Collaborative System Identification via Parameter Consensus

Ivan Papusha

Control and Dynamical Systems, Caltech

American Control Conference
June 4, 2014

“post”-modern control

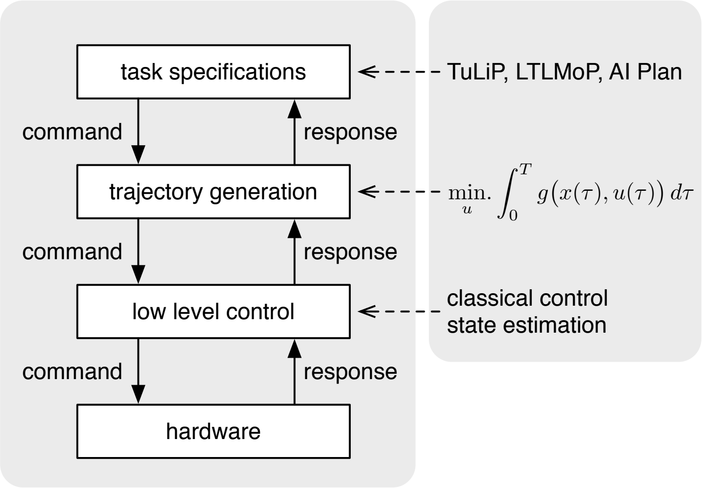

optimal control + supervisory temporal logic

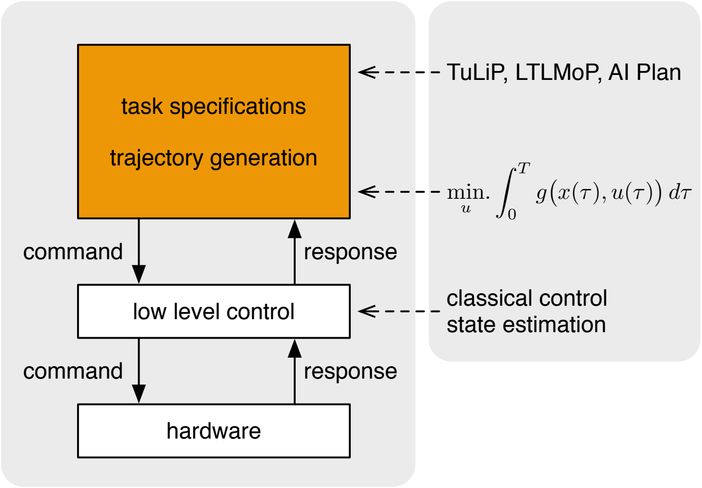

optimal control + adaptation

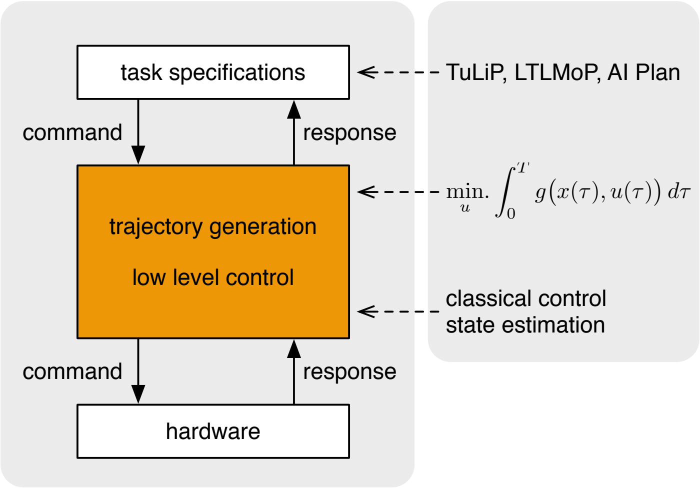

optimal control + adaptation + multiagent

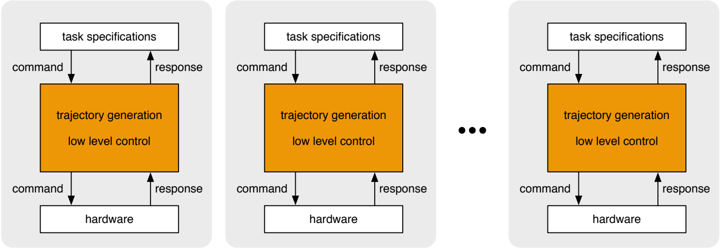

optimal control + adaptation + multiagent + networking

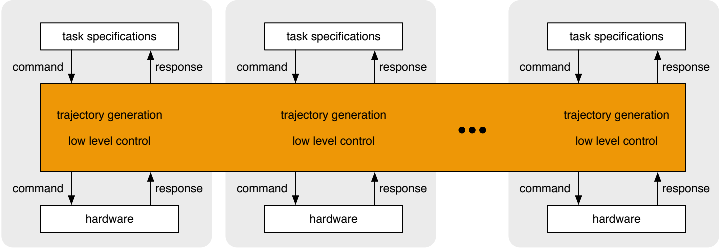

applications of networked adaptive systems

- smartgrid: bootstrapping, disturbance rejection
- circuits: high performance phase locked loops
- robotics: distributed bootstrapping with consensus constraints
- **today**: collaborative system identification

