Distributed Approximate Uniform Global Minimum Sharing

Michelangelo Bin, Thomas Parisini

Imperial College London

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE).

EPSRC

UNIVERSITÀ DEGLI STUDI DI TRIESTE

KIOS

OUTLINE

- Some Motivating Applications
- Problem Formulation and Requirements
- Analysis, Comparisons, and Main Result
- Numerical Results
- A Real Industrial Use Case: Early Results

Reference on background theory and analysis

Joint work with **Michelangelo Bin**:

M. Bin, T. Parisini, ``A Distributed Methodology for Approximate Uniform Global Minimum Sharing," *Automatica*, vol. 131, 109777, 2021 https://doi.org/10.1016/j.automatica.2021.109777 (Open Access)

Automatica 131 (2021) 109777

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

ELSEVIER

A distributed methodology for approximate uniform global minimum sharing ☆

Michelangelo Bin [a,*], Thomas Parisini [a,b,c]

[a] Department of Electrical and Electronic Engineering, Imperial College London, London, UK
[b] KIOS Research and Innovation Center of Excellence, University of Cyprus, Cyprus
[c] Department of Engineering and Architecture at University of Trieste, Trieste, Italy

ARTICLE INFO

Article history:
Received 21 April 2020
Received in revised form 23 March 2021
Accepted 13 May 2021
Available online xxxx

Keywords:
Distributed minimum sharing
Max-consensus
Distributed optimization
Approximate convergence

ABSTRACT

The paper deals with the distributed minimum sharing problem: a set of decision-makers compute the minimum of some local quantities of interest in a distributed and decentralized way by exchanging information through a communication network. We propose an adjustable approximate solution which enjoys several properties of crucial importance in applications. In particular, the proposed solution has good decentralization properties and it is scalable in that the number of local variables does not grow with the size or topology of the communication network. Moreover, a global and uniform (both in the initial time and in the initial conditions) asymptotic stability result is provided towards a steady state which can be made arbitrarily close to the sought minimum. Exact asymptotic convergence can be recovered at the price of losing uniformity with respect to the initial time.

1. Introduction

1.1. Problem description, objectives and context

We consider the problem of computing the minimum of a set of numbers over a network, and we propose a distributed, iterative solution achieving *global* and *uniform*, albeit *approximate*, asymptotic stability. We are given a set $\mathcal{N}$ of N decision makers (or *agents*), where each agent $i \in \mathcal{N}$ is provided with a number $\mathrm{M}_i \in \mathbb{R}_{\geq 0}$ not known a priori by the others. The agents exchange information over a communication network with only a subset of other agents (called their *neighborhood*). The approximate minimum sharing problem consists in the design of an algorithm guaranteeing that each agent asymptotically obtains a "sufficiently good" estimate of the quantity

$$\mathrm{M}^\star := \min_{i\in\mathcal{N}} \mathrm{M}_i. \tag{1}$$

Clearly, "$x_i = \mathrm{M}^\star, \ \forall i \in \mathcal{N}$" is also the *unique* solution to every constrained optimization problem of the form

$$\begin{aligned} &\max \sum_{i\in\mathcal{N}} \psi_i(x_i) \\ &x_i \leq \mathrm{M}_i, \qquad \forall i \in \mathcal{N} \\ &x_i = x_j, \qquad \forall i,j \in \mathcal{N} \end{aligned} \tag{2}$$

obtained with ψ_i, $i \in \mathcal{N}$, continuous and strictly increasing functions. Therefore, the minimum sharing problem is equivalent to the constrained distributed optimization problem (2), thus intersecting the wide research field of distributed optimization (Notarstefano, Notarnicola, & Camisa, 2019).

The problem of computing a minimum (or, equivalently, a maximum) over a network of decision makers is a classical problem in multi-agent control, with applications in distributed estimation and filtering, synchronization, leader election, and computation of network size and connectivity (see, e.g., Bullo, Cortes, and Martinez (2009), Golfar and Ghaisari (2019), Iutzeler, Ciblat, and Jakubowicz (2012), Nejad, Attia, and Raisch (2009), Santoro (2006) and the references therein). Perhaps the most elementary existing algorithms solving the minimum sharing problem are the *FloodMax* (Bullo et al., 2009) and the *Max-Consensus* (Golfar & Ghaisari, 2019; Iutzeler et al., 2012; Nejad et al., 2009). In its simplest form, Max-Consensus[1] requires each

☆ This work has been partially supported by European Union's Horizon 2020 research and innovation programme under grant agreement no. 739551 (KIOS CoE) and by the Italian Ministry for Research in the framework of the 2017 Program for Research Projects of National Interest (PRIN), Grant No. 2017YKXYXJ. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Vijay Gupta under the direction of Editor Christos G. Cassandras.

* Corresponding author.
E-mail addresses: m.bin@imperial.ac.uk (M. Bin), t.parisini@imperial.ac.uk (T. Parisini).

[1] For brevity, we only focus on Max-Consensus. However, the same conclusions applies also to the FloodMax.

https://doi.org/10.1016/j.automatica.2021.109777

MOTIVATING APPLICATIONS

Damping of Traffic Waves ("Stop & Go")

Problem: N self-driving cars must control, in a *decentralised* fashion, a traffic network of mixed human-driven and self-driven vehicles in such a way that there are no "stop-and-go" or other oscillatory behaviours

- No matter which control law self-driving cars will use, eventually they shall agree on a common cruise speed, which is the maximal possible consistent with all constraints
- Each self-driving car has its own maximal speed M_i , *unknown a priori to the others*, and defined by:
  - Its own mechanical or technological constraints
  - Common traffic rules
  - Possibly changing traffic conditions ahead (e.g., the vehicle in front is slow)
- Eventually, the problem boils down to compute, in a *decentralised* way, the quantity

