not precisely constant on $[T+1, T+m]$, and the further effect that we are not at the expectation but only close to it.

To understand the dynamics in (48) we study the system

$$\bar{x}(T+m) = \frac{T}{T+m}\bar{x}(T) + \frac{m}{T+m}\, P(\bar{x}(T)) , \tag{49}$$

and interpret system (48) as a perturbed version thereof. This is the sole purpose of Appendix B, in which we obtain (i) characterizations of the unique fixed point of (49), (ii) characterizations of attractivity properties of neighbourhoods of this fixed point in dependence of the size of $m/(T+m)$, and (iii) the necessary robustness results to extend these attractivity statements to the perturbed system (48). In Appendix C we then bring the stochastic nature of our nonhomogeneous Markov chain into play and use the results of Appendix B to prove almost sure convergence using what is essentially a Lyapunov type argument.

B Deterministic Iteration

In this section we present a collection of stability and robustness results for a deterministic system closely related to the AIMD Markov chain. These results will turn out to be instrumental in the proof of the main result Theorem 4.6.

The first results study a deterministic system defined by successive convex combinations of a point in Σ with the expectation of this point as defined through (13).

Recall, that we assume that $\alpha \in \mathrm{ri}\,\Sigma, \beta \in (0,1)^n$ satisfy the assumption that the quotient $\alpha_i/(1-\beta_i)$ is a constant independent of i. As a consequence the limiting value defined in (13), given the probabilities $\lambda_i(\cdot), i = 1,\dots,n$ at a given point in Σ, reduces to

$$\frac{1}{\sum_{\ell=1}^n \frac{\alpha_i}{\lambda_\ell(x_\ell)(1-\beta_i)}} \begin{bmatrix} \frac{\alpha_1}{\lambda_1(x_1)(1-\beta_1)} \\ \vdots \\ \frac{\alpha_n}{\lambda_n(x_n)(1-\beta_n)} \end{bmatrix} = \frac{1}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \begin{bmatrix} \lambda_1(x_1)^{-1} \\ \vdots \\ \lambda_n(x_n)^{-1} \end{bmatrix}. \tag{50}$$

We thus arrive at the map

$$P : \Sigma \to \Sigma\,, \quad x \mapsto \frac{1}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \begin{bmatrix} \lambda_1(x_1)^{-1} \\ \vdots \\ \lambda_n(x_n)^{-1} \end{bmatrix}. \tag{51}$$

Note that $P(\Sigma) \subset \mathrm{ri}\,\Sigma$ by (A3). We may therefore chose a constant $\delta^- > 0$ such that

$$P(\Sigma) + \overline{B}_1(0,\delta^-) \subset \mathrm{conv}\,P(\Sigma) + \overline{B}_1(0,2\delta^-) \subset \mathrm{ri}\,\Sigma\,. \tag{52}$$

Note that in this instance, and in the following, scalings and sums of sets are in the standard sense of Minkowski sums. Also the factor 2 is an arbitrarily chosen factor that will become useful in later robustness estimates. All that is required is that this factor exceeds 1. Furthermore, we require the constant

$$\delta^+ := \max\{\mathrm{dist}_1(y, P(\Sigma))\ ;\ y \in \Sigma\ \}\,. \tag{53}$$

We will be interested in systems that perform successive convex combinations of the state x and $P(x)$. For $\{\varepsilon_k\}_{k\in\mathbb{N}} \subset (0,1)$ consider the system[5]

$$x(k+1) = (1-\varepsilon_k)x(k) + \varepsilon_k P(x(k))\,. \tag{54}$$

We note the following simple properties of the iteration in (54).

Lemma B.1 *Assume that (A1)–(A3) hold. Let $n \geq 2$, $\alpha \in \mathrm{ri}\,\Sigma_n, \beta \in (0,1)^n$, such that $\alpha_i/(1-\beta_i)$ is independent of i. Then*

(i) There exists a unique point $x^* \in \Sigma$, with $P(x^*) = x^*$.

(ii) The fixed point x^* is characterized by the property

$$x_i^*\lambda_i(x_i^*) = x_j^*\lambda_j(x_j^*) =: \gamma_F\,, \quad \forall i,j \in \{1,\dots,n\}\,. \tag{55}$$

(iii) For every $\varepsilon \in (0,1)$ the point x^* defined in (i) is the unique fixed point of

$$x \mapsto (1-\varepsilon)x + \varepsilon P(x)\,.$$

(iv) For every $x_0 \in \Sigma$ and every sequence $\{\varepsilon_k\}_{k\in\mathbb{N}} \subset (0,1)$ the solution of (54) satisfies $x(k) \gg 0$ for all $k \geq 1$.

[5]This system is an instance of stochastic approximation, and convergence results of [6, Chapter 2] hold under appropriate assumptions. However, we shall require and derive stronger results.

Proof (i) As $P : \Sigma \to \Sigma$ is continuous the existence of a fixed point of P follows from Brouwer's fixed point theorem. If there are two fixed points $x^* \neq y^*$ of P, $x^*, y^* \in \Sigma$, then without loss of generality we may assume $x_1^* > y_1^*$ and $x_2^* < y_2^*$. Using the fixed point property and the definition of P we obtain

$$x_1^*\lambda_1(x_1^*) = \frac{1}{\sum_{\nu=1}^n \lambda_\nu(x_\nu^*)^{-1}} = x_2^*\lambda_2(x_2^*) \quad \text{and} \quad y_1^*\lambda_1(y_1^*) = \frac{1}{\sum_{\nu=1}^n \lambda_\nu(y_\nu^*)^{-1}} = y_2^*\lambda_2(y_2^*) \,.$$

But from Assumption (A2) we have $x_1^*\lambda_1(x_1^*) > y_1^*\lambda_1(y_1^*) = y_2^*\lambda_2(y_2^*) > x_2^*\lambda_2(x_2^*)$. This contradiction completes the proof.
(ii) This follows from a straightforward argument.
(iii) This is an immediate consequence of (i).
(iv) This follows as $P(x) \gg 0$ for all $x \in \Sigma$ by definition and using Assumption (A3). $\square$

In order to simplify notation, we introduce for $\varepsilon \in [0,1]$ the map $R_\varepsilon : \Sigma \to \Sigma$ by

$$R_\varepsilon(x) := (1-\varepsilon)x + \varepsilon P(x) \,. \tag{56}$$

It is clear, that for all $\varepsilon \in (0,1]$ the fixed point x^* of P is also the unique fixed point of R_ε.
In our analysis of the dynamics we require two types of contractive properties of the map R_ε in combination with robustness results. We will also consider set-valued maps of the form

$$\Psi_\varepsilon^\delta(x) := R_\varepsilon(x) + \varepsilon \overline{B}_1(0,\delta) = (1-\varepsilon)x + \varepsilon(P(x) + \overline{B}_1(0,\delta)) \,, \tag{57}$$

where we assume $0 < \delta < \delta^-$. Note that by definition of δ^- this ensures that $\Psi_\varepsilon^\delta(x) \subset \operatorname{ri}\Sigma$. In the following lemma, we analyse properties of the map Ψ_ε^δ by studying individual elements in its image.
The next result describes two important features of the iteration

$$x(k+1) \in \Psi_\varepsilon^\delta(x(k)) \,. \tag{58}$$

On one hand by (i) the iteration converges with rate $(1-\varepsilon)$ to the convex set $\operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$. On the other hand using (ii) if the iteration is perturbed so that all we know that there is a convex combination with some $y \in \Sigma$ then we may bound the increase of the distance to the convex set. Finally, by (iii) the error induced by the perturbation y can be linearly bounded in ε, provided that we are sufficiently far way from $\operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$. For the following statement recall the definition of δ^- in (52).