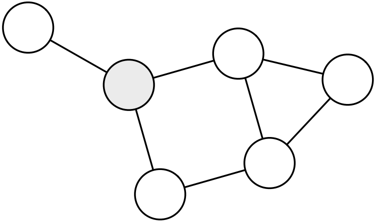

applications of networked adaptive systems

- smartgrid: bootstrapping, disturbance rejection
- circuits: high performance phase locked loops
- robotics: distributed bootstrapping with consensus constraints
- **today**: collaborative system identification

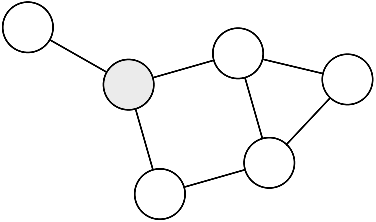

simple example

- input-output model

$$y(t) = \theta u(t)$$

- at each time $t \geq 0$:
  - select input $u(t) \in \mathbf{R}$
  - measure $y(t) \in \mathbf{R}$
- **goal**: determine θ

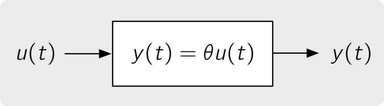

identification approach

- time-varying estimate $\hat{\theta}(t) \in \mathbf{R}$
- simulated output
$$\hat{y}(t) = \hat{\theta}(t)u(t)$$
- **our task**: make simulator match true model
$$(\hat{y}(t) - y(t))^2 \to 0 \quad \text{as} \quad t \to \infty$$

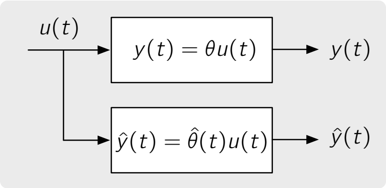

unconstrained minimization

minimize the instantaneous cost

$$J(\hat{\theta}(t)) = \frac{1}{2}(\hat{y}(t) - y(t))^2$$
$$= \frac{1}{2}(\underbrace{\hat{\theta}(t) - \theta}_{\Delta\theta(t)})^2 u(t)^2$$

by gradient descent on $\hat{\theta}(t)$

$$\frac{d}{dt}\hat{\theta}(t) := -\gamma \frac{\partial J}{\partial \hat{\theta}(t)}$$
$$= -\gamma \Delta\theta(t) u(t)^2,$$

where $\gamma > 0$ is the learning rate

gradient learning rule

- gradient rule can be implemented online

$$\frac{d}{dt}\hat{\theta}(t) = -\gamma \Delta\theta(t) u(t)^2$$
$$= -\gamma \underbrace{(\hat{y}(t) - y(t))}_{\Delta y(t)} u(t)$$

- output error: $\Delta y(t)$
- parameter error: $\Delta\theta(t)$
- **fact**: output error (usually) converges, $\Delta y(t) \to 0$ as $t \to \infty$ (proof: Lyapunov argument $V(\Delta\theta) = \Delta\theta^2$)
- **question**: when does parameter error converge?

$$\Delta\theta(t) \overset{?}{\to} 0 \quad \text{as} \quad t \to \infty$$

typical error curves

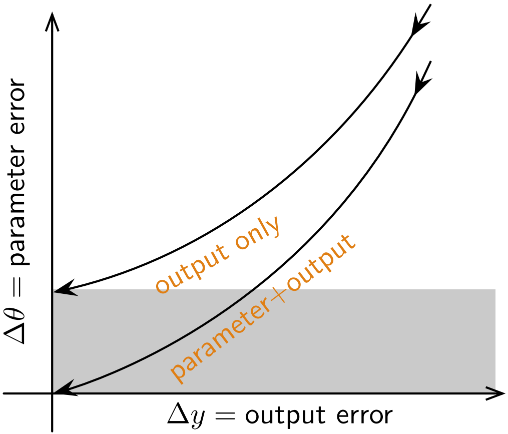

simple condition on parameter convergence

- parameter error dynamics

$$\begin{aligned}\frac{d}{dt}\Delta\theta(t) &= \frac{d}{dt}\left(\hat{\theta}(t) - \theta\right) \\ &= -\gamma\Delta\theta(t)u(t)^2 \\ &\Downarrow \\ \Delta\theta(t) &= \exp\left\{-\gamma\int_0^t u(\tau)^2\,d\tau\right\}\Delta\theta(0)\end{aligned}$$