$$M^\star := \min_i M_i$$

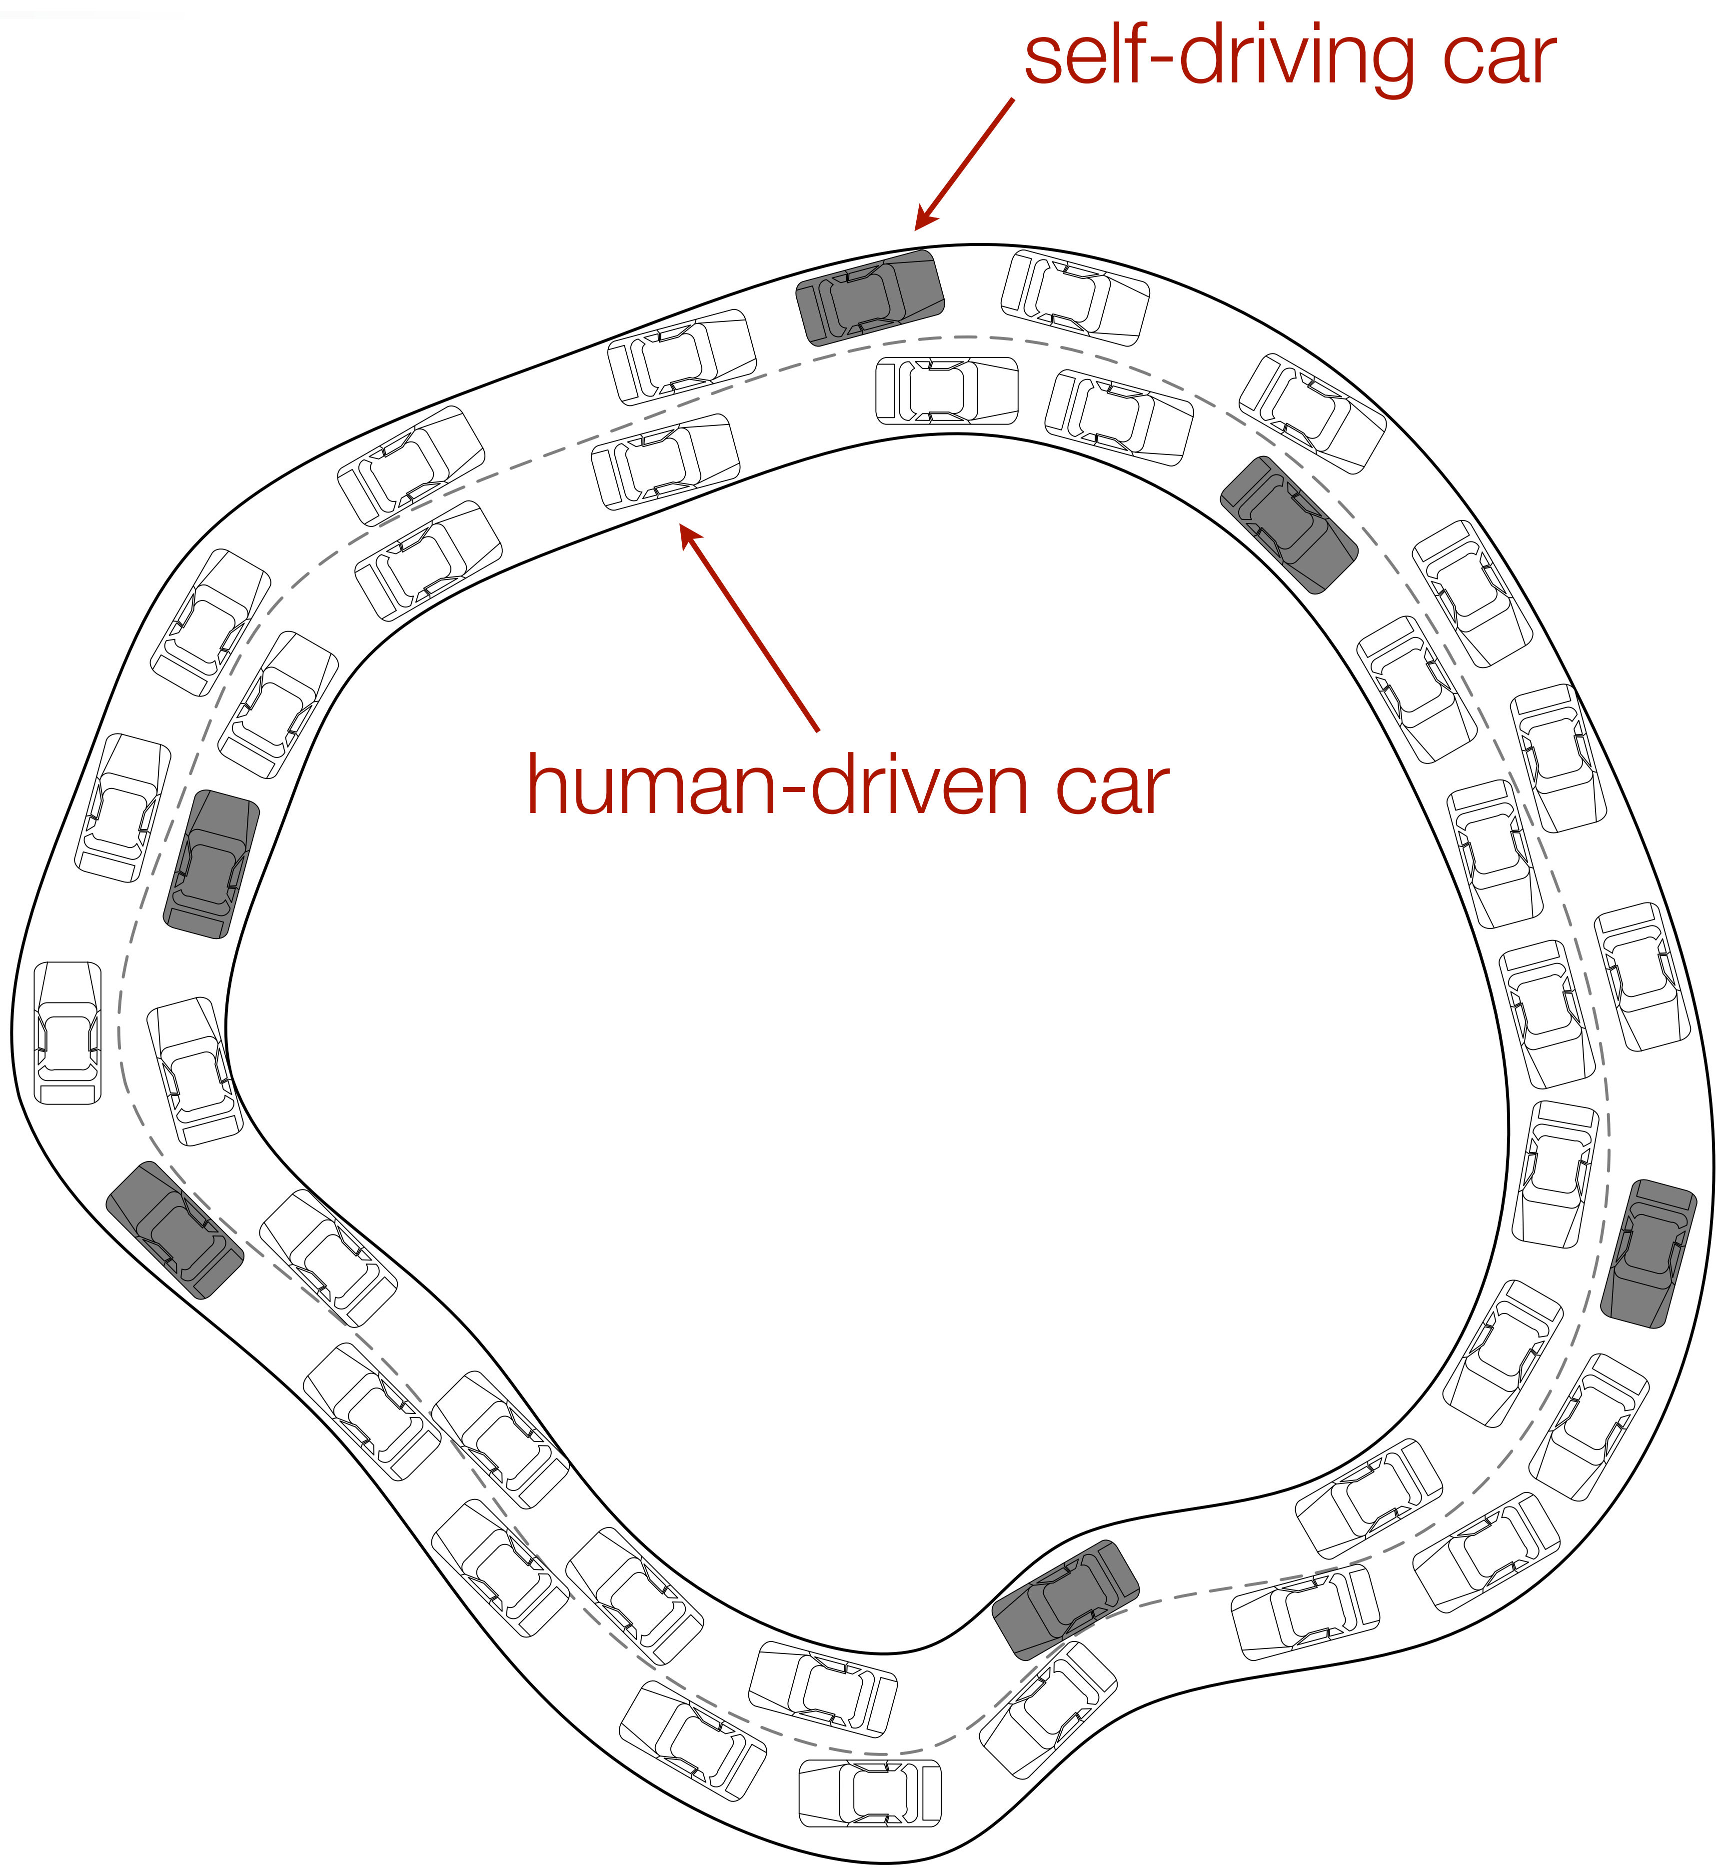

MOTIVATING APPLICATIONS

Leader Election in Swarms

Problem: N agents need to elect (one or more) leader(s) in a *decentralised* way. Typically, each agent owns a number M_i, *unknown a priori to the others*, that may represent, for instance:

- A unique ID
- A rank obtained in the pursuit of some task

- Typical sources of variability in swarms are:
  - The ranks change
  - The topology of communication changes
- A common way to elect a set of leaders in swarms is to choose the ones having the least number M_i. This amounts again to compute in a *decentralised* way:

$$M^\star := \min_i M_i$$

MOTIVATING APPLICATIONS

The COVID-19 Testing Problem

Problem: Who is it best to test first if we want to hinder contagion?

- Assume we are given an interaction network (e.g., obtained from bluetooth data) representing people's recurrent contacts.
- With each node we can associate a ``cost" M_i representing the societal damage if the person it represents gets infected
  - For example:

$$M_i = -\mathrm{PageRank}(i)$$

- The problem again boils down to compute in a *decentralised* way:

$$M^\star := \min_i M_i$$

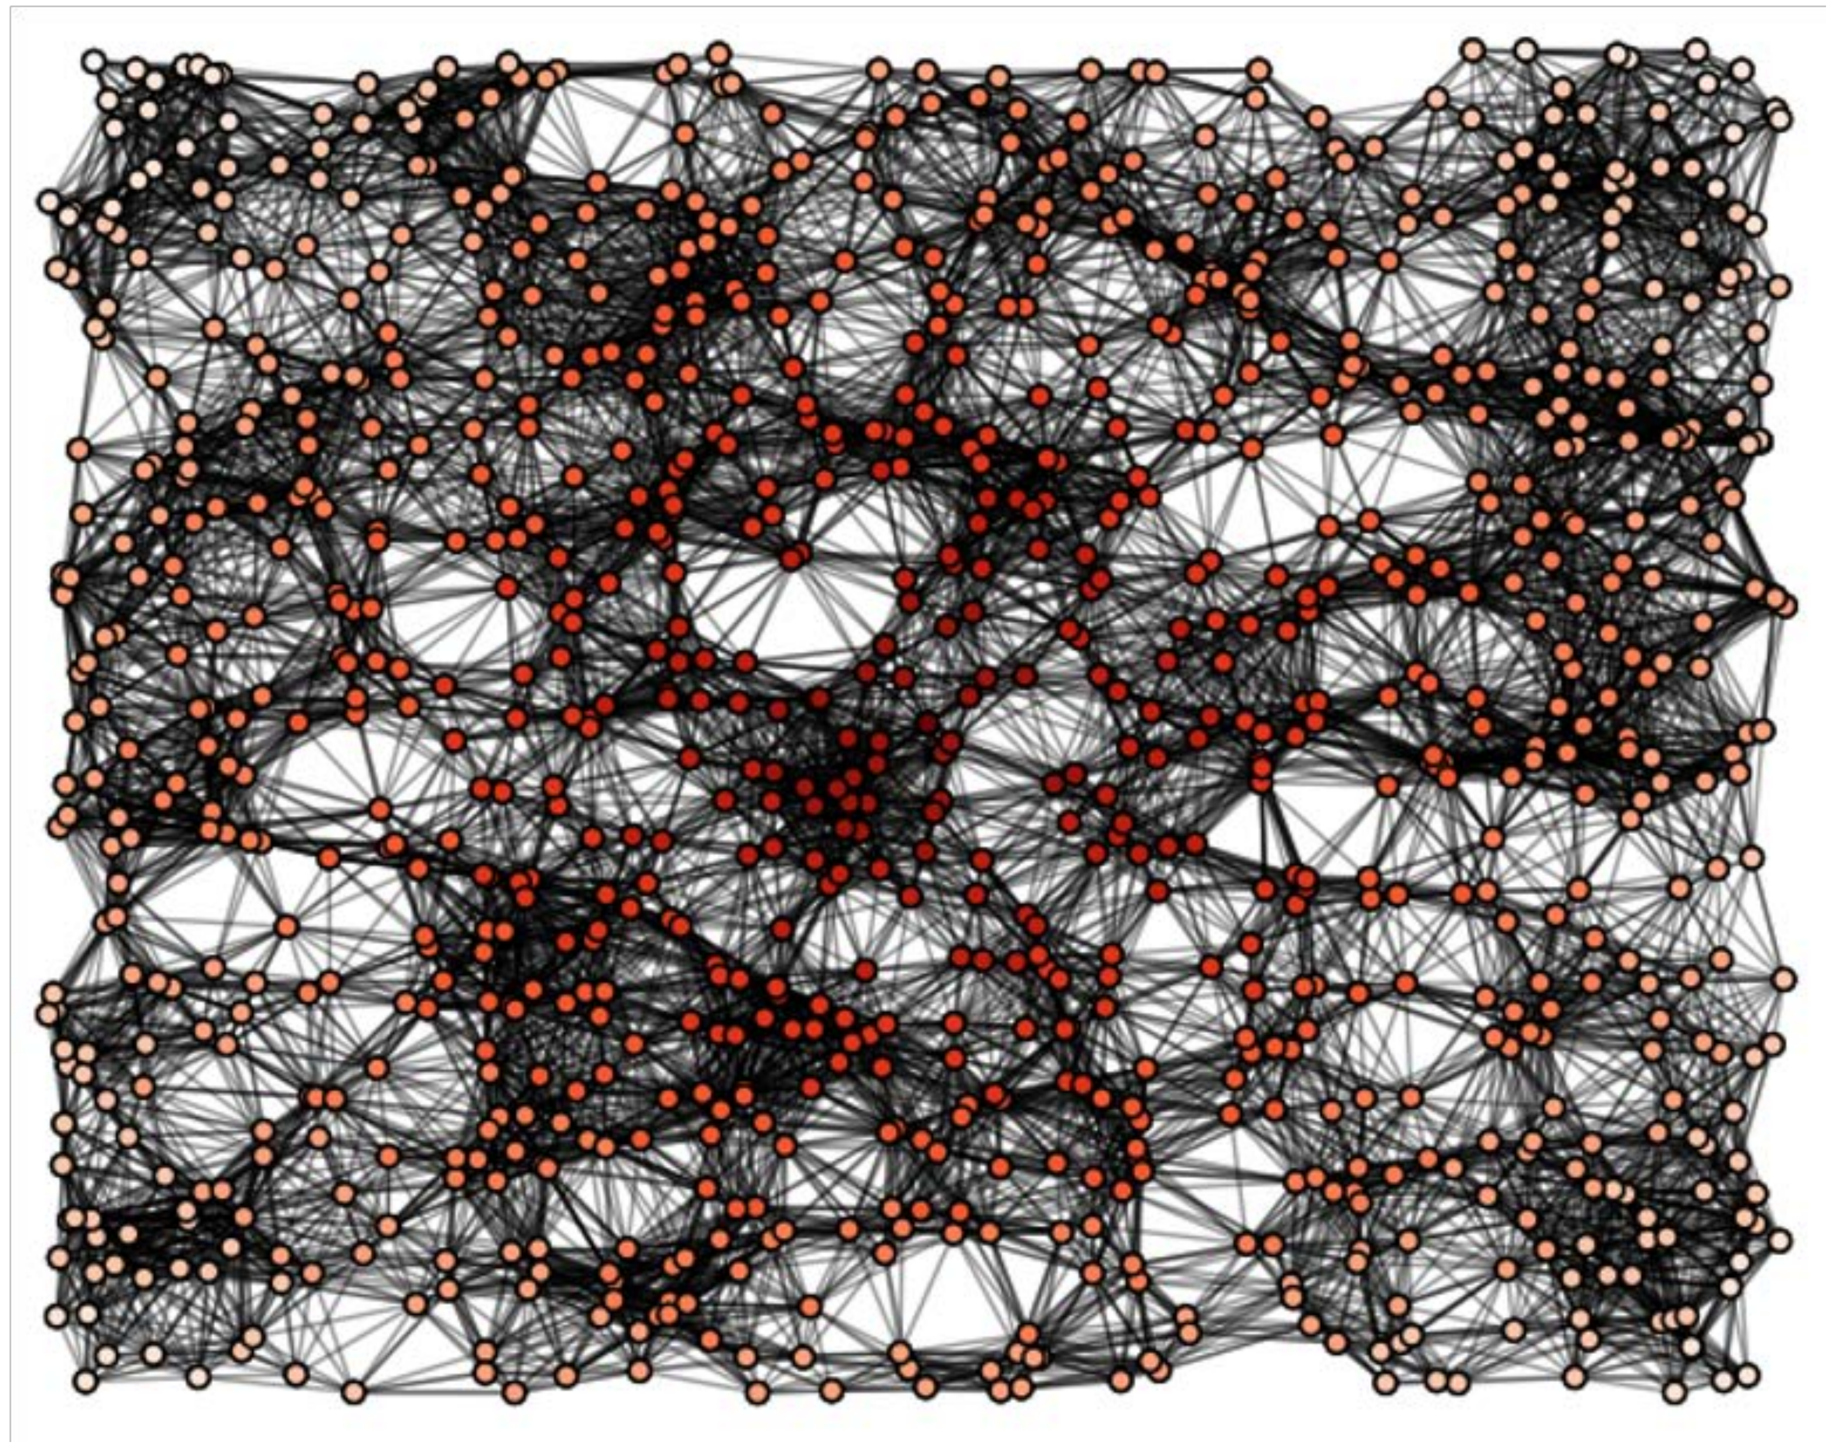

MOTIVATING APPLICATIONS

Actuators Synchronisation (Industry 4.0)