Lemma B.2 *Let $x \in \Sigma$. Then for all $0 < \varepsilon \leq 1$:*

(i) For all $0 < \delta < \delta^-$ and $\Delta \in \mathbb{R}^n, e^\top \Delta = 0, \|\Delta\| \leq \delta$ we have

$$\operatorname{dist}_1(R_\varepsilon(x) + \varepsilon\Delta, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \leq (1-\varepsilon)\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \,. \tag{59}$$

(ii) In view of (53), for all $0 < \delta < \delta^-$ and all $y \in \Sigma$, we have

$$\operatorname{dist}_1((1-\varepsilon)x + \varepsilon y, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \leq (1-\varepsilon)\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) + \varepsilon\delta^+ \,. \tag{60}$$

(iii) For every $0 < \bar{\delta} < \delta^-$ there exists a $C_{\bar{\delta}} > 0$ such that for all $0 < \varepsilon < 1$ and all $0 < \delta < \delta^-$ we have the following implication: If $x \in \Sigma$ satisfies $\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) > \bar{\delta}$ and $y \in \Sigma$, then

$$\operatorname{dist}_1((1-\varepsilon)x + \varepsilon y, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \leq (1 + C_{\bar{\delta}}\varepsilon)\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \,. \tag{61}$$

Proof (i) Let $z \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$ be such that

$$\|x - z\| = \operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \,.$$

Then by convexity $(1-\varepsilon)z + \varepsilon(P(x) + \Delta) \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$ and so

$$\operatorname{dist}_1(R_\varepsilon(x) + \varepsilon\Delta, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) \leq \|(R_\varepsilon(x) + \varepsilon\Delta) - ((1-\varepsilon)z + \varepsilon(P(x) + \Delta))\| = (1-\varepsilon)\|x - z\| \,.$$

(ii) To prove (60) note that for any convex set C, we have $C = (1-\varepsilon)C + \varepsilon C$. Hence,

$$\operatorname{dist}_1((1-\varepsilon)x + \varepsilon y, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) = \\ (1-\varepsilon)\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) + \varepsilon \operatorname{dist}_1(y, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)),$$

which shows the claim by definition of δ^+.

(iii) To prove (61) note that with the assumption $\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)) > \bar{\delta}$ we arrive at

$$\delta^+ \leq \delta^+ \frac{\operatorname{dist}_1(x, \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta))}{\bar{\delta}},$$

and so (61) follows from (60) with an appropriate choice of $C_{\bar{\delta}} > 0$. This completes the proof. □

It is the aim of the following sequence of results to establish similar properties close to the fixed point x^*. To this end we have found it necessary to work with a different metric. The following result is a cornerstone in our proof of the main result.

Theorem B.3 *Let $x^* \in \Sigma$ be the unique fixed point of P, as described in Lemma B.1. For every $\eta > 0$, there is $1 > \varepsilon_0 > 0$ such that for all $0 < \varepsilon < \varepsilon_0$ we have*

$$d_H(x, x^*) \geq \eta \quad \Rightarrow \quad d_H(R_\varepsilon(x), x^*) < d_H(x, x^*). \tag{62}$$

Proof If $x \in \Sigma$ and some entries of x are zero, then $d_H(x, x^*) = \infty$, and also $R_\varepsilon(x) \gg 0$ by construction. Thus the claim follows trivially. In the remainder of the proof we will thus assume that $x \gg 0$. We will make use of the following elementary observations. First we have

$$x \in \operatorname{ri} \Sigma \text{ and } x \neq x^* \quad \Rightarrow \quad \min_j \{x_j / x_j^*\} < 1 < \max_i \{x_i / x_i^*\}. \tag{63}$$

Using this relation it is straightforward to see that for any sequence $\{x_k\} \subset \operatorname{ri} \Sigma$ we have the equivalence

$$e^{d_H}(x_k, x^*) \to 1 \quad \Leftrightarrow \quad \min_j \{x_{kj} / x_j^*\} \to 1 \quad \Leftrightarrow \quad \max_i \{x_{ki} / x_i^*\} \to 1. \tag{64}$$

Recall that the definition yields

$$e^{d_H}(R_\varepsilon(x), x^*) = \frac{\max_i \{R_\varepsilon(x)_i / x_i^*\}}{\min_j \{R_\varepsilon(x)_j / x_j^*\}}. \tag{65}$$

In case $x_i > x_i^*$, we obtain using the constant γ_F defined in (55) and the fact that $x_i \lambda_i(x_i) > \gamma_F$ by assumption (A2) that

$$\begin{aligned} R_\varepsilon(x)_i &= (1-\varepsilon)x_i + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{\lambda_i(x_i)} \\ &= \left((1-\varepsilon) + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{x_i \lambda_i(x_i)} \right) x_i < \left((1-\varepsilon) + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{\gamma_F} \right) x_i . \end{aligned} \tag{66}$$

By a similar argument, if $x_j < x_j^*$ we obtain

$$R_\varepsilon(x)_j = (1-\varepsilon)x_j + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{\lambda_j(x_j)} > \left((1-\varepsilon) + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{\gamma_F} \right) x_j . \tag{67}$$

Combining (66) and (67), we obtain that if $x_i > x_i^*$ and $x_j < x_j^*$, then

$$\frac{R_\varepsilon(x)_i / x_i^*}{R_\varepsilon(x)_j / x_j^*} < \frac{x_i / x_i^*}{x_j / x_j^*} \leq e^{d_H}(x, x^*). \tag{68}$$

Now, for $e^{d_H}(R_\varepsilon(x), x^*)$, there are two possibilities: (i) the value in the defining equality (65) is attained in a pair (i,j) such that $x_i > x_i^*$ and $x_j < x_j^*$, in which case we obtain from (68) that $d_H(R_\varepsilon(x), x^*) <$

$d_H(x, x^*)$, as desired. Alternatively, (ii) the value in the defining equality (65) is attained in a pair (i, j) such that $x_i \le x_i^*$ or $x_j \ge x_j^*$. To complete the proof, it is thus sufficient to show that we may choose $\varepsilon_0 > 0$ small enough so that for $0 < \varepsilon < \varepsilon_0$ the second case can only occur for $x \in B_H(x^*, \eta)$. So assume $\eta > 0$ is fixed. We will make use of three basic inequalities which we derive now. By (64) there exists a constant $r \in (0, 1)$ such that for all $x \in \mathrm{ri}\,\Sigma$ with $d_H(x, x^*) \ge \eta$ we have

$$\min_{j=1,\dots,n} \frac{x_j}{x_j^*} \le r < 1\,. \tag{69}$$

Also, again by (64), there exists a constant $R > 1$ such that for all $x \in \mathrm{ri}\,\Sigma$ with $d_H(x, x^*) \ge \eta$ we have

$$\max_{i=1,\dots,n} \frac{x_i}{x_i^*} \ge R > 1\,. \tag{70}$$

Finally, using (A3) and $\lambda_i(x_i) \in [\lambda_{\min}, 1]$, we have for all $x \in \Sigma$ that

$$0 < c_1 := \frac{\lambda_{\min}}{n} \le P_i(x) = \frac{1}{\sum_{j=1}^n \lambda_j(x_j)^{-1}} \frac{1}{\lambda_i(x_i)} \le \min\left\{1, \frac{1}{n}\frac{1}{\lambda_{\min}}\right\} =: c_2\,. \tag{71}$$