- parameter error converges if $u(t)$ is **persistently exciting**:

$$\lim_{t\to\infty}\int_0^t u(\tau)^2\,d\tau = +\infty$$

checking the memoryless system

- choose input $u(t) = c$, where $c \neq 0$ is a real constant

$$\begin{aligned}\lim_{t\to\infty}\int_0^t u(\tau)^2\,d\tau &= \lim_{t\to\infty}\int_0^t c^2\,d\tau \\ &= \lim_{t\to\infty} c^2 t \\ &= +\infty \quad \checkmark\end{aligned}$$

- excitation condition:

$$u(t) = c \text{ is persistently exciting} \quad \Leftrightarrow \quad c \neq 0$$

- persistence of excitation guarantees parameter convergence

multiple agent identification model

- n agents labeled $i = 1, \ldots, n$
- at time $t \geq 0$, agent i can measure $x_i(t) \in \mathbf{R}^q$ and $y_i(t) \in \mathbf{R}$
- regressor: $\phi : \mathbf{R}^q \to \mathbf{R}^p$
- parameters: $\theta \in \mathbf{R}^p$
- true output:

$$y_i(t) = \theta^T \phi(x_i(t)), \quad i = 1, \ldots, n$$

- simulated output:

$$\hat{y}_i(t) = \hat{\theta}_i(t)^T \phi(x_i(t)), \quad i = 1, \ldots, n$$

- **goal**: parameter convergence $\|\theta_i(t) - \theta\| \to 0$ for all $i = 1, \ldots, n$.

multiple agent identification model

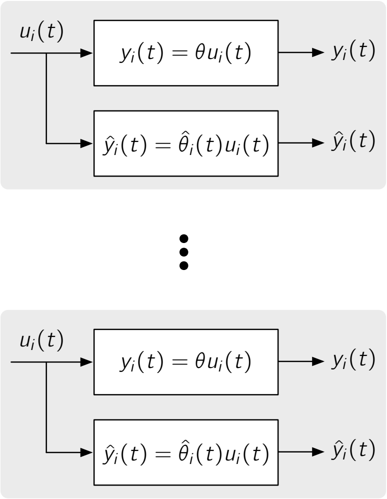

multiple agent consensus scheme

- each agent's parameter estimate is a sum of two terms

$$\frac{d}{dt}\hat{\theta}_i = \underbrace{-\gamma\phi(x_i)(\hat{y}_i - y_i)}_{\text{local information}} + \underbrace{\sum_{j \in \mathcal{N}_i} a_{ij}(\hat{\theta}_j - \hat{\theta}_i)}_{\text{neighboring information}}$$

- can be implemented **online**
- **respects** network communication structure

interpretations of consensus scheme

- gradient descent on instantaneous cost

$$J(\hat{\theta}_1, \ldots, \hat{\theta}_n) = \underbrace{\sum_{i=1}^{n} (\hat{y}_i(t) - y_i(t))^2}_{\text{identification objective}} + \underbrace{\sum_{\{v_i, v_j\} \in \mathcal{E}} \frac{1}{2} a_{ij} \|\hat{\theta}_j(t) - \hat{\theta}_i(t)\|_2^2}_{\text{disagreement objective}}$$

- distributed PD control
- dynamical model fusion (*cf.* sensor fusion)
- augmented Lagrangian flow

$$\begin{array}{ll} \text{minimize} & \displaystyle\sum_{i=1}^{n} (\hat{y}_i(t) - y_i(t))^2 \\ \text{subject to} & \hat{\theta}_j(t) - \hat{\theta}_i(t) = 0, \quad i, j = 1, \ldots, n \end{array}$$

convergence

candidate Lyapunov function:

$$V(\Delta\theta) = \sum_{i=1}^{n} \Delta\theta_i^T \Delta\theta_i$$

require:

- connected communication graph $\mathcal{G}$
- bounded (uniformly cts) regressors
- **collective** persistence of excitation

rate determined by:

- algebraic connectivity of $\mathcal{G}$
- minimum level of collective persistence of excitation

collective persistence of excitation

proof idea:

- error dynamics are (for $\theta, \theta_i \in \mathbf{R}^1$)

$$\frac{d}{dt}\Delta\theta(t) = -(\underbrace{L}_{\text{rank } n-1} + \gamma\Phi(t))\Delta\theta(t)$$