Problem: N ``smart electrical actuators" connected in series need to *agree* on a common speed, which must be as high as possible to maximise production/reduce energy losses

- Each drive has its own maximal reference velocity M_i which depends on several factors such as:
  - Technological constraints (torque, power, rate limiters, etc.)
  - The load resistant torque on which the actuator is working
- A solution that is robust allows for:
  - Automatic change-overs
  - Adaptation to different loads and process setup
  - Plug-and-play replacements
- The problem again boils down to compute in a *decentralised* way:

$$M^\star := \min_i M_i$$

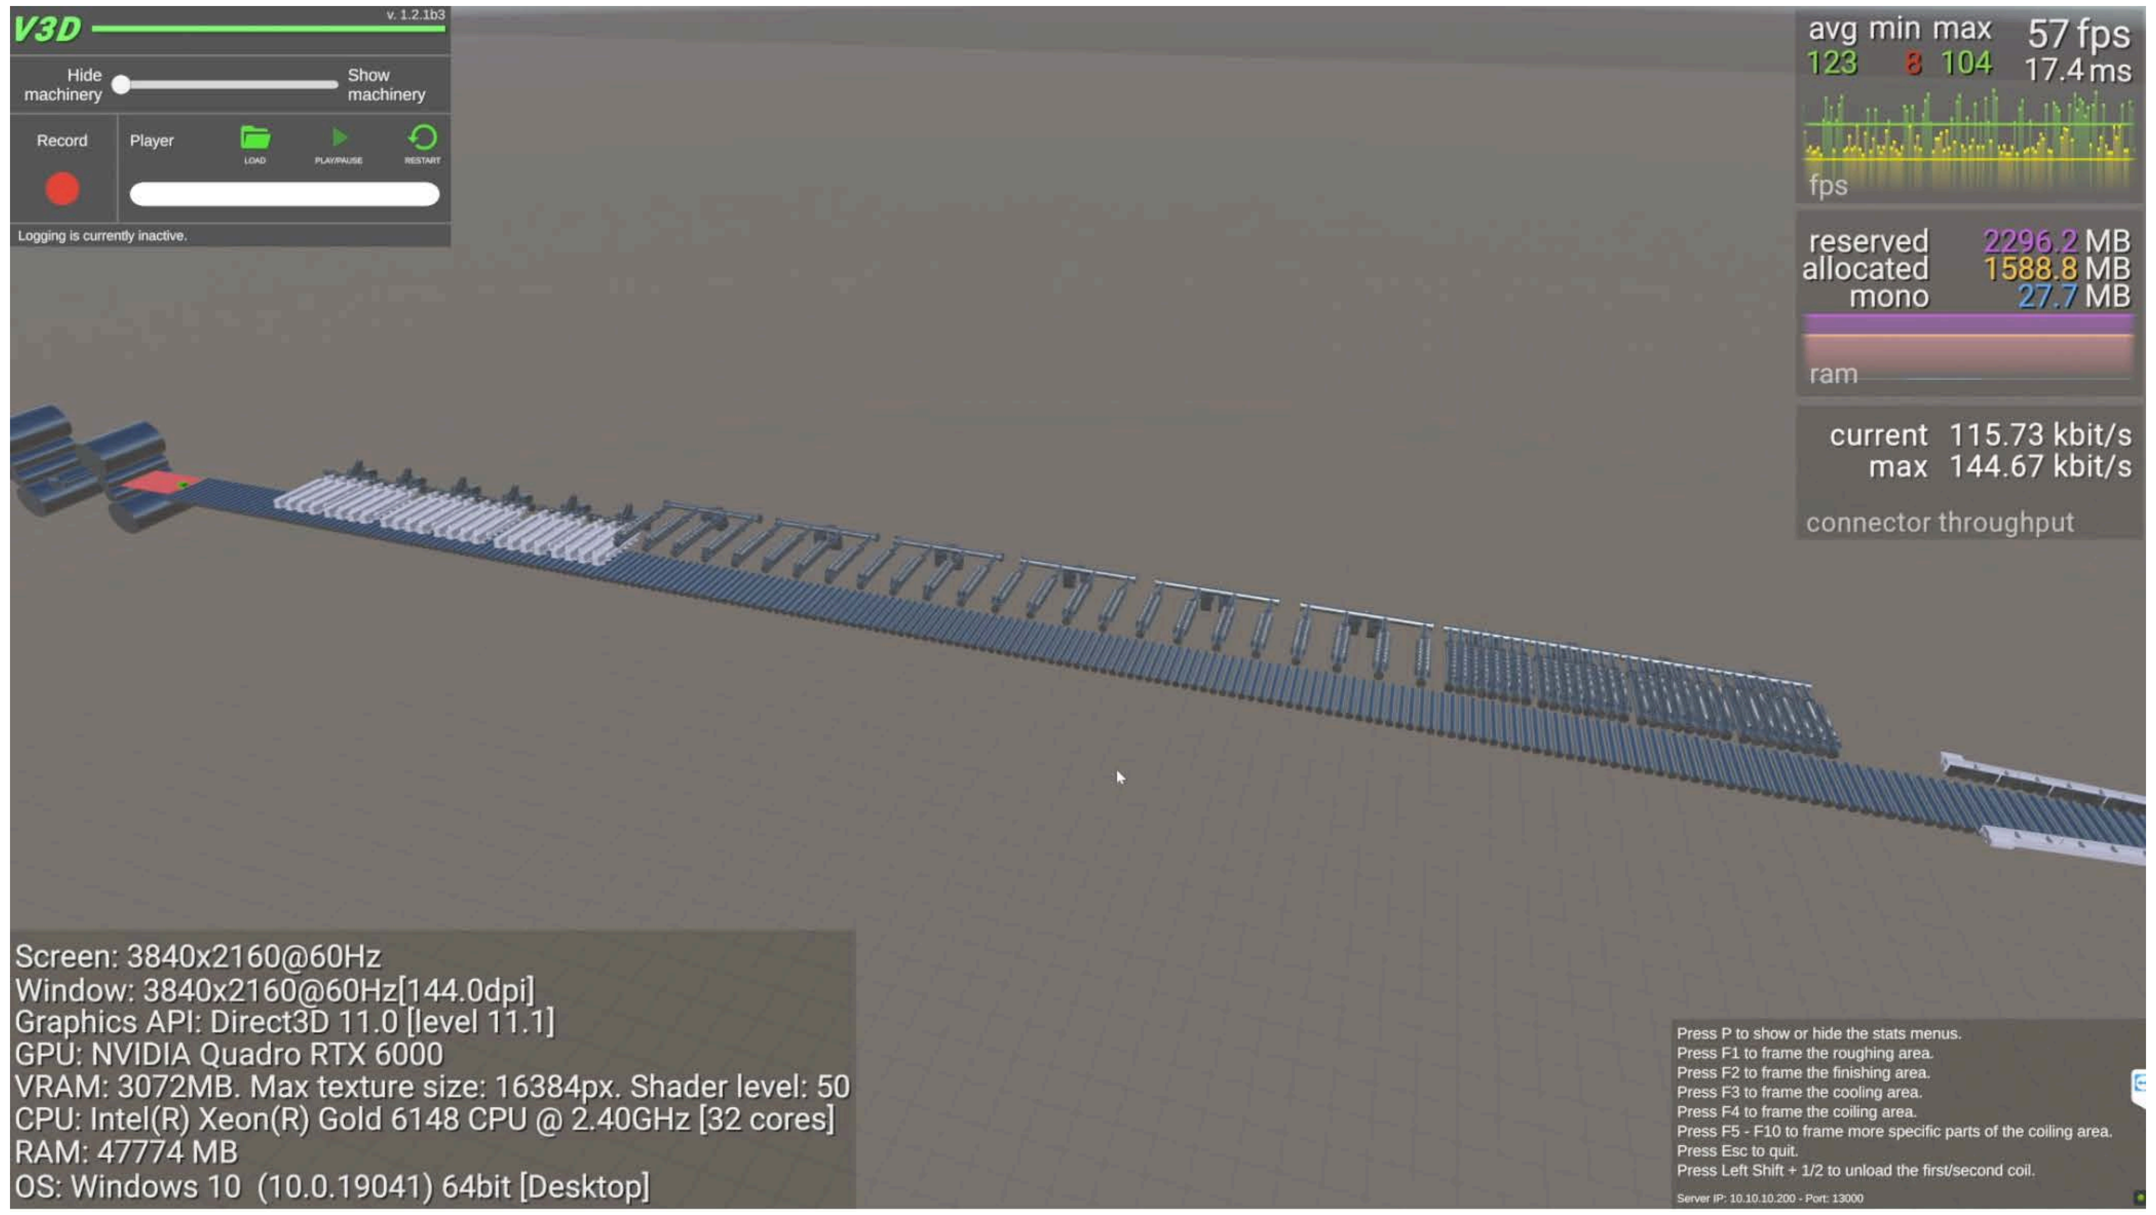

PROBLEM FORMULATION AND REQUIREMENTS

Basic Problem

- N decision makers (agents) connected through a communication network
- Agent i has a number M_i *unknown a priori to the other agents*
- Agents want to compute *iteratively* and in a *decentralised way*

$$M^\star := \min_i M_i$$

Requirements (hence, Objectives):

- **Robustness** to changing environments: we are interested in the case where numbers M_i and the communication topology may exhibit some changes at runtime
  - Calls for **global asymptotic stability**
- **Decentralisation**: the update laws uses only **local** information and does **not** depend from centralised quantities
- **Scalability**: the number of variables stored by each agent does not depend on the **size** of the network an its **topology**
- **Uniformity**: the convergence rate must not change (vanish)

"ALGORITHMIC" APPROACHES

The Max-Consensus Algorithm (1)

- Agent i stores a local estimate x_i of $M^\star$

  and updates it as

$$x_i^{t+1} = \min_{j \in [i]} x_j^t$$

  where $[i]$ denotes the set of neighbours of i in the communication topology

- If initialised as

$$x_i^{t_0} = M_i$$

  the algorithm converges in *finite time*. Namely:

$$\exists\, T > 0 \text{ such that } x_i^t = M^\star \,, \forall\, t \geq T$$

The Max-Consensus Algorithm (2)

The algorithm suffers from a significant **drawback** considering our application-driven requirements:

- Suppose that at some time $t_1 > T$ all M_i suddenly change to some value $M_i^{\text{new}} > M^\star$
- Then:
  - ✦ $M^\star$ changes to $M^\star_{\text{new}} = \min_i M_i^{\text{new}}$
  - ✦ at time t_1 one has $x_i^{t_1} \neq M_i^{\text{new}}\ ,\ \forall\, i$
  - ✦ shifting the time, we can see the problem as a new run with $t_0^{\text{new}} = t_1$ in which $x_i^{t_0} = M_i$ does not hold
  - ✦ it can be shown that $x_i^t = x_i^T\ ,\ \forall\, t \geq T$ and hence $x_i^t \not\longrightarrow M^\star_{\text{new}}$
  - ✦ Namely: Max-Consensus **fails to track the new minimum**

The Max-Consensus Algorithm (3)

Max-Consensus fails because **convergence is not global** in the initial conditions:

- From on observer theoretic point of view:
  - In the first run, the Max-Consensus converges to a subset of the consensus manifold $\mathbb{C}$:

$$x^t = (x_1^t, \ldots, x_N^t) \longrightarrow (M^\star, \ldots, M^\star) = M^\star \cdot (1, \ldots, 1) \in \mathbb{C} = \mathrm{span}(1, \ldots, 1)$$

  - But, every point of $\mathbb{C}$ is a fixed point of the update law since

$$c = \min(c, \ldots, c)\,, \ \forall\, c$$

  - Thus, if $x^t \in \mathbb{C}$ and $M^\star$ changes, its new value (the new desired steady-state) is **not observable**!
- Other ``algorithmic" approaches, such as *FloodMax*, *Yo Yo*, *MegaMerger*, have a similar drawback

OPTIMISATION APPROACHES

$M^\star := \min_i M_i$ is the unique solution to every constrained optimisation problem of the form:

$$\max \sum_i \psi_i(x_i) \quad \text{subject to} \quad x_i \leq M_i \text{ and } x_i = x_j \,,\ \forall\, i,j$$

where $\psi_i(x_i)$ are continuous and strictly increasing functions

We can borrow approaches from **distributed optimisation**:

- Consensus-based (sub)gradient and second order methods
  - ✦ Global convergence, unconstrained or vanishing step-size (**non uniform convergence rate**)
- Projected and primal-dual methods with inequality constraints
  - ✦ Handle constrained problems, require initialisation (**non-global convergence**)
- Gradient-Tracking methods
  - ✦ Uniform convergence, unconstrained problems + require initialisation (**non-global convergence**)
- \`\`Node-based" ADMM approaches
  - ✦ Uniform convergence, require initialisation (**non-global convergence**)
- \`\`Edge-based" ADMM approaches
  - ✦ Uniform and Global convergence, **not scalable**

Every agent must store and update a certain number of variables for every different connection with other agents

A problem for **scalability** (number of variable not constant in the size of the net) and for **robustness** (what happens if the connection topology changes?)