Define $x^*_{\min} := \min_{i=1,\dots,n} x_i^* > 0, x^*_{\max} := \max_{i=1,\dots,n} x_i^* > 0$ and define $\varepsilon_1 > 0$ by

$$\varepsilon_1 := \frac{1 - r}{1 - r + \frac{c_2}{x^*_{\min}} - \frac{c_1}{x^*_{\max}}}\,. \tag{72}$$

We claim that the assumptions that $d_H(x, x^*) \ge \eta$, and $0 < \varepsilon < \varepsilon_1$ imply that

$$\frac{R_\varepsilon(x)_j}{x_j^*} = \min_{\ell=1,\dots,n} \frac{R_\varepsilon(x)_\ell}{x_\ell^*} \quad \Rightarrow \quad x_j < x_j^*\,. \tag{73}$$

So fix $x \in \mathrm{ri}\,\Sigma, d_H(x, x^*) \ge \eta$ and an index j so that $x_j/x_j^* \le r$, where we have used (69). Let $\ell \in \{1, \dots, n\}$ be any index for which $x_\ell \ge x_\ell^*$. Then we obtain using (69), (70), (71) as well as the definition of $x^*_{\min}, x^*_{\max}$ that

$$\frac{R_\varepsilon(x)_j}{x_j^*} = \frac{(1-\varepsilon)x_j + \varepsilon P_j(x)}{x_j^*} \le (1-\varepsilon)r + \varepsilon\frac{c_2}{x^*_{\min}} \overset{(72)}{<} (1-\varepsilon) + \varepsilon\frac{c_1}{x^*_{\max}} \le \frac{(1-\varepsilon)x_\ell + \varepsilon P_\ell(x)}{x_\ell^*}\,.$$

This shows (73). With a similar calculation we obtain that with

$$\varepsilon_2 := \frac{R - 1}{R - 1 + \frac{c_2}{x^*_{\min}} - \frac{c_1}{x^*_{\max}}}\,, \tag{74}$$

the conditions $d_H(x, x^*) \ge \eta$, and $0 < \varepsilon < \varepsilon_2$ imply that

$$\frac{R_\varepsilon(x)_i}{x_i^*} = \max_{\ell=1,\dots,n} \frac{R_\varepsilon(x)_\ell}{x_\ell^*} \quad \Rightarrow \quad x_i > x_i^*\,. \tag{75}$$

Summarizing, we obtain the following result from (73) and (75): If $0 < \varepsilon < \varepsilon_0 := \min\{\varepsilon_1, \varepsilon_2\}$, then $d_H(x, x^*) \ge \eta$ implies that the premise of (68) is satisfied for the pair (i, j) of indices such that $(R_\varepsilon(x_i)/x_i^*)/(R_\varepsilon(x_j)/x_j^*) = e^{d_H}(x, x^*)$. Then (68) shows the claim and the proof is complete. □

Corollary B.4 *Let $x^* \in \Sigma$ be the unique fixed point of P, as described in Lemma B.1. For every $\eta > 0$, there is $1 > \varepsilon_0 > 0$ and a constant $C_\eta > 0$ such that for all $0 < \varepsilon < \varepsilon_0$ we have*

$$d_H(x, x^*) \ge \eta \quad \Rightarrow \quad e^{d_H}(R_\varepsilon(x), x^*) < (1 - C_\eta\varepsilon)e^{d_H}(x, x^*)\,, \tag{76}$$

or equivalently,

$$d_H(x, x^*) \ge \eta \quad \Rightarrow \quad d_H(R_\varepsilon(x), x^*) < d_H(x, x^*) + \log(1 - C_\eta\varepsilon) < d_H(x, x^*) - C_\eta\varepsilon\,. \tag{77}$$

Proof Fix $\eta > 0$, let $\varepsilon > 0$ be the constant constructed in the proof of Theorem B.3. Also let r, R be the constants corresponding to η as given by (69), (70). Using Assumptions (A1), (A2) and Lemma B.1, there are constant $L_1 < \gamma_F < L_2$ such that

$$x_j \lambda_j(x_j) \le L_1 \quad \text{for all } j = 1, \dots, n, \text{ and all } x_j \in [0,1] \text{ such that } \frac{x_j}{x_j^*} \le r,$$

$$x_i \lambda_i(x_i) \ge L_2 \quad \text{for all } i = 1, \dots, n, \text{ and all } x_i \in [0,1] \text{ such that } \frac{x_i}{x_i^*} \ge R.$$

With this notation we can refine the inequalities (66) and (67). Namely, if $d_H(x, x^*) \ge \eta$ and if i, j are indices such that

$$e^{d_H}(x, x^*) = \frac{x_i / x_i^*}{x_j / x_j^*}, \tag{78}$$

then with (69), (70) we have $x_i/x_i^* \ge R$, $x_j/x_j^* \le r$. We obtain following the steps of (66) and (67)

$$\begin{aligned} e^{d_H}(R_\varepsilon(x), x^*) \le \frac{R_\varepsilon(x)_i / x_i^*}{R_\varepsilon(x)_j / x_j^*} &< \frac{(1-\varepsilon) + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{L_2}}{(1-\varepsilon) + \frac{\varepsilon}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}} \frac{1}{L_1}} \, e^{d_H}(x, x^*) \\ &= \left(1 - \varepsilon \, \frac{\frac{1}{L_1} - \frac{1}{L_2}}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}(1-\varepsilon) + \frac{\varepsilon}{L_1}} \right) \, e^{d_H}(x, x^*) \,. \end{aligned}$$

The term on the right hand side may be bounded by

$$< (1 - C_\eta \varepsilon) \, e^{d_H}(x, x^*) \,,$$

where

$$C_\eta = \min_{\varepsilon \in [0,1]} \left\{ \frac{\frac{1}{L_1} - \frac{1}{L_2}}{\sum_{\nu=1}^n \lambda_\nu(x_\nu)^{-1}(1-\varepsilon) + \frac{\varepsilon}{L_1}} \right\} > 0 \,.$$

Note that C_η depends on η as the choice of r, R is a function of η and these constants in turn determine possible values for L_1, L_2. The final claim follows from a simple application of the logarithm and by using a standard inequality. ☐

We also need the following two robustness results. The first concerns the perturbed averaged system (58), while the second yields a bound on the worst case behaviour of convex combination with arbitrary points in Σ.