- for $\Delta\theta \to 0$, bound in every direction $w \in \mathbf{R}^n$

$$w^T \left(\frac{1}{t - t_0} \int_{t_0}^{t} L + \gamma\Phi(\tau)\, d\tau \right) w > 0$$

- **collective PE**: there exist positive real numbers $m_1, m_2 > 0$ such that for all $t_0 \geq 0$ and $t > t_0$ the matrix inequality

$$m_2 I \succeq \frac{1}{t - t_0} \int_{t_0}^{t} \sum_{i=1}^{n} \phi_i(\tau)\phi_i(\tau)^T\, d\tau \succeq m_1 I$$

excitation can be moved around

the following all imply parameter convergence:

- enlightened: a few ϕ_i are persistently exciting,
- total: every ϕ_i is persistently exciting,
- intermittent: there exists an unbounded sequence of times $t_1, t_2, \ldots$ such that some ϕ_i obeys the collective PE condition in each interval $[t_k, t_{k+1}]$,
- collaborative: none of the ϕ_i is persistently exciting, but the collective PE condition still holds.

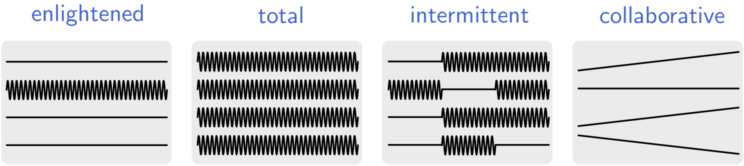

example: collaborative PE (w/o and w/ consensus)

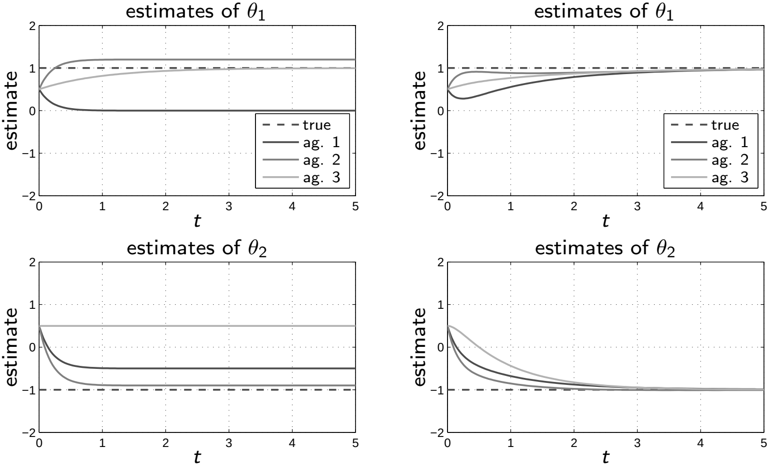

example: collaborative PE error curves

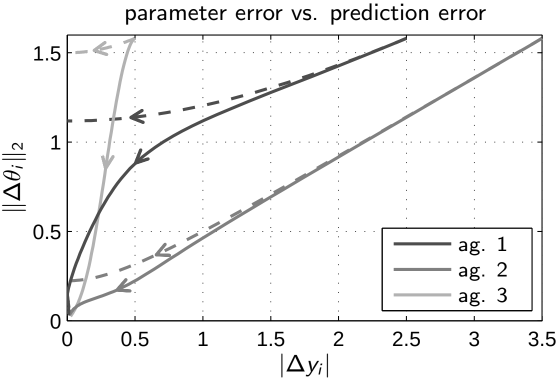

rate bound

take direction $w = \underbrace{\alpha \mathbf{1}/\sqrt{n}}_{\text{consensus subspace}} + \sum_{j=2}^{n} \beta_j v_j$

rate $\geq \inf_{|\alpha|\leq 1} \max \left\{ \lambda_2(1-\alpha^2), \gamma \frac{\alpha^2}{n} m_1 - 2\gamma m_2 \sqrt{\frac{\alpha^2}{n}(1-\alpha^2)} \right\}$

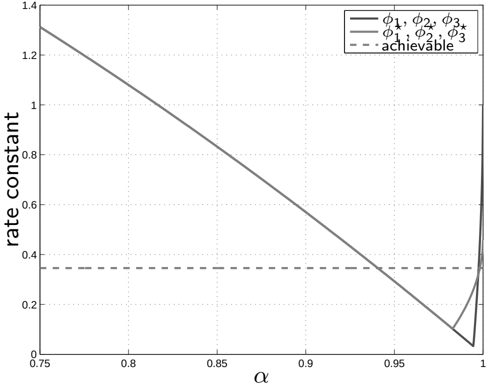

model reference adaptive control

- n van der pol (nonlinear) oscillators

$$\ddot{x}_i = -x_i + \mu(1 - x_i^2)\dot{x}_i + u_i, \quad i = 1, \ldots, n$$

- reference model for each oscillator (place poles at $-1 \pm j$)

$$\ddot{x}_i^{\text{ref}} = -2(x_i^{\text{ref}} + \dot{x}_i^{\text{ref}}), \quad i = 1, \ldots, n$$