THE ``BLANKET" IS SHORT!

- ``Algorithmic" Approaches:
  - scalable but not robust to changes
- Optimisation approaches:
  - scalable but not robust to changes
  - scalable but not uniform
  - robust, uniform, but not scalable

\+ optimisation approaches typically use centralised parameters and so they do not have good decentralisation properties

Question: can we have **all** properties?

Yes (Our Answer), provided that we compromise on exact convergence.

Our design yields:

- Approximate (to an arbitrary precision), uniform and global convergence + scalability + stability

OR:

- Exact, global but non-uniform convergence + scalability + stability with good decentralisation properties

Stability is never proved in other approaches but is crucial as it ensures that small variations lead to small consequences

STANDING ASSUMPTIONS

- **Connectedness**: the communication network is $I^\star$ - connected. That is, every node is connected to at least one node in the set:

$$I^\star := \arg\min_i M_i = \text{set of all agents with } M_i = M^\star$$

- **Consistency**: every agent knows a lower bound $\mu_i > 0$ for $M^\star$

Since $M^\star$ is a centralised quantity, the latter assumption undermines decentralisation!

- Fortunately, in applications it is usually not a problem to know such a μ_i
- Moreover, agents need not know other agents' μ_i
- Furthermore, the assumption is not necessary if one seeks ε-approximate convergence. Indeed, in this case one can simply put $\mu_i = \varepsilon$ regardless the value of $M^\star$

THE UPDATE LAWS

Each agent stores an estimate x_i of $M^\star$ and changes it according to the following update law:

$$x_i^+ = \Pi_{[\mu_i, M_i]}\left[e^{h_i^t} x_i + k_i \sum_{j\in[i]} (x_j - x_i)\right]$$

where:

- $\Pi_{[a,b]}(\cdot) := \min\{b, \max\{a, \cdot\}\}$
- $k_i \in \left(0, \frac{1}{\text{card}([i]) - 1}\right]$
- h_i is a signal to be designed later on

- Agents can choose the gain k_i independently to one another and the same holds for h_i
  - good **decentralisation**
  - no need of having a column/row stochastic matrix
- Only one variable is stored (like the Max-Consensus) independently from the network size or topology
  - good **scalability**

CHOICE OF h_i

The signals h_i must be chosen to be "exciting enough" → two options:

1. **Sufficiency of excitation**:

Definition 5 (Sufficiency of Excitation). *With* $t_0 \in \mathbb{N}$*, the family* $(h_i)_{i \in \mathcal{N}}$*, is said to be* sufficiently exciting from t_0 *if there exist* $\underline{h}(t_0) > 0$ *and* $\Delta(t_0) \in \mathbb{N}_{\geq 1}$ *such that, for each* $m \in \mathbb{N}_{\geq 1}$ *satisfying*

$$m \leq \frac{1}{\underline{h}(t_0)} \log\left(\frac{\mathrm{M}^\star}{\underline{\mu}}\right) \qquad (15)$$

and each $i \in \mathcal{N}$*, there exists at least one* $s_i \in \{t_0 + 1 + (m-1)\Delta(t_0), \ldots, t_0 + m\Delta(t_0)\}$ *such that* $h_i^{s_i} \geq \underline{h}(t_0)$.

→ $\underline{\mu} := \min_i \mu_i \implies h_i^t \to 0$

This enables exact convergence but ruins uniformity

In qualitative terms, given an initial time t_0, sufficiency of excitation implies that the signals h_i^t are positive "frequently enough" for a "large enough" amount of time succeeding t_0. When $(h_i)_{i\in\mathbb{N}}$ is sufficiently exciting from every t_0, and independently on it, then we say that $(h_i)_{i\in\mathbb{N}}$ enjoys the **uniformity** of excitation property.

2. **Uniformity of excitation**:

Definition 6 (Uniformity of Excitation). *The family* $(h_i)_{i \in \mathcal{N}}$ *is said to be* uniformly exciting *if it is sufficiently exciting from every* t_0*, with* $\underline{h}$ *and* Δ *not dependent on* t_0.

→ Stronger condition, does **not** allow

$$h_i^t \to 0$$

ACHIEVING SUFFICIENCY OF EXCITATION

- The definition of sufficiency of excitation involves centralised quantities $\underline{\mu}, M^\star, \Delta, \underline{h}$
- However, we can achieve it in a fully decentralised way!

Lemma: Suppose that, for each i , the signals satisfy:

$$\sum_{t\in\mathbb{N}} h_i^t = \infty$$

Then $(h_i)_{i\in\mathbb{N}}$ is sufficiently exciting.

If, in addition

$$\sum_{t\in\mathbb{N}} (h_i^t)^2 < \infty$$

Then: $h_i^t \to 0$

Definition 5 (Sufficiency of Excitation). *With $t_0 \in \mathbb{N}$, the family $(h_i)_{i\in\mathcal{N}}$, is said to be* sufficiently exciting *from t_0 if there exist $\underline{h}(t_0) > 0$ and $\Delta(t_0) \in \mathbb{N}_{\geq 1}$ such that, for each $m \in \mathbb{N}_{\geq 1}$ satisfying*

$$m \leq \frac{1}{\underline{h}(t_0)} \log\left(\frac{\mathrm{M}^\star}{\underline{\mu}}\right) \qquad (15)$$

and each $i \in \mathcal{N}$, there exists at least one $s_i \in \{t_0 + 1 + (m-1)\Delta(t_0), \ldots, t_0 + m\Delta(t_0)\}$ such that $h_i^{s_i} \geq \underline{h}(t_0)$.

- Classical *Robbins-Monro*'s vanishing step-size conditions for stochastic gradient descent.
- In our setting, this just gives a decentralised choice to achieve **both** sufficiency of excitation and vanishing signals

ACHIEVING UNIFORMITY OF EXCITATION

Definition 6 (Uniformity of Excitation). *The family $(h_i)_{i\in\mathcal{N}}$ is said to be* uniformly exciting *if it is sufficiently exciting from every t_0, with $\underline{h}$ and Δ not dependent on t_0.*

- The definition of uniformity of excitation involves centralised quantities as well
- However, we can achieve it in a fully decentralised way too:

Lemma 1. *Suppose that, for each $i \in \mathcal{N}$, h_i is periodic and there exists $t \in \mathbb{N}$ for which $h_i^t > 0$. Then, the family $(h_i)_{i\in\mathcal{N}}$ is uniformly exciting.*

- The periods can be different and chosen independently → every agent can choose an arbitrary **periodic** signal
- This connects to typical **dithering signals** in extremum seeking and output feedback in presence of unobservable steady states

MAIN CONVERGENCE RESULT

Theorem 1. *Under Assumptions 1 and 2, consider the update laws* (13), *in which* k_i *satisfies* (14). *Suppose that, for a given* $t_0 \in \mathbb{N}$, *the family* $(h_i)_{i \in \mathcal{N}}$ *is sufficiently exciting from* t_0 *in the sense of Definition 5. Then, the following claims hold:*

1. *There exists* $t^\star = t^\star(t_0)$ *such that every solution* x *to* (9) *starting at* t_0 *satisfies*

$$
\begin{aligned}
x_i^t &\geq \mathrm{M}^\star, & \forall t \geq t^\star(t_0),\ \forall i \in \mathcal{N} \setminus I^\star \\
x_i^t &= \mathrm{M}^\star, & \forall t \geq t^\star(t_0),\ \forall i \in I^\star,
\end{aligned}
$$

with $I^\star$ *given by* (11).

If **sufficiency of excitation** holds, in finite time agents with $M_i = M^\star$ reach the equilibrium $M^\star$ and stay there forever.

MAIN CONVERGENCE RESULT

Theorem 1. *Under Assumptions 1 and 2, consider the update laws (13), in which k_i satisfies (14). Suppose that, for a given $t_0 \in \mathbb{N}$, the family $(h_i)_{i\in\mathcal{N}}$ is sufficiently exciting from t_0 in the sense of Definition 5. Then, the following claims hold:*

2. *For each $\epsilon > 0$, there exists $\delta(\epsilon) > 0$ such that, if*

$$\limsup_{t\to\infty} h_i^t \le \delta(\epsilon), \quad \forall i \in \mathcal{N}, \tag{16}$$

then each solution x starting at t_0 satisfies

$$\lim_{t\to\infty} |x_i^t - \mathrm{M}^\star| \le \epsilon, \quad \forall i \in \mathcal{N}. \tag{17}$$

In particular, the set

$$\mathcal{A}_\epsilon := \prod_{i\in\mathcal{N}} [\mathrm{M}^\star, \min\{\mathrm{M}^\star + \epsilon, \mathrm{M}_i\}]$$

is globally attractive for (9).

3. *If the family $(h_i)_{i\in\mathcal{N}}$ is uniformly exciting in the sense of Definition 6, then $\mathcal{A}_\epsilon$ is globally uniformly attractive.*

For every arbitrary precision ε , one can choose h_i to guarantee

$$x_i^t \to M^\star + O(\varepsilon)$$

globally for all agents

If, in addition, $(h_i)_{i\in\mathbb{N}}$ is **uniformly exciting**, then convergence is **uniform**!

MAIN CONVERGENCE RESULT

Theorem 1. *Under Assumptions 1 and 2, consider the update laws (13), in which k_i satisfies (14). Suppose that, for a given $t_0 \in \mathbb{N}$, the family $(h_i)_{i\in\mathcal{N}}$ is sufficiently exciting from t_0 in the sense of Definition 5. Then, the following claims hold:*

4. *If all the signals h_i are non-zero and periodic (with possibly different periods), then there exists a compact set $\mathcal{A}^u_\epsilon \subset \mathcal{A}_\epsilon$ which is globally uniformly attractive and stable, hence, globally uniformly asymptotically stable.*
5. *If*

$$\lim_{t\to\infty} h_i^t = 0, \quad \forall i \in \mathcal{N}$$

then, the set $\mathcal{A}$, given by (10), is globally attractive for (9), i.e.

$$\lim_{t\to\infty} x_i^t = \mathrm{M}^\star, \quad \forall i \in \mathcal{N}.$$

If $(h_i)_{i\in\mathbb{N}}$ are non-zero and periodic, then there exists a compact attractor $\subset M^\star + O(\varepsilon)$ which is also stable (hence **uniformly globally asymptotically stable**)

If $(h_i)_{i\in\mathbb{N}}$ are only **sufficiently exciting** and

$$h_i^t \to 0$$

then **exact convergence** holds. That is:

$$x_i^t \to M^\star \,,\ \forall i$$

However, in this case **convergence is not uniform**

NUMERICAL RESULTS

Example 1: changing conditions with uniform convergence

- **Topology (a)**
  - Agents 1,2,3,4
  - $(M_1,\ M_2,\ M_3,\ M_4) = (10,\ 12,\ 13,\ 13),\ M^\star = M_1 = 10$
  - $h_i\ :\ \text{square wave}$
- **Topology (b)**
  - Agents 1,2,3,4 + 5,6
  - $(M_5,\ M_6) = (7,\ 11)$
  - $(M_1,\ M_3) = (11,\ 13),\ M^\star = M_5 = 7$
- **Topology (c)**
  - Agents 2,3 leave the network
  - $(M_1, M_4, M_5, M_6) = (12, 16, 11, 16)\ ,\ M^\star = M_5 = 11$
- **Topology (c)**
  - $M_4 = 8,\ M^\star = M_4 = 8$