Lemma B.5 *Let $x^* \in \Sigma$ be the unique fixed point of P, as described in Lemma B.1. Consider $\delta^- > 0$ as defined in (52). There exists a constant $K > 0$ such that for all $0 < \delta < \delta^-$, $\varepsilon \in (0,1)$, all $x \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$, and all $\Delta \in \mathbb{R}^n$ with $e^\top \Delta = 0$ and $\|\Delta\|_1 \le \delta$ we have*

$$e^{d_H}(R_\varepsilon(x) + \varepsilon\Delta, x^*) - e^{d_H}(R_\varepsilon(x), x^*) \le K \varepsilon \delta \,. \tag{79}$$

Proof The assumption on δ yields that $\operatorname{conv} P(\Sigma) + \overline{B}_1(0, 2\delta) \subset \operatorname{ri} \Sigma$, see (52). By definition we have

$$e^{d_H}(R_\varepsilon(x) + \varepsilon\Delta, x^*) - e^{d_H}(R_\varepsilon(x), x^*) = \frac{\max_i \{R_\varepsilon(x)_i + \varepsilon\Delta_i / x_i^*\}}{\min_j \{R_\varepsilon(x)_j + \varepsilon\Delta_j / x_j^*\}} - \frac{\max_i \{R_\varepsilon(x)_i / x_i^*\}}{\min_j \{R_\varepsilon(x)_j / x_j^*\}}$$

assuming the i, j are chosen so that the maximum, resp. minimum is attained for the perturbed term, we may continue

$$\le \frac{(R_\varepsilon(x)_i + \varepsilon\delta) R_\varepsilon(x)_j / x_i^* x_j^* - R_\varepsilon(x)_i (R_\varepsilon(x)_j - \varepsilon\delta) / x_i^* x_j^*}{(R_\varepsilon(x)_j - \varepsilon\delta) R_\varepsilon(x)_j / x_j^* x_j^*} = \varepsilon\delta \frac{(R_\varepsilon(x)_j + R_\varepsilon(x)_i) / x_i^*}{(R_\varepsilon(x)_j - \varepsilon\delta) R_\varepsilon(x)_j / x_j^*} \,.$$

To complete the proof, we need to show that the factor of $\varepsilon\delta$ in the expression on the right can be uniformly bounded for all $x \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$. By assumption, $\operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$ is a compact subset of $\operatorname{ri}\Sigma$, so that all entries of x and $R_\varepsilon(x)$ are bounded away from 0. Furthermore, the terms $R_\varepsilon(x)_j - \varepsilon\delta$ are bounded away from 0, because for arbitrary indeces $j' \neq j$ we have $R_\varepsilon(x) - \varepsilon\delta e_j + \varepsilon\delta e_{j'} \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0, 2\delta) \subset \operatorname{ri}\Sigma$. Thus the factor of $\varepsilon\delta$ in the final expression may be bounded by a constant, as the denominator is bounded away from 0. This constant only depends on δ^-. This proves the claim. ☐

Corollary B.6 *Let $x^* \in \Sigma$ be the unique fixed point of P, as described in Lemma B.1. For a given $\eta > 0$, let $1 > \varepsilon_0 > 0$ and a $C_\eta > 0$ be the constants of Corollary B.4 such that (76) and (77) hold. Let*

$$\delta^* := \min\left\{ \frac{C_\eta e^\eta}{2K}, \delta^- \right\}. \tag{80}$$

Then for every $0 < \delta < \delta^*$, all $0 < \varepsilon < \varepsilon_0$ and all $x \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$ and all $\Delta \in \mathbb{R}^n$, $e^\top \Delta = 0, \|\Delta\| \le \delta$ we have $R_\varepsilon(x) + \varepsilon\Delta \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$ and

$$d_H(x, x^*) \ge \eta \quad \Rightarrow \quad e^{d_H}(R_\varepsilon(x) + \varepsilon\Delta, x^*) < (1 - \frac{C_\eta}{2}\varepsilon) e^{d_H}(x, x^*), \tag{81}$$

or equivalently,

$$d_H(x, x^*) \ge \eta \quad \Rightarrow \quad d_H(R_\varepsilon(x) + \varepsilon\Delta, x^*) < d_H(x, x^*) + \log\left(1 - \frac{C_\eta}{2}\varepsilon\right) < d_H(x, x^*) - \frac{C_\eta}{2}\varepsilon. \tag{82}$$

Proof The first claim $R_\varepsilon(x) + \varepsilon\Delta \in \operatorname{conv} P(\Sigma) + \overline{B}_1(0,\delta)$ is obvious by convexity. Under the assumptions we may apply Corollary B.4 to obtain that $d_H(x, x^*) \ge \eta$ implies

$$e^{d_H}(R_\varepsilon(x), x^*) < (1 - C_\eta\varepsilon) e^{d_H}(x, x^*). \tag{83}$$

Thus with an application of Lemma B.5 we obtain

$$e^{d_H}(R_\varepsilon(x) + \varepsilon\Delta, x^*) \le e^{d_H}(R_\varepsilon(x), x^*) + K\varepsilon\delta \le (1 - \frac{C_\eta}{2}\varepsilon) e^{d_H}(x, x^*) + \varepsilon\left(-\frac{C_\eta}{2} e^\eta + K\delta\right).$$

By assumption the last term on the right hand side is negative and we obtain (81). The final claim (82) is then obvious. ☐

Lemma B.7 *Let $x^* \in \Sigma$ be the unique fixed point of P, as described in Lemma B.1. There exists a constant $C > 0$ such that for all $x, y \in \Sigma$ with $x \gg 0$ and for all $\varepsilon \in [0,1)$ we have*

$$d_H((1-\varepsilon)x + \varepsilon y, x^*) \le d_H(x, x^*) + \log\left(1 + C\frac{\varepsilon}{1-\varepsilon}\right). \tag{84}$$

In particular, for any $0 < \varepsilon_0 < 1$ there is a constant C_0 such that for all $x, y \in \Sigma$ with $x \gg 0$ and for all $\varepsilon \in [0, \varepsilon_0)$ we have

$$d_H((1-\varepsilon)x + \varepsilon y, x^*) \le d_H(x, x^*) + C_0\varepsilon. \tag{85}$$

Proof Let $x, y \in \Sigma$ be arbitrary with $x \gg 0$. Then we obtain

$$\begin{aligned} e^{d_H}((1-\varepsilon)x + \varepsilon y, x^*) &= \frac{\max_i\{((1-\varepsilon)x_i + \varepsilon y_i)/x_i^*\}}{\min_j\{((1-\varepsilon)x_j + \varepsilon y_j)/x_j^*\}} \le \frac{\max_i\{((1-\varepsilon)x_i + \varepsilon)/x_i^*\}}{\min_j\{(1-\varepsilon)x_j/x_j^*\}} \\ &\le e^{d_H}(x, x^*) + \frac{\varepsilon}{1-\varepsilon}\frac{\max_i\{1/x_i^*\}}{\min_j\{x_j/x_j^*\}} = e^{d_H}(x, x^*) + \frac{\varepsilon}{1-\varepsilon}\frac{\max_i\{1/x_i^*\}}{\min_j\{x_j/x_j^*\}}\frac{\max_i\{x_i/x_i^*\}}{\max_i\{x_i/x_i^*\}} \end{aligned}$$

and using that $\max_i\{x_i/x_i^*\} \ge 1$ and $x_i^* \ge x_{\min}^*$ we obtain

$$\le e^{d_H}(x, x^*)\left(1 + \frac{\varepsilon}{1-\varepsilon}\frac{1}{x_{\min}^*}\right).$$

The claim (84) now follows by taking the logarithm and defining C appropriately. Then (85) follows as $1/(1-\varepsilon)$ is bounded on an interval of the form $[0, \varepsilon_0]$ for $\varepsilon_0 < 1$. ☐

C Proof of the Main Result

In the following result we will make use of the following simple fact concerning series of random variables.

