Multirobot Systems

The consensus problem and applications

Advanced topics related to the consensus problem

Master Program in Robotics, Graphics and Computer Vision
Departamento de Informática e Ingeniería de Sistemas
Universidad de Zaragoza

Departamento de Informática e Ingeniería de Sistemas
Universidad Zaragoza

In this lecture

- The consensus problem: (Sketch of) proofs of convergence
- Effects of lower/higher connectivity
- Results for different graphs (disconnected, time varying topologies)
- Gossip consensus

(Recall) The consensus protocol in discrete time

- **Compact form**: $\mathbf{x}(\mathrm{k}+1) = \mathrm{W}\,\mathbf{x}(\mathrm{k})$

- Undirected graphs: weight matrices that are "**doubly stochastic**"

$$\mathrm{W}\,\mathbf{1} = \mathbf{1} \qquad \mathbf{1}^T\mathrm{W} = \mathbf{1}^T \qquad \mathbf{1} = (1, \dots, 1)^T$$

- For instance, is the Perron matrix (obtained from the Laplacian) doubly stochastic?:

$$\mathrm{W} = \mathrm{I} - \alpha\,\mathrm{L}$$ (with α positive 0 < α < 1/(2N))

- The Laplacian and Degree matrices (Framework):

$$L = \Delta - A \qquad \Delta = diag(d_i) = diag\left(\sum_{j=1}^{N} A_{ij}\right) \qquad \mathrm{L}\,\mathbf{1} = \mathbf{0}$$

$$\mathrm{W}\,\mathbf{1} = (\mathrm{I} - \alpha\,\mathrm{L})\mathbf{1} = \mathbf{1} - \alpha\,\mathrm{L}\mathbf{1} = \mathbf{1} - \mathbf{0} = \mathbf{1}$$

- Yes, it is stochastic (and doubly stochastic since it is symmetric)

(Sketch of) proof of convergence

- **Formal** definition of (average) consensus:

$$\lim_{k\to\infty} \mathbf{x}(\mathrm{k}) = \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0) = \frac{\left(x_1(0) + x_2(0) + \cdots + x_N(0)\right)}{N}\mathbf{1} = \begin{bmatrix}\mathrm{avg}\\ \vdots \\ \mathrm{avg}\end{bmatrix}$$

- We try to prove that: $\lim_{k\to\infty} \mathbf{x}(\mathrm{k}) = \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0)$

  for connected undirected graphs

- Is the "**average kept**"? (recall that $\mathbf{1}^T \mathrm{W} = \mathbf{1}^T$) Proof:

$$\mathbf{x}(\mathrm{k}+1) = \mathrm{W}\,\mathbf{x}(\mathrm{k})$$

$$\frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(\mathrm{k}+1) = \frac{\mathbf{1}\mathbf{1}^T}{N}\mathrm{W}\,\mathbf{x}(\mathrm{k}) = \frac{\mathbf{1}\mathbf{1}^T}{N}\,\mathbf{x}(\mathrm{k}) = \frac{\mathbf{1}\mathbf{1}^T}{N}\,\mathrm{W}\,\mathbf{x}(\mathrm{k}-1) = \cdots = \frac{\mathbf{1}\mathbf{1}^T}{N}\,\mathbf{x}(0)$$

- So <u>yes</u>, the average is kept

(Sketch of) proof of convergence

- **Error** vector: $\mathbf{e}(\mathrm{k}) \triangleq \mathbf{x}(\mathrm{k}) - \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0)$

- Does the error vector have zero-mean? => Yes!

$$\frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{e}(\mathrm{k}) \triangleq \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(\mathrm{k}) - \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0) = \mathbf{0}$$

- Thus: $\left(\mathrm{W} - \frac{\mathbf{1}\mathbf{1}^T}{N}\right)\mathbf{e}(\mathrm{k}) = \mathrm{W}\,\mathbf{e}(\mathrm{k})$

- Error dynamics (combining all the previous ideas):

$$\mathbf{e}(\mathrm{k}+1) = \mathbf{x}(\mathrm{k}+1) - \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0) = \mathrm{W}\,\mathbf{x}(\mathrm{k}) - \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0) =$$

$$\mathrm{W}\left(\mathbf{e}(\mathrm{k}) + \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0)\right) - \frac{\mathbf{1}\mathbf{1}^T}{N}\mathbf{x}(0) = \mathrm{W}\mathbf{e}(\mathrm{k}) = \left(\mathrm{W} - \frac{\mathbf{1}\mathbf{1}^T}{N}\right)\mathbf{e}(\mathrm{k})$$

(Sketch of) proof of convergence

Images: created by the lecturers of the course

- **Error** dynamics:

$$\mathbf{e}(\mathrm{k}) = \left(\mathrm{W} - \frac{\mathbf{1}\mathbf{1}^T}{N}\right)\mathbf{e}(\mathrm{k}-1) = \cdots = \left(\mathrm{W} - \frac{\mathbf{1}\mathbf{1}^T}{N}\right)^k \mathbf{e}(0)$$