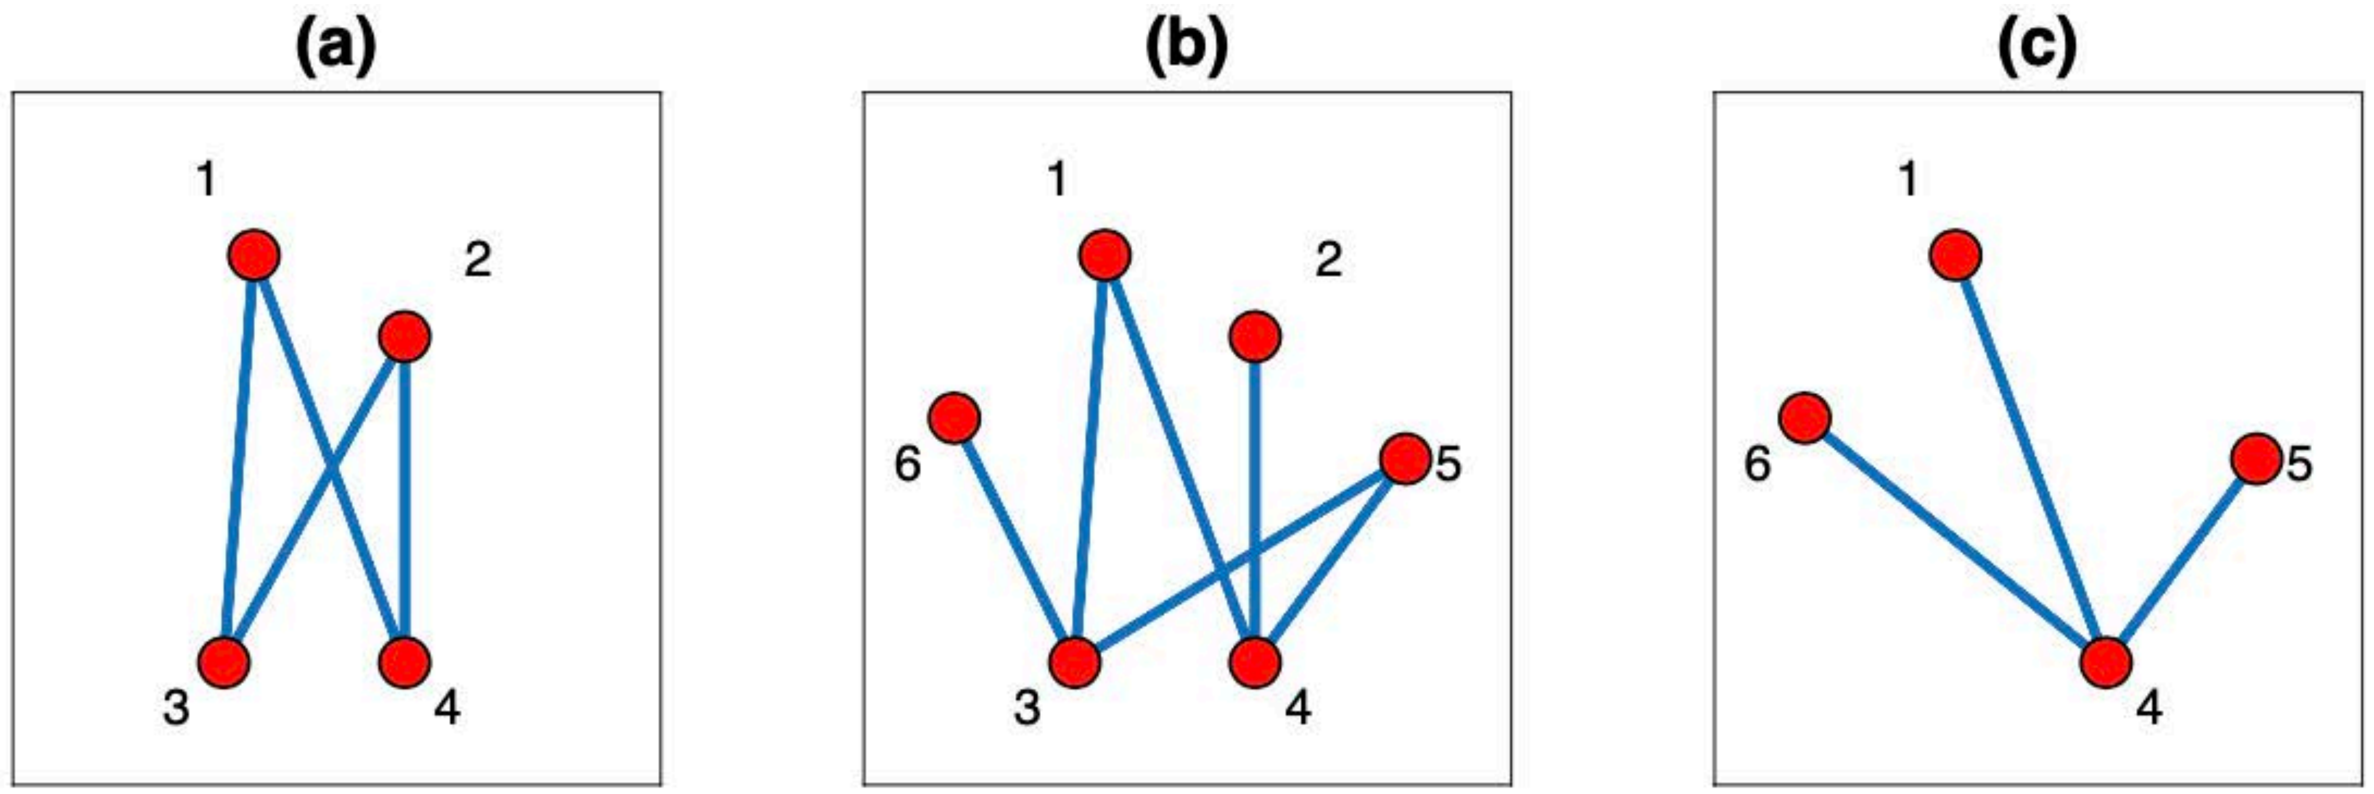

NUMERICAL RESULTS

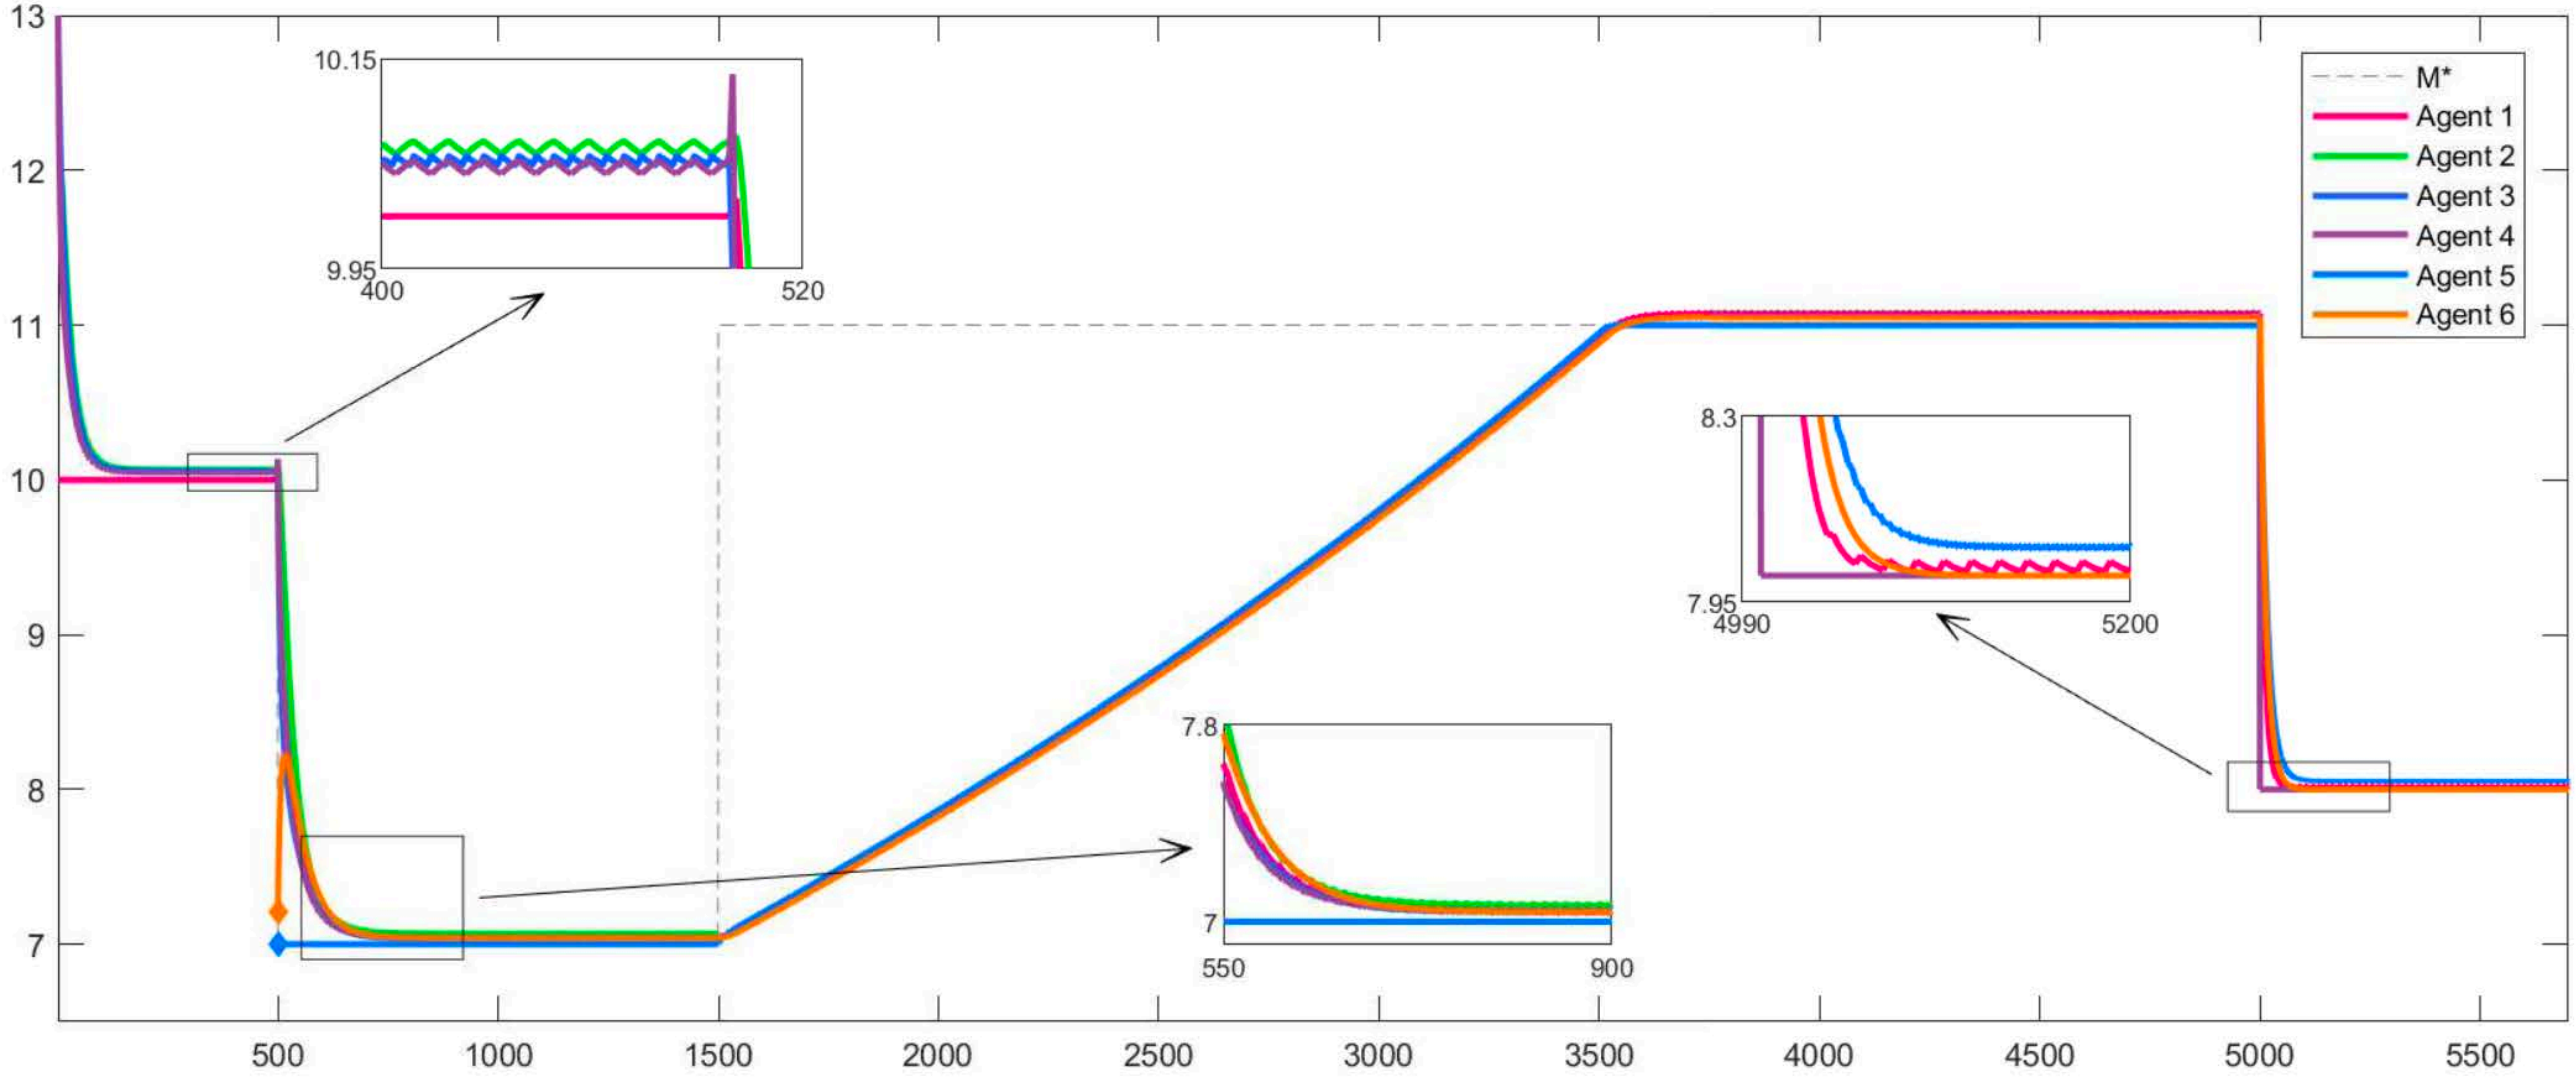

NUMERICAL RESULTS

Max-Consensus fails!