Lemma C.1 *Let X_k be a sequence of independent, identically distributed, real-valued random variables, satisfying $\mathbb{P}(X_1 = a) = p_1, \mathbb{P}(X_1 = b) = p_2$, where $a < 0 < b, 1 = p_1 + p_2$ and $\mathbf{E}(X_1) =: r < 0$. Let $\{\varepsilon_k\}_{k\in\mathbb{N}}$ be a sequence of positive real numbers, that is square summable, but not summable. Then*

$$\sum_{k=1}^{L} \varepsilon_k X_k \to -\infty\,, \quad a.s.\ as\ L \to \infty\,. \tag{86}$$

Furthermore,

$$\lim_{\ell\to\infty} \sup_{L\geq 1} \sum_{k=\ell}^{\ell+L} \varepsilon_k X_k = 0\,, \quad a.s. \tag{87}$$

Proof As $\mathrm{VAR}(\varepsilon_k X_k) = \varepsilon_k^2 \mathrm{VAR}(X_k)$, the variance of the variables in the series (86) is square summable, and so by [5, Theorem 22.6] the series

$$\sum_{k=1}^{\infty} \varepsilon_k (X_k - \mathbf{E}(X_k))$$

converges almost surely to a finite value. By assumption the series $\sum_{k=1}^{\infty} \varepsilon_k\, \mathbf{E}(X_k) = r \sum_{k=1}^{\infty} \varepsilon_k$ diverges to $-\infty$ and the claim follows. To prove the second claim, consider

$$\sum_{k=\ell}^{\ell+L} \varepsilon_k X_k \leq \sum_{k=\ell}^{\ell+L} \varepsilon_k (X_k - \mathbf{E}(X_k))\,.$$

Again by [5, Theorem 22.6] the partial sums on the right are almost surely partial sums of a convergent series. Then the Cauchy criterion says that there are only finitely many $\ell \in \mathbb{N}$ such that the sum exceeds a given $C > 0$. This shows "≤ 0" in (87). Equality follows from the case $L = 1$. □

In the proof we also need a continuity result extending Lemma 2.2 to the family of Markov chains with fixed probability $z_0 \in \Sigma$. In the following result we use the notation $\mathbb{P}_{z_0}$ to indicate a probability statement for the Markov chain (11) with fixed probability $\lambda = \lambda(z_0)$.

Lemma C.2 *Assume that (A1)–(A3) hold. Consider the family of Markov chains* (11) *with fixed probability $\lambda = \lambda(z_0)$, parametrized by z_0.*
For all $\bar{\delta} > 0$ and all $\theta \in (0, 1]$ there exists an $k_0 \in N$ such that for all $z_0 \in \Sigma$

$$\mathbb{P}_{z_0}\left(\left\|\overline{S}(k_0) - P(z_0)e^\top\right\|_1 > \bar{\delta}\right) < \theta\,. \tag{88}$$

Proof Fix $\varepsilon, \delta > 0$ and $\hat{z}_0 \in \Sigma$. By Lemma 2.2 there exists an $k_0(\hat{z}_0)$ such that

$$\mathbb{P}_{\hat{z}_0}\left(\left\|\overline{S}(k_0) - P(\hat{z}_0)e^\top\right\|_1 > \delta\right) < \varepsilon\,. \tag{89}$$

Now the map $z_0 \mapsto P(z_0)$ is continuous by assumptions (A1) and (A3). Furthermore, the probability of the linear operator $\overline{S}(k)$ varies continuously with z_0. Thus (89) holds on a neighbourhood of $\hat{z}_0$ with the constants $2\delta, 2\varepsilon$ replacing δ, ε. As Σ is compact, it is covered by finally many of such neighbourhoods. With this argument, and as ε, δ are arbitrary, we see that there are finitely many $k_{01}, \ldots, k_{0d}, d \in \mathbb{N}$ such that for every $z_0 \in \Sigma$ there is an k_{0j} such that (89) holds with $k_{0j}, j \in \{1, \ldots, d\}$.
The final claim the follows from an application of Tchebycheff's inequality as follows. For $z_0 \in \Sigma$, $k_0 \in \mathbb{N}$ consider the real valued random variable $X(k_0) = \left\|\overline{S}(k_0) - P(z_0)e^\top\right\|_1$. Note that $X(k_0) \in [0, 2]$, as $\overline{S}(k_0)$ and $P(z_0)e^\top$ are column stochastic. Thus trivially, $\mathrm{VAR}(X(k_0)) \leq 4$. Also, if (89) holds then it follows that

$$\mathbf{E}(X(k_0)) \leq (1 - \varepsilon)\delta + 2\varepsilon\,.$$

We note that the latter inequality is independent of a particular z_0 and just depends on the fact that k_0 is chosen so that (89) holds. Fix $\overline{\delta}, \theta > 0$. With the previous observation we may choose $\varepsilon, \delta > 0$ such that $(1-\varepsilon)\delta + 2\varepsilon < \overline{\delta}/2$. Given $z_0 \in \Sigma$ we may thus choose a k_0 such that (89) holds and it follows that

$$\mathbf{E}(X(k_0)) < \frac{\overline{\delta}}{2}\,. \tag{90}$$

Denoting $\Pi(k) = A(k-1)\dots A(0), k \in \mathbb{N}$ we have for multiples of k_0 that

$$\begin{aligned} X(\ell k_0) &\leq \frac{1}{\ell}\sum_{\nu=0}^{\ell-1}\left\|\frac{1}{k_0}\sum_{j=0}^{k_0-1} A(\nu\ell+j)\cdots A(\nu\ell)\Pi(\nu\ell) - P(z_0)e^\top\right\|_1 \\ &\leq \frac{1}{\ell}\sum_{\nu=0}^{\ell-1}\left\|\frac{1}{k_0}\sum_{j=0}^{k_0-1} A(\nu\ell+j)\cdots A(\nu\ell) - P(z_0)e^\top\right\|_1 \|\Pi(\nu\ell)\|_1 \\ &= \frac{1}{\ell}\sum_{\nu=0}^{\ell-1}\left\|\frac{1}{k_0}\sum_{j=0}^{k_0-1} A(\nu\ell+j)\cdots A(\nu\ell) - P(z_0)e^\top\right\|_1 . \end{aligned}$$

By the independence assumption on the $A(j)$ we see that the final term is the average of ℓ independent copies of $X(k_0)$, the variance of which is bounded by $4/\ell$. Let $\{X_\nu(k_0)\}_{\nu\in\mathbb{N}}$ be a sequence of independent copies of $X(k_0)$. It follows from Tchebycheff's inequality that for all $\ell \geq 4\sqrt{2}/(\theta\sqrt{\overline{\delta}})$ we have

$$\mathbb{P}\left(\left\|\overline{S}(\ell k_0) - P(z_0)e^\top\right\|_1 > \overline{\delta}\right) \leq \mathbb{P}\left(\frac{1}{\ell}\sum_{\nu=0}^{\ell-1} X_\nu(k_0) > \mathbf{E}(X(k_0)) + \frac{\overline{\delta}}{2}\right) < \theta\,.$$

As the previous argument only depends on the validity of (89), it holds uniformly for all $z_0 \in \Sigma$ for which the choice of k_0 guarantees (89). The proof is completed, by choosing a sufficiently large common multiple of $k_{01}, \dots, k_{0d}$. ☐

Proof (of Theorem 4.6) In the proof, we make extensive use of the deterministic system discussed in Appendix B. We will show that for T sufficiently large the behaviour of $\bar{z}(T)$ is well approximated by the deterministic system.