- regressors

$$\phi(x_i) = (1 - x_i^2)\dot{x}_i, \quad i = 1, \ldots, n$$

- adaptation: two control gains per agent & $\mu > 0$
- consensus on μ only

model reference adaptive control

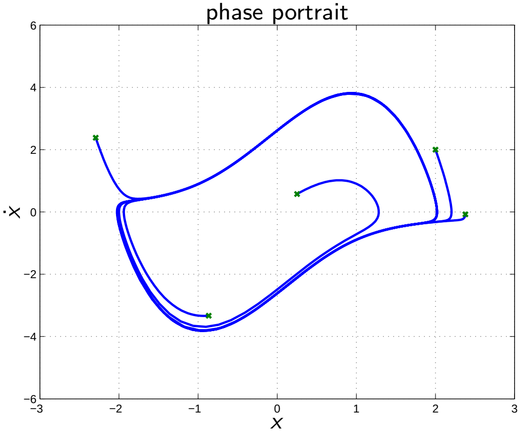

random initial conditions, $n = 5$ agents, open loop

model reference adaptive control

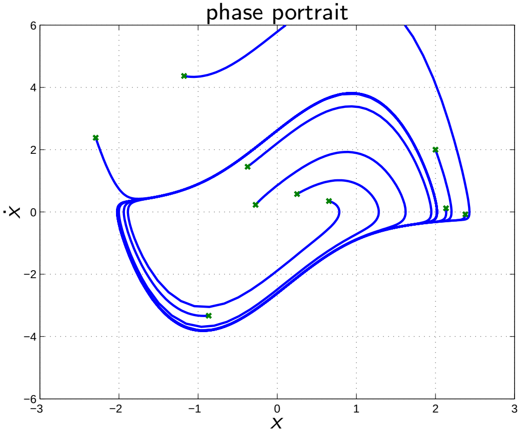

random initial conditions, $n = 10$ agents, open loop

model reference adaptive control

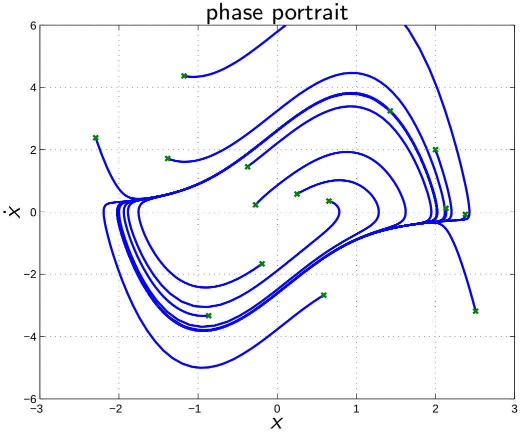

random initial conditions, $n = 15$ agents, open loop

model reference adaptive control

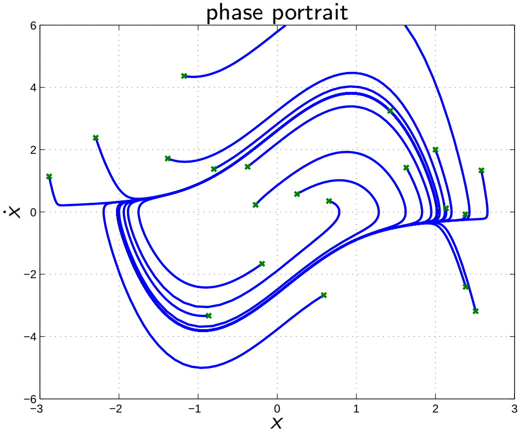

random initial conditions, $n = 20$ agents, open loop

model reference adaptive control

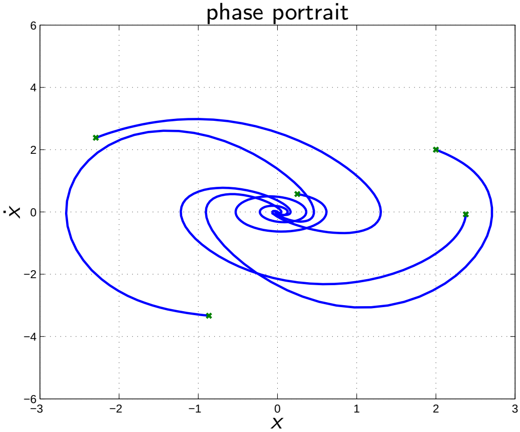

random initial conditions, $n = 5$ agents, MRAC

model reference adaptive control