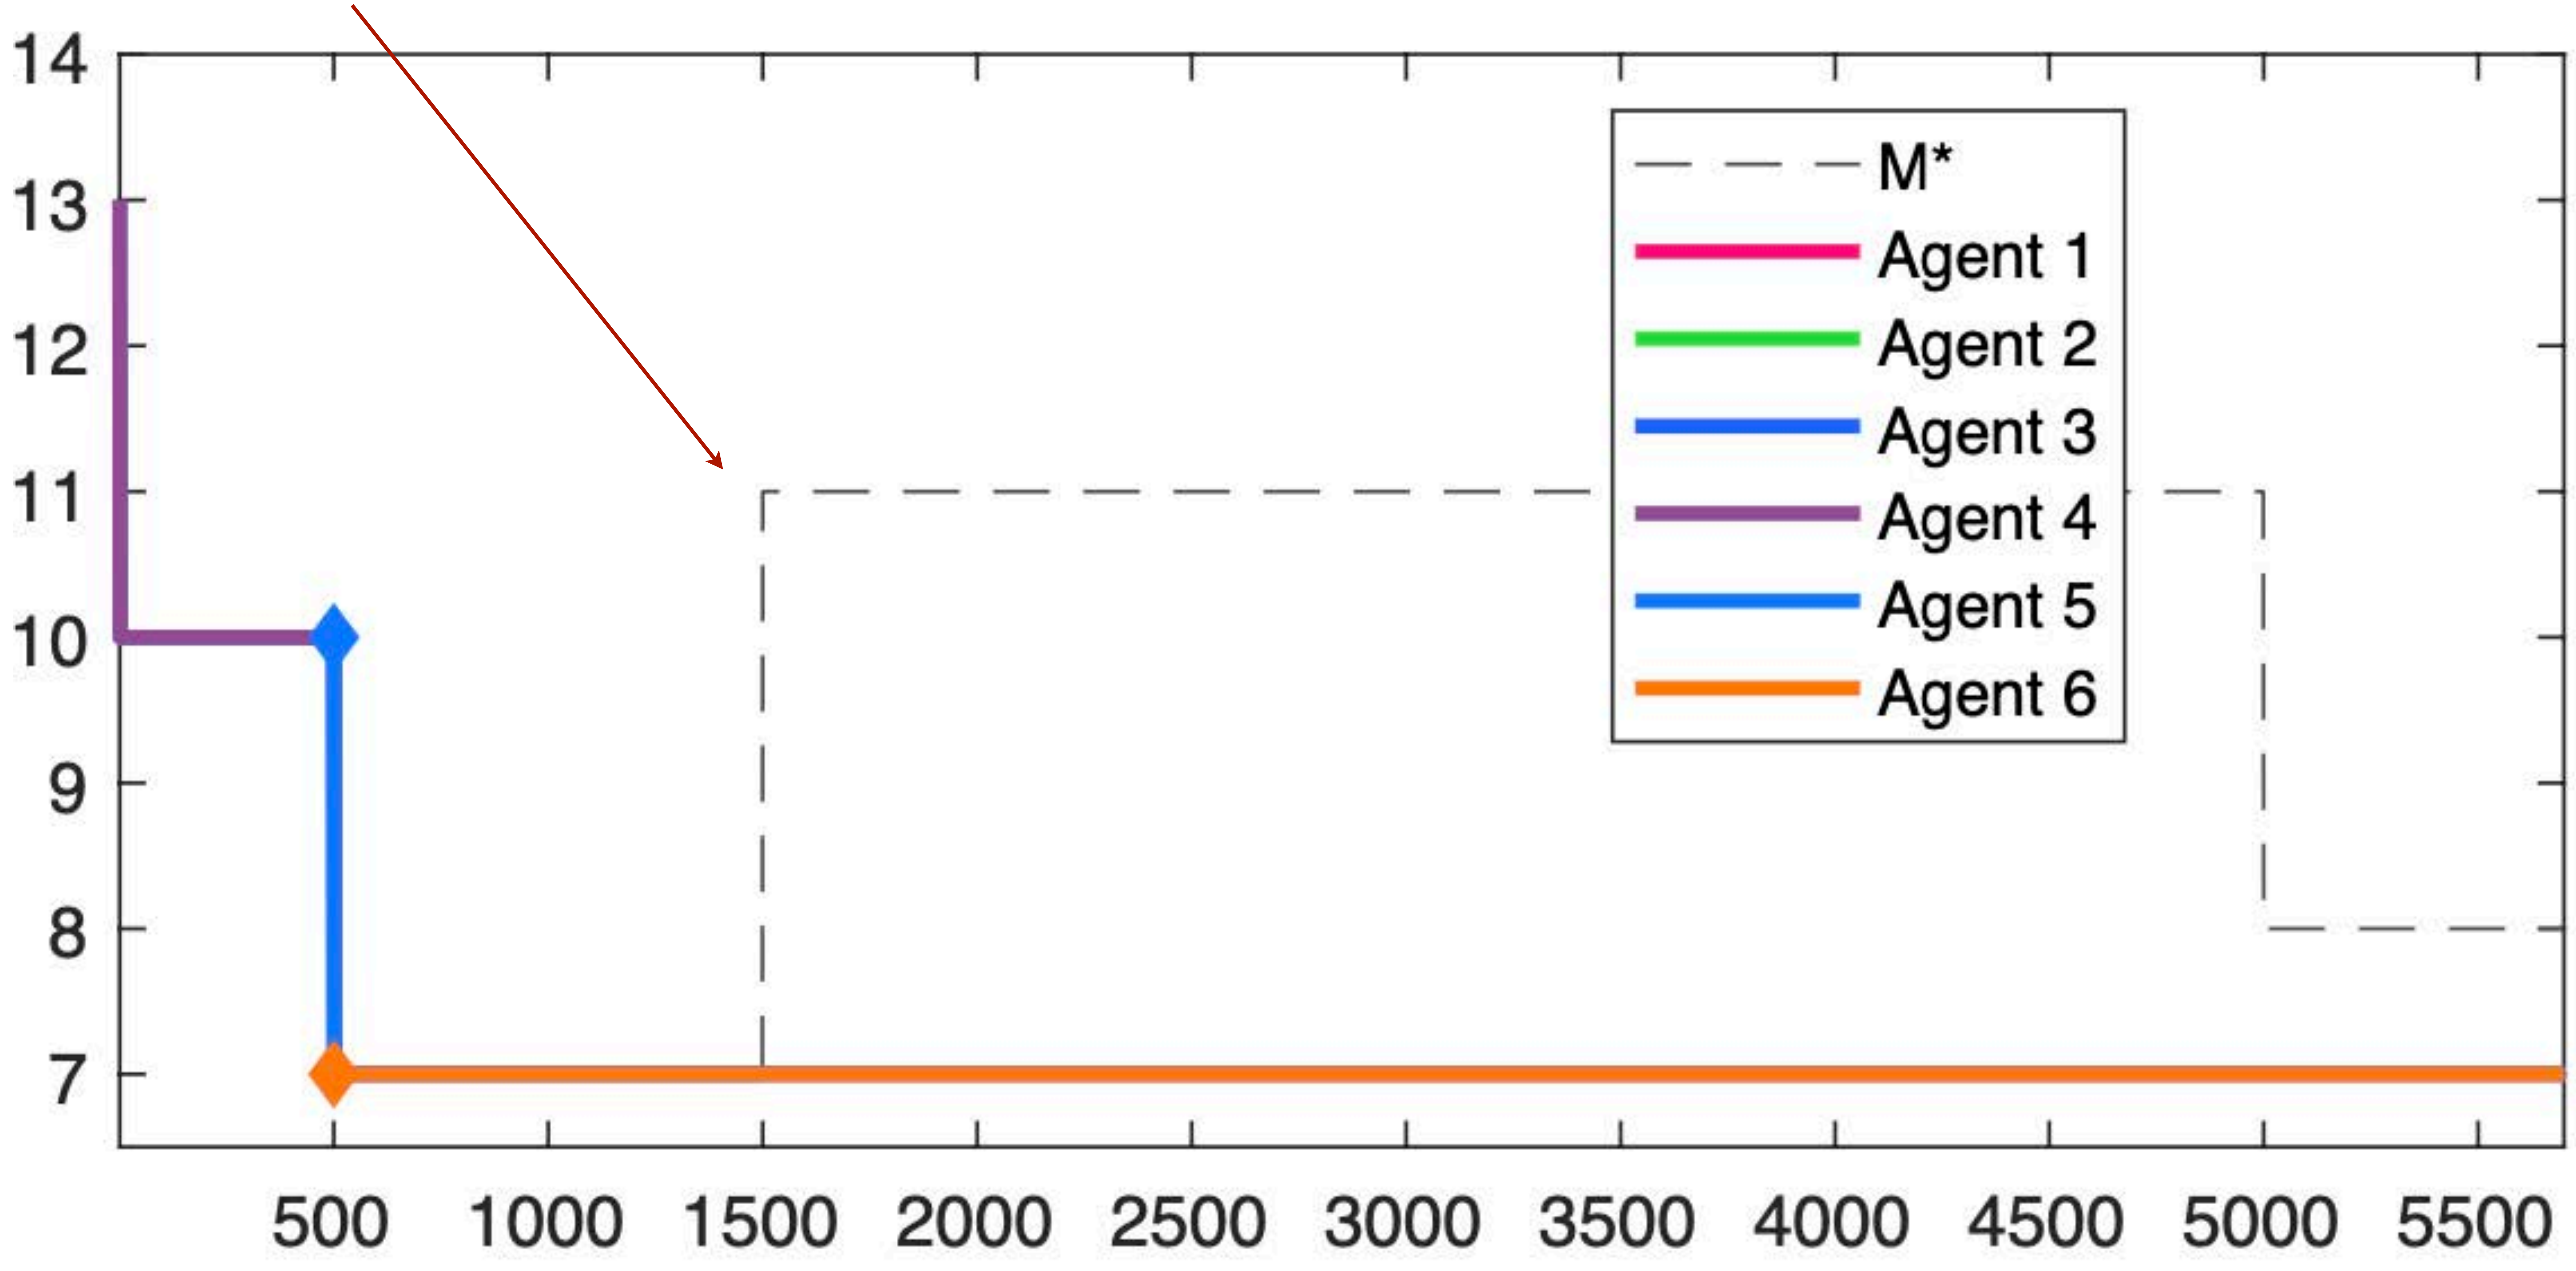

NUMERICAL RESULTS

Example 2: comparisons of scenarios with different $(h_i)_{i\in\mathbb{N}}$ and non-uniform convergence