We assume that the constants δ^-, δ^+ from (52), (53) have been fixed. We will use the notation $z(T; z_0)$, resp. $\bar{z}(T; z_0)$ to indicate the initial condition for the random variable $z(T)$, resp. its second component vector $\bar{z}(T)$. Similarly, the notation $z(T+m; z(T))$ indicates the conditioning of $z(T+m)$ on a certain value at time T, etc.
Fix $\eta > 0$. We aim to show that almost surely the sample path $\bar{z}(T) \in B_1(x^*, \eta)$ for all T large enough. As $\eta > 0$ is arbitrary this will show the claim.
To attain our goal, we perform the following sequence of choices:

(i) For the constant η pick $\varepsilon_0 > 0$ and $C_\eta > 0$ according to Corollary B.4, so that (77) is satisfied for all $0 < \varepsilon < \varepsilon_0$.

(ii) Let C_0 be the constant guaranteed by Lemma B.7 satisfying (85) for all $0 < \varepsilon < \varepsilon_0$.

(iii) Let $K > 0$ be the constant given by Lemma B.5.

(iv) Choose δ^*, $\bar{\delta}$ according to (80), so that

$$\delta^* := \min\left\{\frac{C_\eta e^\eta}{2K}, \delta^-\right\} \quad \text{and} \quad 0 < \bar{\delta} < \delta^*/3\,, \tag{91}$$

so that Corollary B.6 and Lemma B.2 (iii) are applicable. Let $C_{\bar{\delta}} > 0$ be the constant guaranteed by Lemma B.2 (iii).

(v) Pick $\theta \in (0,1)$ so that

$$-(1-\theta) + \theta(1 + C_{\bar{\delta}}) < 0\,, \quad \text{and} \quad -(1-\theta)C_\eta + \theta C_0 < 0\,.$$

(vi) We now appeal to Lemma C.2 to determine the length of the (short) averaging period discussed in the preamble.

Using Lemma C.2 and (15), pick $m \in \mathbb{N}$ such that for all $z_0 \in \Sigma$ the Markov chain (11) with fixed probability $\lambda = \lambda(z_0)$ satisfies for all x_0 in Σ that

$$\mathbb{P}\left(\left\|\frac{1}{m}\sum_{t=1}^{m} x(t; x_0) - P(z_0)\right\|_1 > \bar{\delta}\right) < \frac{\theta}{2}\,. \tag{92}$$

(vii) Pick $T_0 \in \mathbb{N}$ such that for all $T \geq T_0$ we have $m/(T+m) < \varepsilon_0$ and so that for the Markov chain (38) with place dependent probabilities we have for the average of the first component that

$$\mathbb{P}\left(\left\|\frac{1}{m}\sum_{t=1}^{m} z_1(T+t; [z_1(T), \bar{z}(T)]) - P(\bar{z}(T))\right\|_1 > \bar{\delta}\right) < \theta\,. \tag{93}$$

This is possible as this inequality is a perturbed version of (92): indeed, with increasing T the variation of $\bar{z}(T+t), 1 \leq t \leq m$ (i.e. in the first m steps after time T) becomes arbitrarily small. More precisely, for $t = 1\ldots, m$ we have by definition

$$\|\bar{z}(T) - \bar{z}(T+t)\|_1 \leq \frac{2m}{T}\,.$$

Thus as $T \to \infty$ the place dependent probabilities of the Markov chain (38) that are considered on the interval $[T, T+m]$ converge to the fixed probabilities $\lambda(\bar{z}(T))$. The claim then follows from (92) by continuity of the probability functions λ_i (see (A1)).

Let $T \geq T_0$, so that by construction $\varepsilon := m/(T+m) < \varepsilon_0$. We will study the evolution of the value $\bar{z}(T+km) \mapsto \bar{z}(T+(k+1)m)$, $k \in \mathbb{N}$. This is given by

$$\bar{z}(T+(k+1)m) = \frac{T+km}{T+(k+1)m}\bar{z}(T+km) + \frac{m}{T+(k+1)m}\left(\frac{1}{m}\sum_{t=1}^{m} z_1(T+km+t; z(T+km))\right).$$

For ease of notation we define $P_{\text{co}}(\bar{\delta}) := \operatorname{conv}(P(\Sigma) + \overline{B}_1(0,\bar{\delta}))$, $\tau(k) := T+km$ and $\varepsilon_k := m/\tau(k+1)$, so that the previous equation can be expressed as[6]

$$\bar{z}(\tau(k+1)) = (1-\varepsilon_k)\bar{z}(\tau(k)) + \varepsilon_k\left(\frac{1}{m}\sum_{t=1}^{m} z_1(\tau(k)+t; z(\tau(k)))\right). \tag{94}$$

At this point the reader should recognize the structure of the discrete iteration we have analysed in Section B and notice that by (vii) we have a high probability that $\bar{z}(\tau(k+1))$ is close to $R_{\varepsilon_k}(\bar{z}(\tau(k)))$.

Note that the constants have been chosen so that both Lemma B.2 and Corollary B.6 are applicable. We will use the estimates obtained in Lemma B.2 to show that trajectories starting in Σ will reach the set $P_{\text{co}}(2\bar{\delta})$ in a finite number of steps. We then show that for trajectories starting in the strict superset $P_{\text{co}}(3\bar{\delta})$ the estimates of Corollary B.6 yield that we reach the set $B_H(x^*,\eta)$ again in a finite number of steps; almost surely.

Step 1: More precisely, we will first show that (the first hitting time)

$$\sigma_1(z(T)) := \min\{k \in \mathbb{N}\,;\ \bar{z}(\tau(k); z(T)) \in P_{\text{co}}(2\bar{\delta})\ \}$$

is almost surely finite. Obviously, if $\bar{z}(T) \in P_{\text{co}}(2\bar{\delta})$ there is nothing to show. Appealing to Lemma B.2 (i) and the choice made in (vii), we have that if $\operatorname{dist}_1(\bar{z}(\tau(k)), P_{\text{co}}(\bar{\delta})) > \bar{\delta}$, then

$$\mathbb{P}\Big(\operatorname{dist}_1(\bar{z}(\tau(k+1)), P_{\text{co}}(\bar{\delta})) \leq (1-\varepsilon_k)\ \operatorname{dist}_1(\bar{z}(\tau(k)), P_{\text{co}}(\bar{\delta}))\Big) \geq 1-\theta\,.$$

[6]This chain is not a typical stochastic approximation chain due to the interdependence between z_1 and $\bar{z}$—cf. (38) (28). It is similar in form to the two-timescales process of [6, Chapter 6], but does not satisfy the convergence conditions therein.