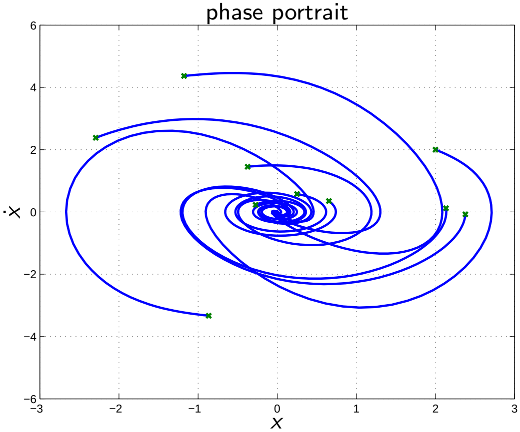

random initial conditions, $n = 10$ agents, MRAC

model reference adaptive control

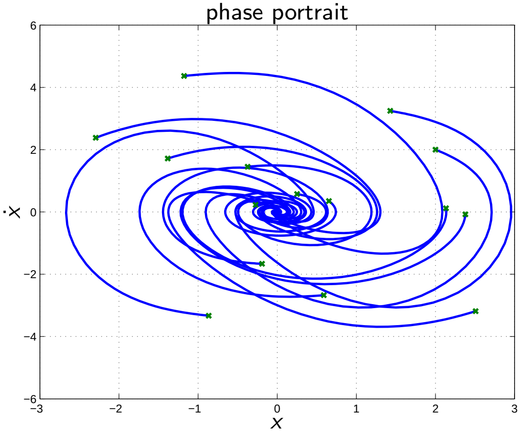

random initial conditions, $n = 15$ agents, MRAC

model reference adaptive control

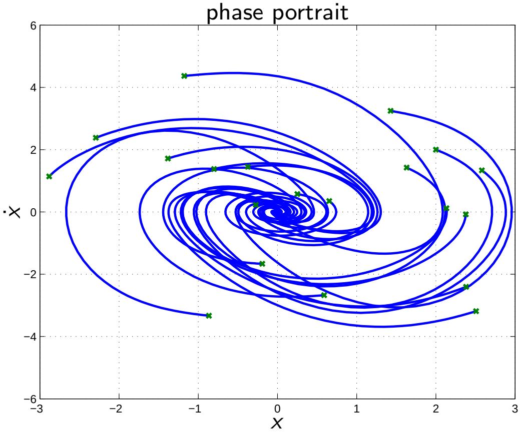

random initial conditions, $n = 20$ agents, MRAC

model reference adaptive control

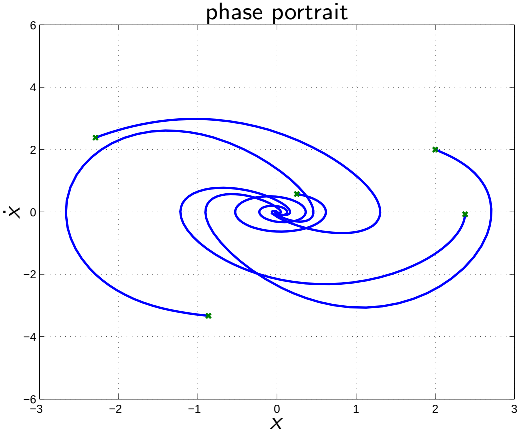

random initial conditions, $n = 5$ agents, MRAC + μ-consensus

model reference adaptive control

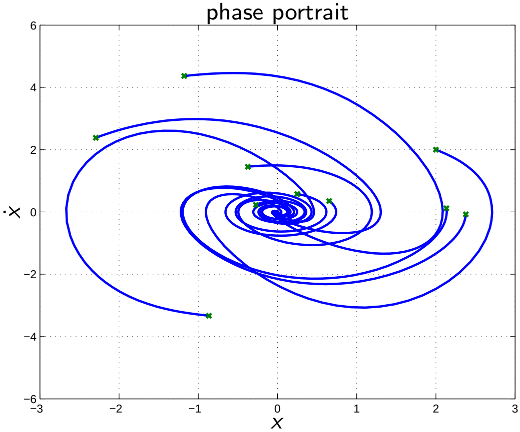

random initial conditions, $n = 10$ agents, MRAC + μ-consensus

model reference adaptive control

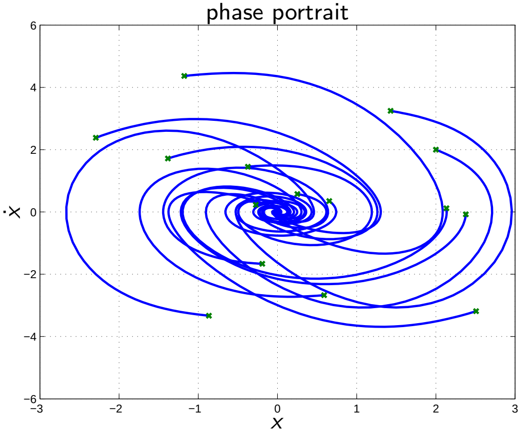

random initial conditions, $n = 15$ agents, MRAC + μ-consensus

model reference adaptive control

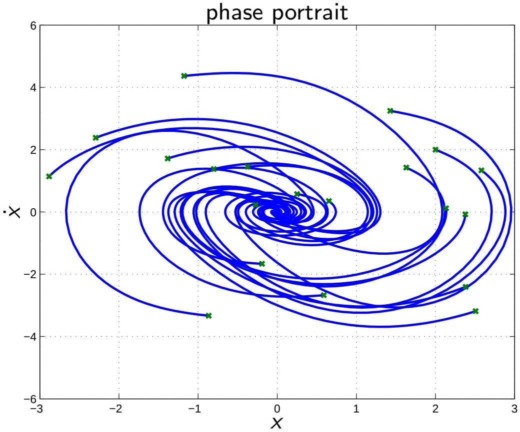

random initial conditions, $n = 20$ agents, MRAC + μ-consensus