- **Topology (a)**
  - Agents 1,2,3,4
  - $(M_1, M_2, M_3, M_4) = (3, 6, 9, 15)\,,\ M^\star = M_1 = 3$
  - $h_i\ :\ \text{square wave}$
- **Topology (a)**
  - $M_1 = 15\,,\ M^\star = M_2 = 6$
- **Topology (a)**
  - $M_2 = 15\,,\ M^\star = M_3 = 9$
- **Topology (a)**
  - $M_3 = 12\,,\ M^\star = M_3 = 12$
- **Topology (a)**
  - $M_3 = 15\,,\ M^\star = M_1 = M_2 = M_3 = M_4 = 15$

(a)

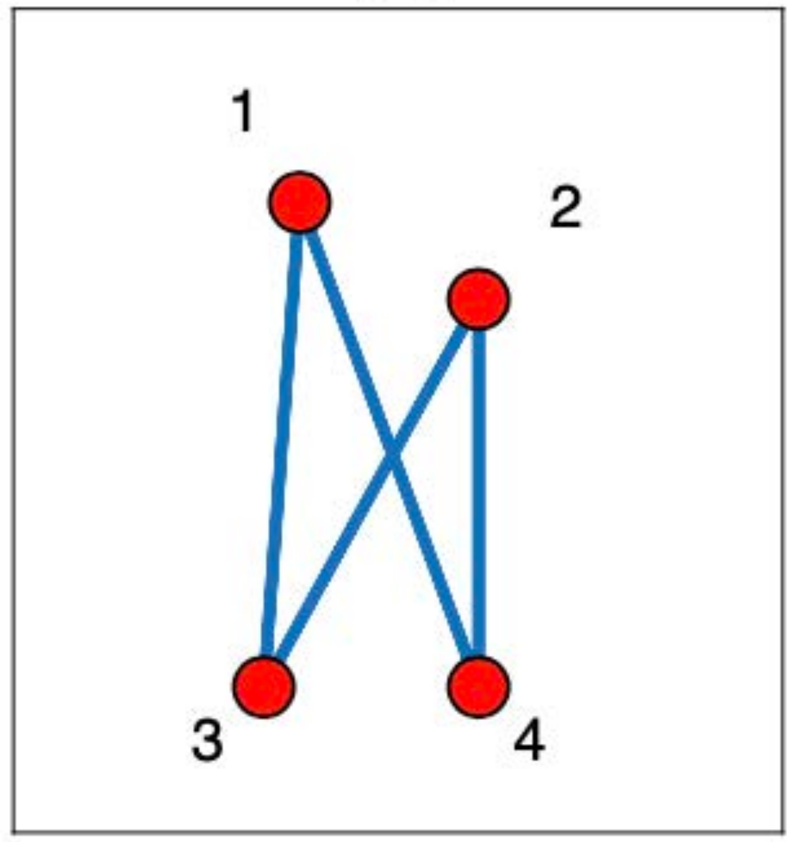

NUMERICAL RESULTS

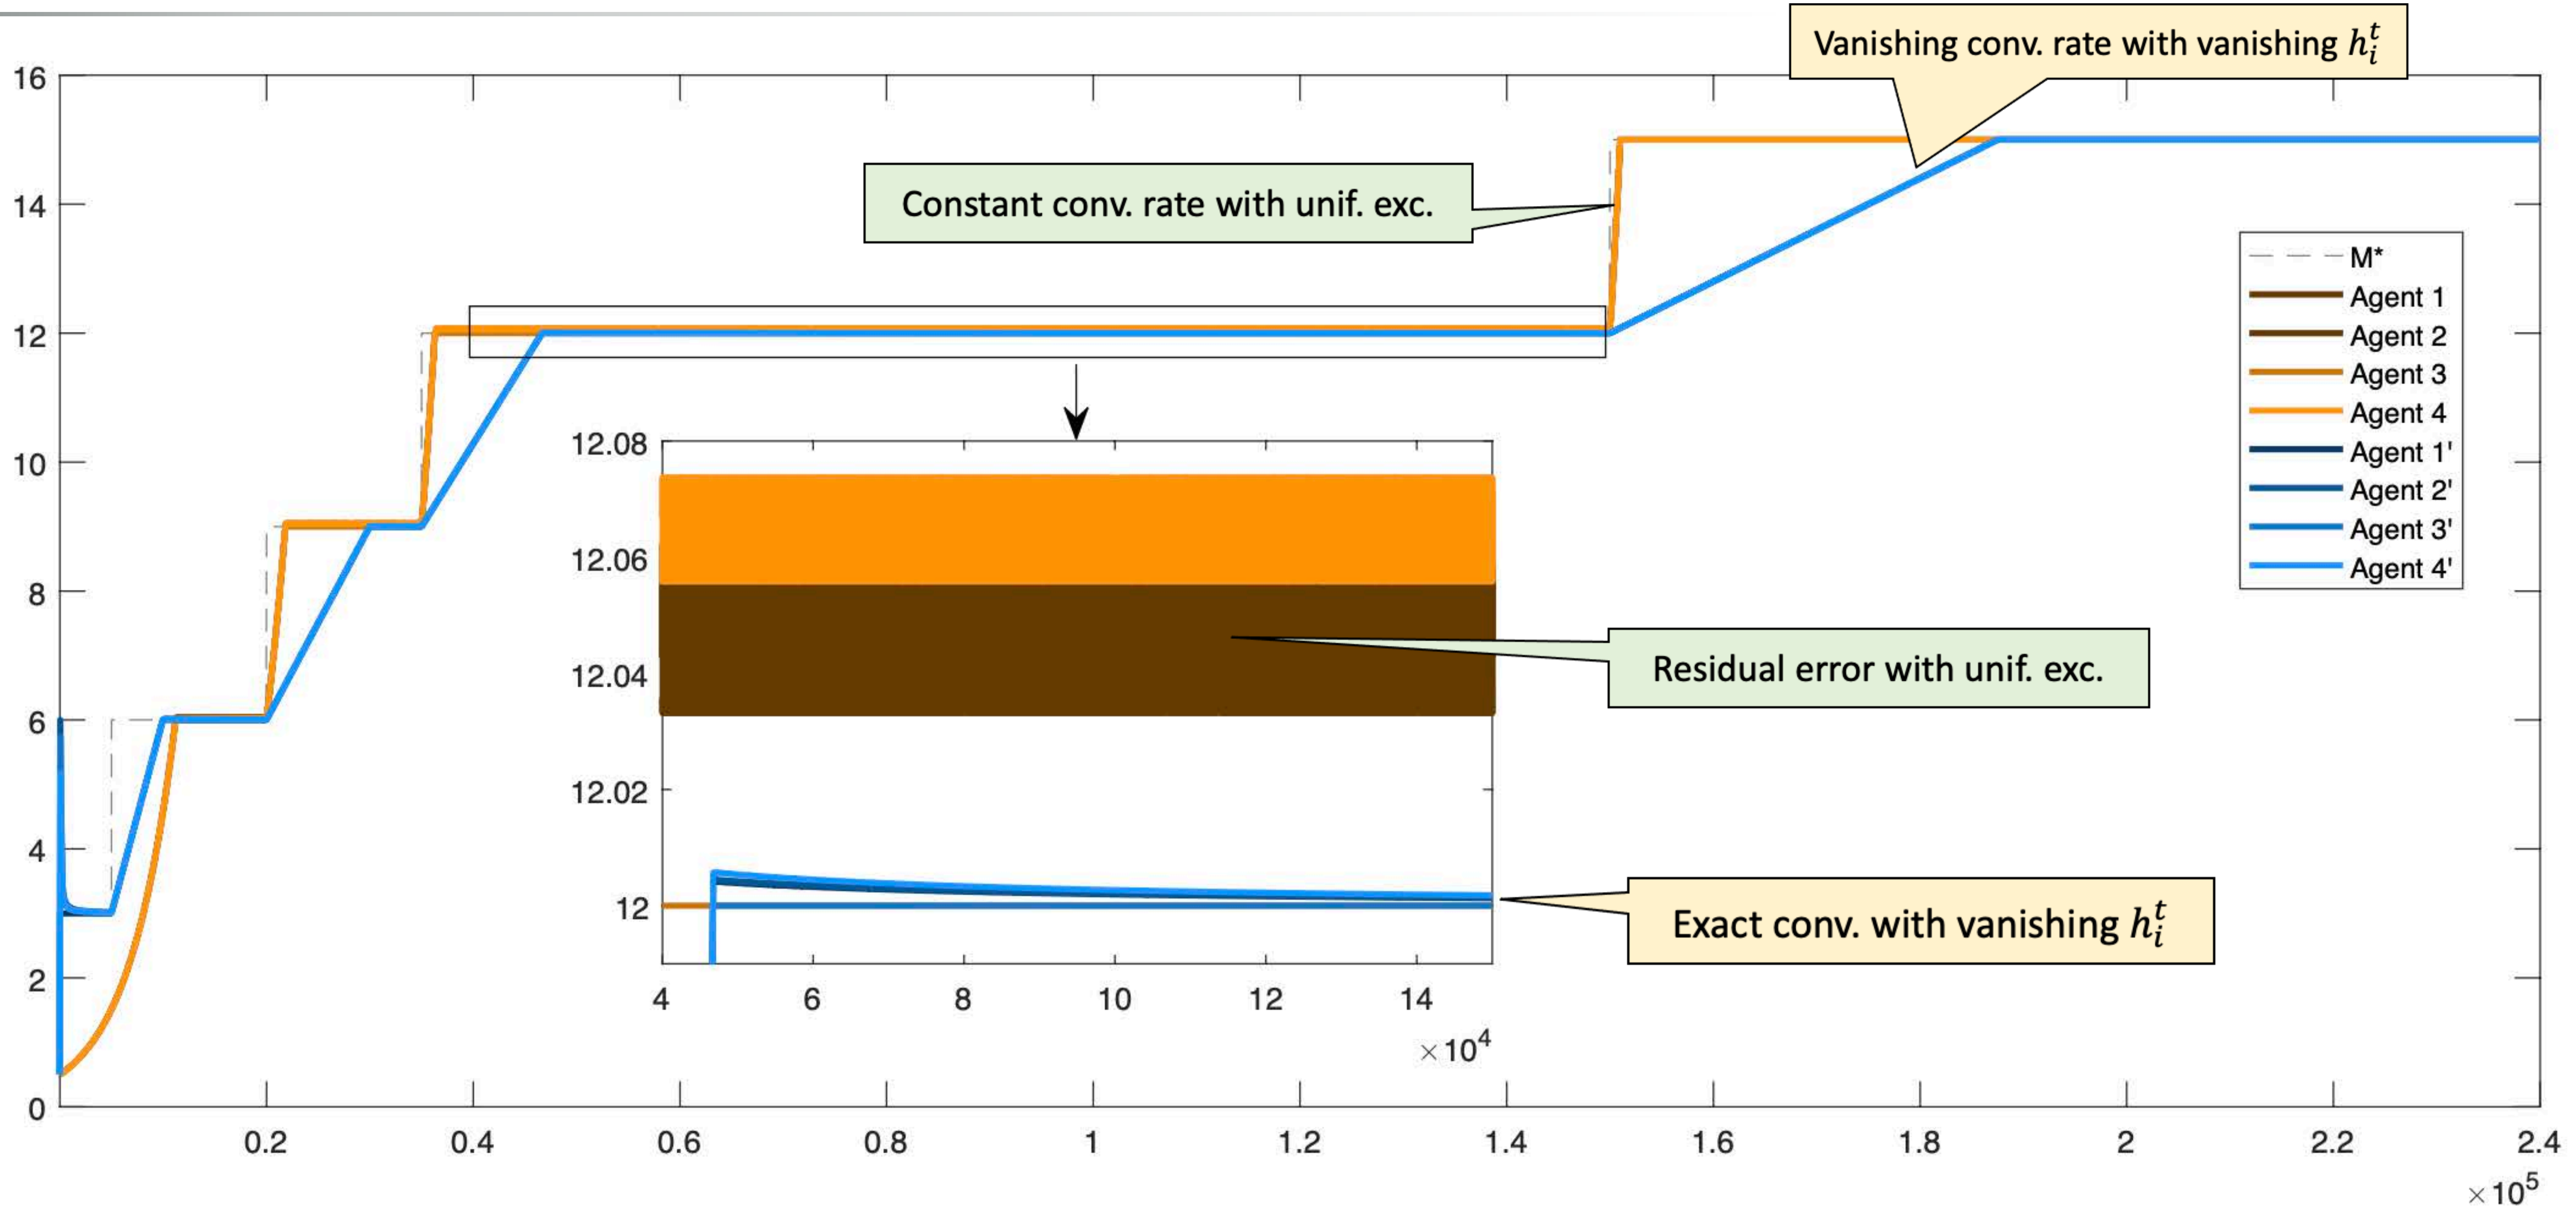

NUMERICAL RESULTS

Max-Consensus fails!