- Norm of the error vector

$$||\mathbf{e}(\mathrm{k})|| \leq \cdots \leq= \left\|\left(\mathrm{W} - \frac{\mathbf{1}\mathbf{1}^T}{N}\right)\right\|^k \|\mathbf{e}(0)\|$$

- With more space/time, it can be proved (Laplacian entries + basic norm rules) that

$$\left\|\left(\mathrm{W} - \frac{\mathbf{1}\mathbf{1}^T}{N}\right)\right\| < 1$$

(Proof completed)

$\mathrm{k} = 0..10$

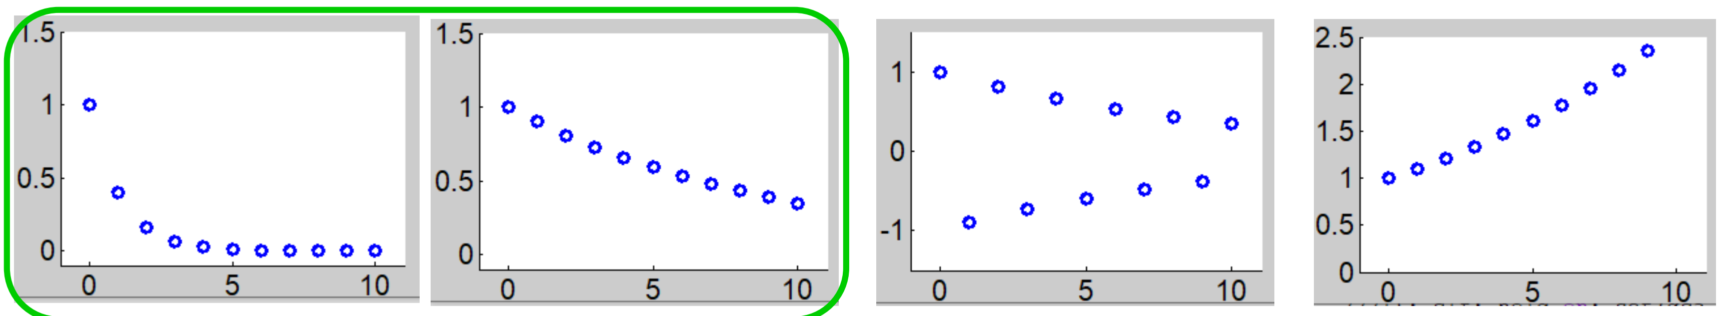

0.4^k $\quad 0.9^k$ $\quad (-0.9)^k$ $\quad 1.01^k$

Effects of lower/higher connectivity

Images: created by the lecturers of the course

- Adding links → higher connectivity

- Deleting links
- Adding robots (with few additional links)

} → lower connectivity

↓ ↑ $\left\| \left(W - \frac{\mathbf{1}\mathbf{1}^T}{N} \right) \right\|$ < 1

$$||\mathbf{e}(k)|| \leq \cdots \leq= \left\| \left(W - \frac{\mathbf{1}\mathbf{1}^T}{N} \right) \right\|^k ||\mathbf{e}(0)||$$

k= 0..10

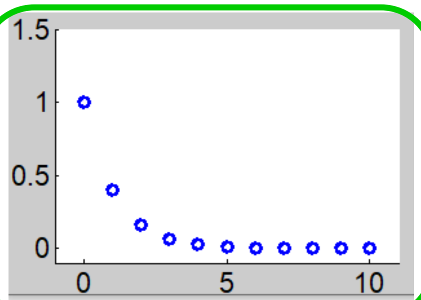

0.4^k

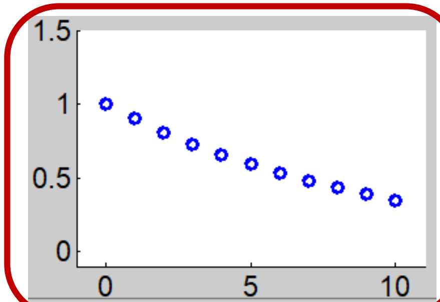

0.9^k

Seen in Exercise 1 (colab & python)

Results on disconnected graphs

Images: created by the lecturers of the course

Seen in Exercise 1 (colab & python)