model reference adaptive control

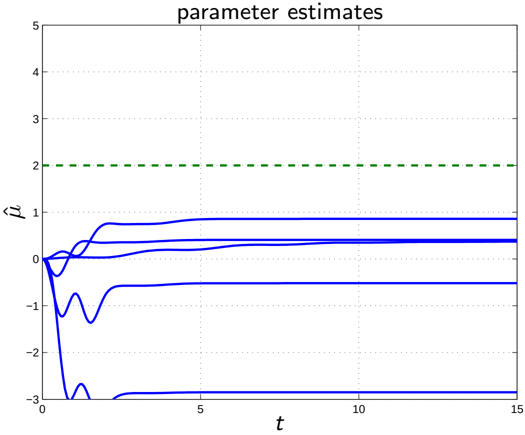

random initial conditions, $n = 5$ agents, MRAC

model reference adaptive control

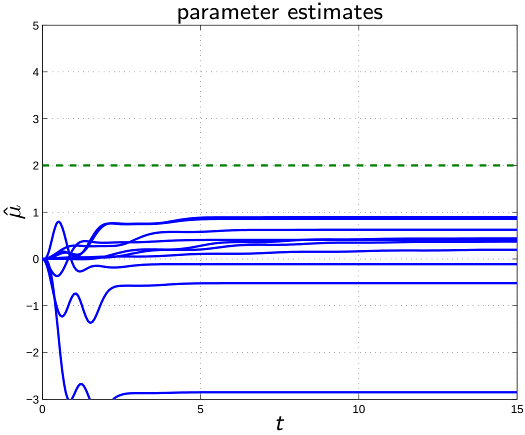

random initial conditions, $n = 10$ agents, MRAC

model reference adaptive control

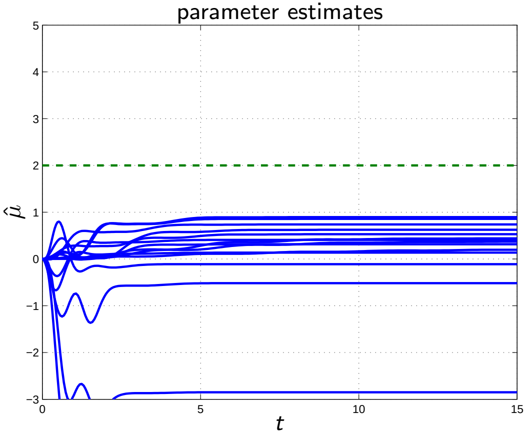

random initial conditions, $n = 15$ agents, MRAC

model reference adaptive control

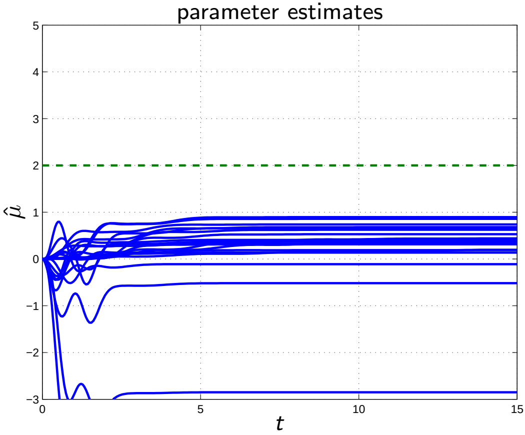

random initial conditions, $n = 20$ agents, MRAC

model reference adaptive control

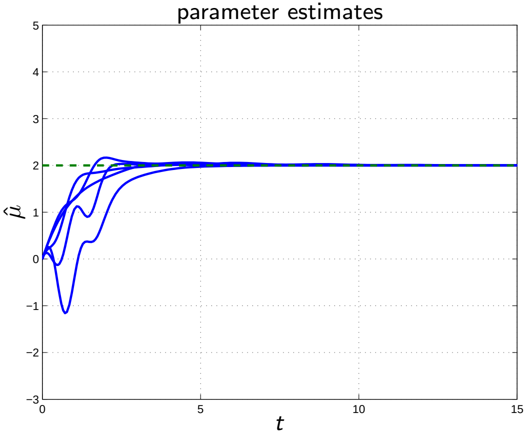

random initial conditions, $n = 5$ agents, MRAC + μ-consensus

model reference adaptive control

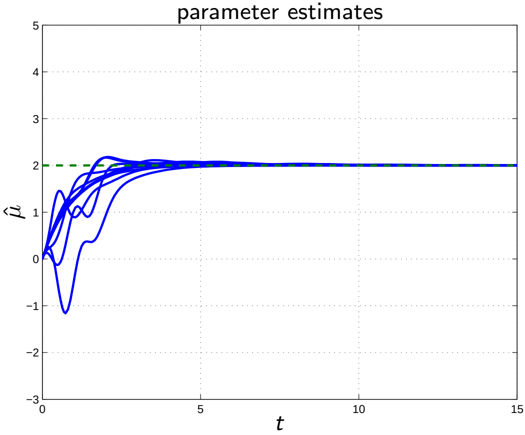

random initial conditions, $n = 10$ agents, MRAC + μ-consensus

model reference adaptive control