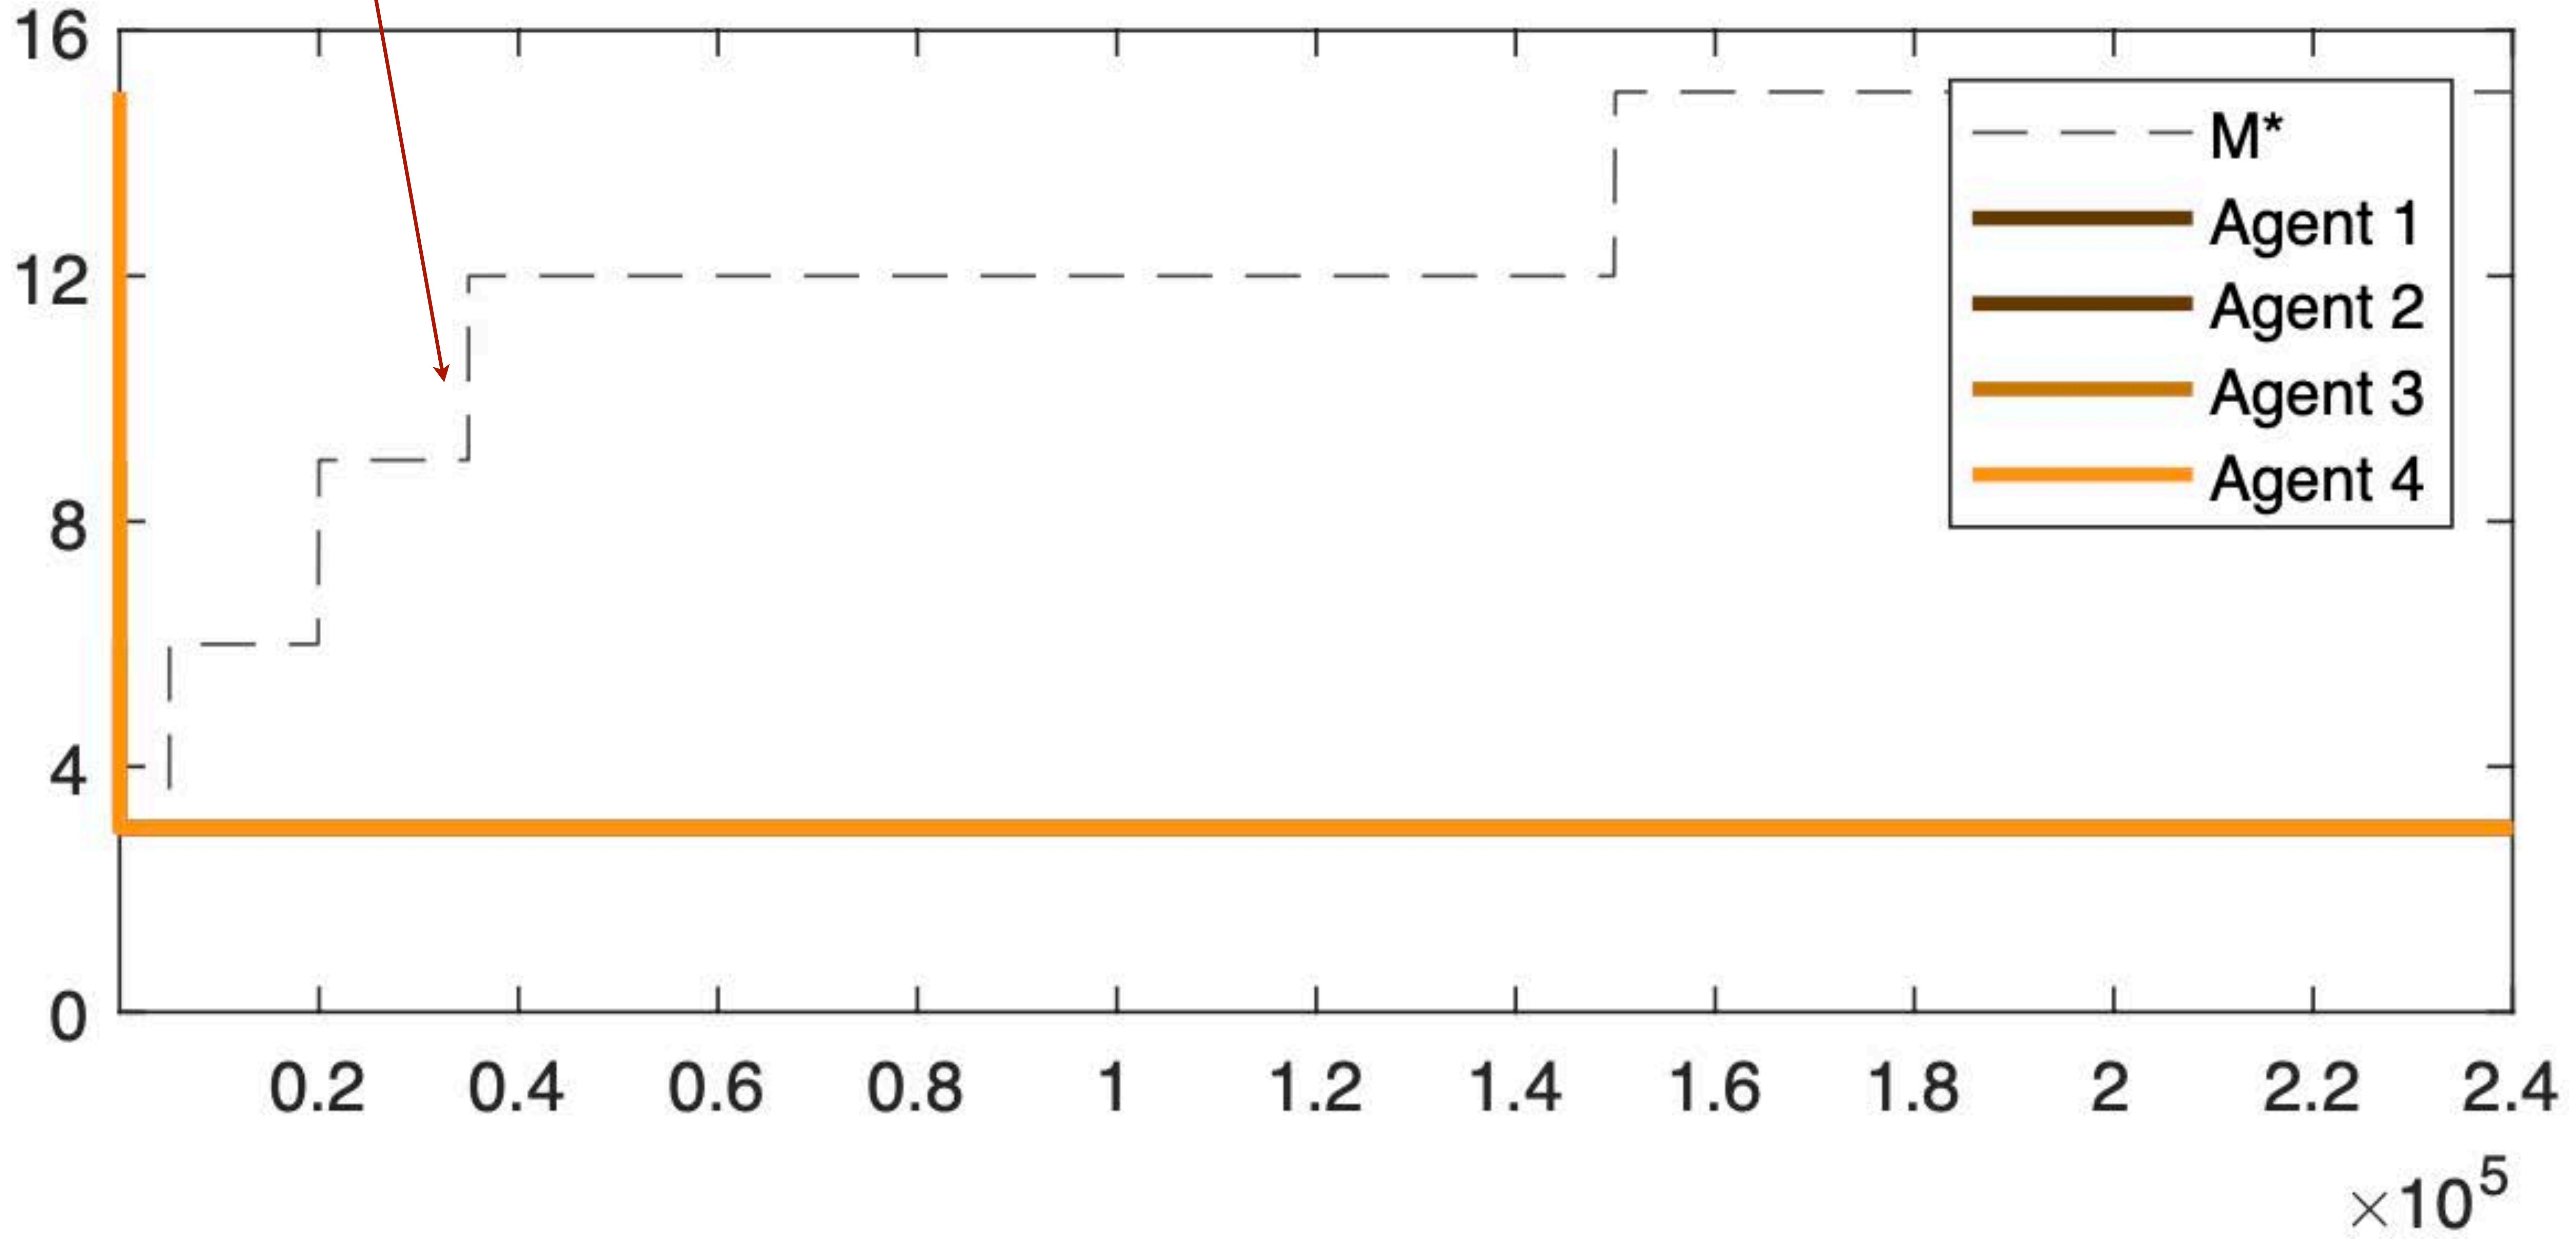

DISTRIBUTED MINIMUM SHARING: TAKEAWAYS

- Ideally, we'd like solutions ensuring altogether:
  - Decentralisation
  - Scalability
  - Exact convergence
  - Uniform convergence
  - Global convergence
  - Stability
- The **blanket is short** and we need to **compromise**
- Our solution enjoys ``almost all" properties. The compromise is between:
  - Approximate but uniform convergence
  - Exact but non-uniform convergence
- This compromise goes a little deeper:
  - "faster" convergence $\Longleftrightarrow$ "larger" steady state errors
- Adaptive selection of $(h_i)_{i \in \mathbb{N}}$ can reduce significantly the impact of this compromise

A SLIGHTLY BIGGER PICTURE

``Dithering" signals are widespread in control and always lead to a compromises:

- Extremum seeking: dithering allows to estimate the descent direction
- Control toward sets where observability is lost: dithering allows to keep uniform observability
- Regularisation in least-squares problem: can be seen as dithering, it makes problems well-conditioned
- Watermarking in cybersecurity: avoiding replay attacks
- ...

Thank you!