In the complementary event, which happens with probability of at most θ we have by Lemma B.2 that

$$\mathrm{dist}_1(\bar{z}(\tau(k+1)), P_{\mathrm{co}}(\bar{\delta})) \le (1 + \varepsilon_k C_{\bar{\delta}})\, \mathrm{dist}_1(\bar{z}(\tau(k)), P_{\mathrm{co}}(\bar{\delta}))$$

Combining these two observations we see that for $\tau(k) < \sigma_1(\bar{z}(T))$ we have that

$$\mathrm{dist}_1(\bar{z}(\tau(k)), P_{\mathrm{co}}(\bar{\delta})) \le \left(\prod_{\ell=1}^{k} a_\ell \right) \mathrm{dist}_1(\bar{z}(T), P_{\mathrm{co}}(\bar{\delta})), \tag{95}$$

where $a_\ell, \ell \in \mathbb{N}$ is a random variable that has the value $(1 - \varepsilon_\ell)$ with probability $1 - \theta$ and the value $(1 + \varepsilon_\ell C_\delta)$ with probability θ. By construction the random variables a_ℓ are independent, as the bounds obtained do not depend on the particular sample path of the Markov chain.
To be able to apply Lemma C.1 we first note that for all ℓ large enough we have $\log(1+\varepsilon_\ell C_\delta) < \varepsilon_\ell(1+C_\delta)$. By the choice of θ in (vi), we obtain for all ℓ large enough that $\mathbf{E}(\log a_\ell) \le (-(1-\theta) + \theta(1+C_\delta))\, \varepsilon_\ell$. Lemma C.1 thus implies that $\sum_{\ell=1}^{k} \log a_\ell \to -\infty$, almost surely. Thus almost surely we have

$$\lim_{k\to\infty} \mathrm{dist}_1(\bar{z}(\tau(k)), P_{\mathrm{co}}(\bar{\delta})) = 0\,,$$

provided that $\bar{z}(\tau(k)) \notin P_{\mathrm{co}}(2\bar{\delta})$ for all k. This is of course impossible, and so almost surely $\bar{z}(\tau(k)) \in P_{\mathrm{co}}(2\bar{\delta})$ for a finite k.

Step 2: Similarly, if $\bar{z}(\tau(k)) \in P_{\mathrm{co}}(3\bar{\delta})$, then by Corollary B.6 and the choice made in (vii) we have

$$\mathbb{P}\left(d_H(\bar{z}(\tau(k+1)), x^*) < d_H(\bar{z}(\tau(k)), x^*) - \frac{C_\eta}{2}\varepsilon_k \right) \ge 1 - \theta\,.$$

On the other hand with probability of at most θ we have by Lemma B.7 that

$$d_H(\bar{z}(\tau(k+1)), x^*) \le d_H(\bar{z}(\tau(k)), x^*) + C_0 \varepsilon_k\,. \tag{96}$$

In a similar fashion to the first step, as long as $\bar{z}(\tau(k)) \in P_{\mathrm{co}}(3\bar{\delta})$ and $d_H(\bar{z}(\tau(k)), x^*) > \eta$, we have

$$d_H(\bar{z}(\tau(\sigma_1(\bar{z}(T)+k)), x^*) \le d_H(\bar{z}(\sigma_1(\bar{z}(T)+k)), x^*) + \sum_{\ell=1}^{k} b_\ell\,, \tag{97}$$

where b_ℓ is a random variable that takes the value $-\varepsilon_\ell C_\eta$ with probability $(1-\theta)$ and the value $\varepsilon_\ell C_0$ with probability θ. As before, Lemma C.1 ensures that $\sum_\ell b_\ell$ diverges to $-\infty$, almost surely. Note that it is always possible to leave the set $P_{\mathrm{co}}(3\bar{\delta})$ with a small probability. In this case Step 1 can be applied again, so that we re-enter the set $P_{\mathrm{co}}(2\bar{\delta})$, almost surely. Now by (95) the process of entering $P_{\mathrm{co}}(2\bar{\delta})$ and subsequently leaving $P_{\mathrm{co}}(3\bar{\delta})$ requires that for some partial sum we have

$$\sum_{k=\ell}^{\ell+L} \log(a_k) \ge \bar{\delta}\,.$$

By Lemma C.1, with probability 1, this happens only a finite number of times. Consequently, almost surely a sample path will reach $B_H(x^*, \eta)$.

Step 3: Finally, to obtain almost sure convergence, we need to show that almost surely

$$\bar{z}(\tau(k)) \in B_H(x^*, \eta)\,, \quad \text{for all } k \text{ large enough.} \tag{98}$$

To this end we repeat the choices made in (i) - (vii) for the value $\eta/2$. Thus we can conclude that almost surely a sample path enters $B_H(x^*, \eta/2)$. If we assume that the sample path leaves $B_H(x^*, \eta)$ at some later time, then again by Steps 1 and 2 it will almost surely re-enter $B_H(x^*, \eta/2)$. The question is thus whether it is possible that infinitely often the sample path exits the ball $B_H(x^*, \eta)$ given that it was previously within the ball $B_H(x^*, 3\eta/4)$. In view of (97) this amounts to saying that

$$\sum_{k=\ell}^{\ell+L} b_k > \frac{\eta}{4}$$

for pairs $(\ell, L) \in \mathbb{N}^2$ with arbitrarily large ℓ. By Lemma C.1 this almost surely does not happen. This shows (98). The proof is complete by noting that the small variations of $\bar{z}$ on the intervals $\tau(k), \ldots, \tau(k+1)$ do not destroy stability. Indeed, if $\bar{z}(\tau(k)) \in B_H(x^*, \eta/2)$ for all k large enough, then also $\bar{z}(\tau(k)+j) \in B_H(x^*, \eta)$ for $j = 1, \ldots, m$, provided k is large enough. □

References

[1] D. Angeli and P.-A. Kountouriotis. A stochastic approach to dynamic-demand refrigerator control. *IEEE Transactions on Control Systems Technology*, 20(3):581 –592, May 2012.

[2] M. Barnsley, S. Demko, J. Elton, and J. Geronimo. Invariant measures for Markov processes arising from iterated function systems with place-dependent probabilities. *Annales de l'Institut Henri Poincaré (B) Probabilités et Statistiques*, 24(3):367–394, 1988.

[3] P. Bianchi, G. Fort, and W. Hachem. Performance of a distributed stochastic approximation algorithm. *IEEE Transactions on Information Theory*, 59(11):7405–7418, Nov 2013.

[4] B. Biegel, P. Andersen, T. Pedersen, K. Nielsen, J. Stoustrup, and L. Hansen. Smart grid dispatch strategy for on/off demand-side devices. In *2013 European Control Conference (ECC)*, pages 2541–2548, July 2013.

[5] P. Billingsley. *Probability and Measure*. John Wiley & Sons, 1995. 3rd ed.

[6] V. S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.

[7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[8] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl. Demand dispatch. *IEEE Power and Energy Magazine*, 8(3):20 – 29, May - June 2010.

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

[10] K. Clement, E. Haesen, and J. Driesen. Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids. In *IEEE/PES Power Systems Conference and Exposition (PSCE).*, pages 1 – 7, March 2009.

[11] K. Clement-Nyns, E. Haesen, and J. Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, 25(1):371 – 380, February 2010.