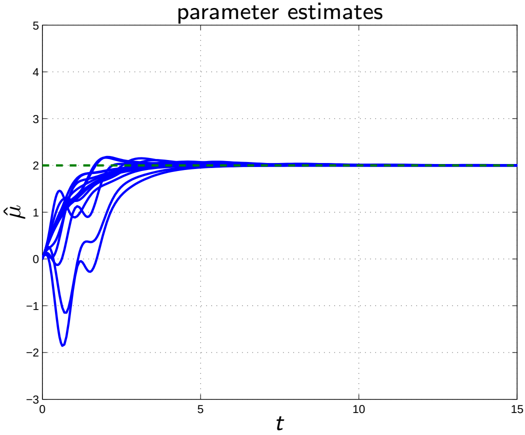

random initial conditions, $n = 15$ agents, MRAC + μ-consensus

model reference adaptive control

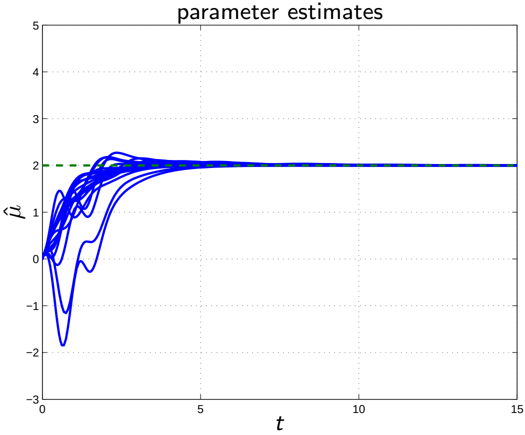

random initial conditions, $n = 20$ agents, MRAC + μ-consensus

summary

- simple idea: defined by

$$\hat{\theta}^{(t+1)} := \text{classical update rule} + \text{consensus}$$

- fundamentally nonlinear analysis and tools (mature theory)
- **future directions**:
  - quantitative analysis of noise effects (often) unchanged
  - engineer systems where the network does not fight adaptation
  - adaptation: graceful degradation when network fails
  - network: source of extra performance and robustness

thanks!

more information:

- Ivan Papusha: `www.cds.caltech.edu/~ipapusha`
- Richard M. Murray: `www.cds.caltech.edu/~murray`

funding:

- NDSEG
- Powell Foundation
- STARnet/TerraSwarm