[12] M. Corless and R. Shorten. Analysis of a general nonlinear increase resource allocation algorithm with application to network utility maximization. submitted, March 2014.

[13] E. Crisostomi, M. Liu, M. Raugi, and R. Shorten. Stochastic distributed algorithms for power generation in a microgrid. *IEEE Transactions on Smart Grid*, submitted in 2014.

[14] S. Deilami, A. Masoum, P. Moses, and M. Masoum. Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile. *IEEE Transactions on Smart Grid*, 2(3):456 – 467, September 2011.

[15] J. Duchi, A. Agarwal, and M. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606, 2012.

[16] J. H. Elton. An ergodic theorem for iterated maps. *Ergodic Theory Dynam. Systems*, 7(4):481–488, 1987.

[17] L. P. Fernández, T. S. Róman, R. Cossent, C. Domingo, and P. Frías. Assessment of the impact of plug-in electric vehicles on distribution networks. *IEEE Transactions on Power Systems*, 26(1):206 – 213, February 2011.

[18] P. Finn, C. Fitzpatrick, and M. Leahy. Increased penetration of wind generated electricity using real time pricing; demand side management. In *IEEE International Symposium on Sustainable Systems and Technology (ISSST)*, pages 1 – 6, May 2009.

[19] S. Floyd. High speed TCP for large congestion windows. Technical report, Internet draft draft-floyd-tcp-highspeed-02.txt, work in progres, February 2003.

[20] B. Hajek and G. Gopalakrishnan. A framework for studying demand in hierarchical networks. Preliminary draft, 2004.

[21] A. Harris. Charge of the electric car. *Engineering and Technology*, 4(10), June 2000.

[22] D. J. Hartfiel. *Nonhomogeneous Matrix Products*. World Scientific, Singapore, 2002.

[23] V. Jacobson. Congestion avoidance and control. In *Proc. of SIGCOMM*, 1988.

[24] B. Johansson, M. Rabi, and M. Johansson. A randomized incremental subgradient method for distributed optimization in networked systems. *SIAM Journal on Control and Optimization*, 20(3):1157–1170, 2009.

[25] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.

[26] K. Kar, S. Sarkar, and L. Tassiulas. A simple rate control algorithm for max total user utility. In *Proc. 2001 IEEE INFOCOM, Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1, pages 133–141 vol.1, 2001.

[27] T. Kelly. On engineering a stable and scalable TCP variant. Technical report, Cambridge University Engineering Department Technical Report CUED/F-INFENG/TR.435, 2002.

[28] S. Kumar, S.-J. Park, and S. Sitharama Iyengar. A loss-event driven scalable fluid simulation method for high-speed networks. *Computer Networks*, 54(1):112 – 132, 2010.

[29] N. Li, L. Chen, and S. Low. Optimal demand response based on utility maximization in power networks. In *2011 IEEE Power and Energy Society General Meeting.*, pages 1 – 8, July 2011.

[30] S. Low, F. Paganini, and J. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 32(1):28–43, 2002.

[31] S. H. Low, L. L. Peterson, and L. Wang. Understanding TCP Vegas: a duality model. *Journal of the ACM (JACM)*, 49(2):207–235, 2002.

[32] G. Mann, R. T. McDonald, M. Mohri, N. Silberman, and D. Walker. Efficient large-scale distributed training of conditional maximum entropy models. *Advances in Neural Information Processing Systems*, 22:1231–1239, 2009.

[33] J. Masaki, G. Nishantha, and Y. Hayashida. Development of a high-speed transport protocol with tcp-reno friendliness. *International Conference on Advanced Communication Technology, ICACT*, 1:174 – 179, 2010.

[34] J. Mathieu, M. Kamgarpour, J. Lygeros, and D. Callaway. Energy arbitrage with thermostatically controlled loads. In *2013 European Control Conference (ECC)*, pages 2519–2526, July 2013.

[35] S. Molnar, S. Balazs, and T. Tuan. A comprehensive TCP fairness analysis in high speed networks. *Computer Communications*, 32:1460–1484, 2009.

[36] A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

[37] G. Putrus, P. Suwanapingkarl, D. Johnston, E. Bentley, and M. Narayana. Impact of electric vehicles on power distribution networks. In *IEEE Vehicle Power and Propulsion Conference (VPPC).*, pages 827 –831, September 2009.

[38] S. S. Ram, A. Nedić, and V. V. Veeravalli. Distributed stochastic subgradient projection algorithms for convex optimization. *Journal of optimization theory and applications*, 147(3):516–545, 2010.

[39] U. G. Rothblum and R. Shorten. Nonlinear aimd congestion control and contraction mappings. *SIAM Journal on Control and Optimization*, 46(5):1882–1896, 2007.

[40] S. Sayeef, S. Heslop, D. Cornforth, T. Moore, S. Percy, J. K. Ward, A. Berry, and D. Rowe. Solar intermittency: Australia's clean energy challenge, June 2012. accessed at `"http://www.csiro.au/Organisation-Structure/Flagships/Energy-Flagship/Solar-Intermittency-Report.aspx"` on 5.3.2014.

[41] S. Shafiei, H. Rasmussen, and J. Stoustrup. Model predictive control for a thermostatic controlled system. In *2013 European Control Conference (ECC)*, pages 1559–1564, July 2013.

[42] R. Shorten, C. King, F. Wirth, and D. Leith. Modelling TCP congestion control dynamics in drop-tail environments. *Automatica*, 43(3):441–449, 2007.

[43] R. Shorten, F. Wirth, and D. Leith. A positive systems model of TCP-like congestion control: Asymptotic analysis. *IEEE/ACM Transactions on Networking*, 14:616–629, 2006.

[44] R. Srikant. *Internet congestion control*, volume 14 of *Control theory*. Birkhäuser Boston Inc., Boston, MA, 2004.

[45] R. Stanojevic and R. Shorten. Distributed dynamic speed scaling. In *Proc. 2010 IEEE INFOCOM*, pages 1–5, San Diego, CA, March 2010.

[46] S. Stüdli, E. Crisostomi, R. Middleton, and R. Shorten. A flexible distributed framework for realising electric and plug-in hybrid vehicle charging policies. *International Journal of Control*, 85(8):1130 – 1145, 2012.

[47] K. Tan, J. Song, Q. Zhang, and M. Sridharan. A compound TCP approach for high-speed and long distance networks. Technical report, Microsoft Technical Report, 2007.

[48] F. Wirth, R. Stanojevic, R. Shorten, and D. Leith. Stochastic equilibria of AIMD communication networks. *SIAM Journal on Matrix Analysis and Applications*, 28(3):703–723, 2006.

[49] L. Xu, K. Harfoush, and I. Rhee. Binary increase congestion control (BIC) for fast long-distance networks. In *Proc. of IEEE INFOCOM 2004*, pages 2514–2524, Hong Kong, 2004.

[50] J. Y. Yu and S. Mannor. Asymptotics of efficiency loss in competitive market mechanisms. In *Proceedings IEEE INFOCOM*, 2006.

[51] M. Zinkevich, M. Weimer, A. J. Smola, and L. Li. Parallelized stochastic gradient descent. In *Advances in Neural Information Processing Systems*, volume 23, pages 2595–2603, 2010.