$W =$ 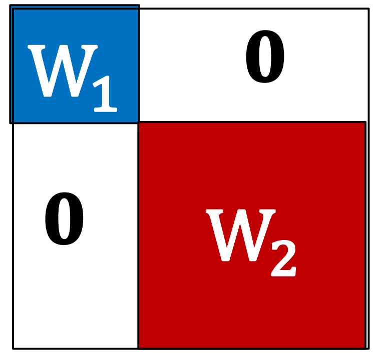

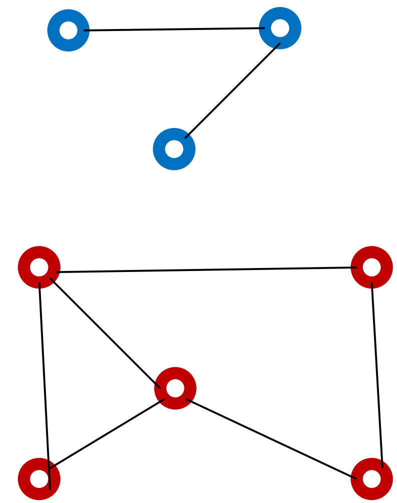

$$\left\| \left(W_1 - \frac{\mathbf{1}\mathbf{1}^T}{N} \right) \right\| < 1 \qquad \left\| \left(W_2 - \frac{\mathbf{1}\mathbf{1}^T}{N} \right) \right\| < 1$$

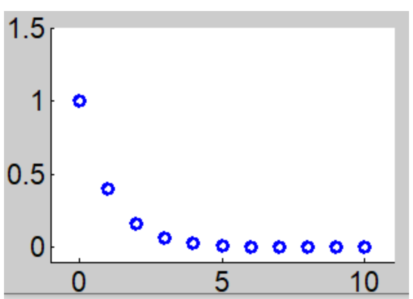

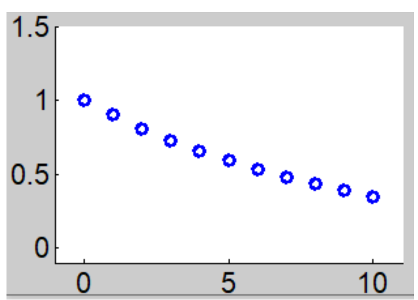

k= 0..10 $\quad 0.4^k \qquad 0.9^k$

Independent systems:

- ✓ Different rates
- ✓ Different consensus values

Results on time-varying graphs

Images: created by the lecturers of the course

Seen in Exercise 1 (colab & python)

$$W = W_1 \; W_2 \; W_3$$

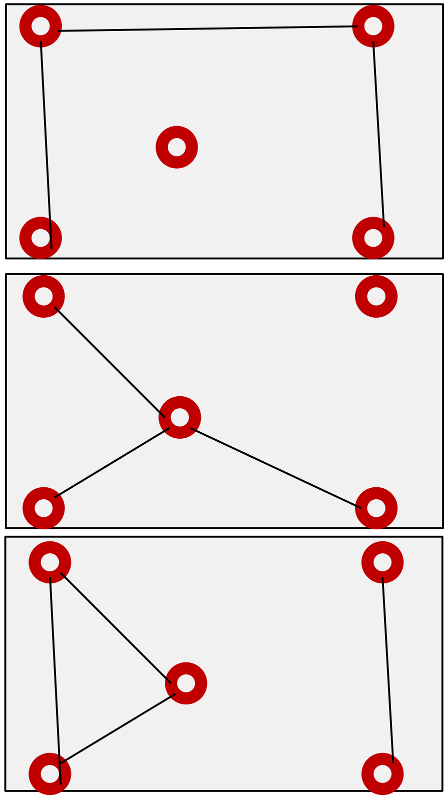

If the union of the graphs that occur *infinitely often* is *jointly connected*

Convergence rate? More, e.g., existence of an interval of *joint connectivity*

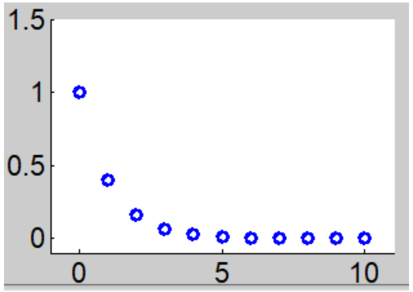

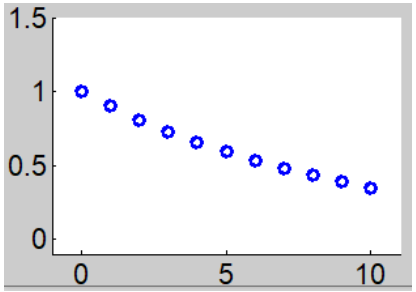

$\bigcup =$

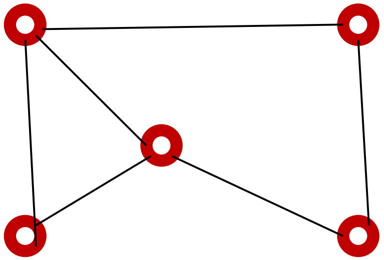

k= 0..10 $\quad 0.4^k \quad 0.9^k$

Gossip consensus

$\mathcal{U}$ =

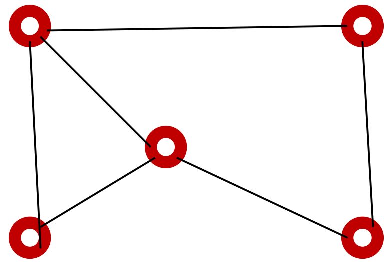

Seen in Exercise 1 (colab & python)

Select 1 link randomly

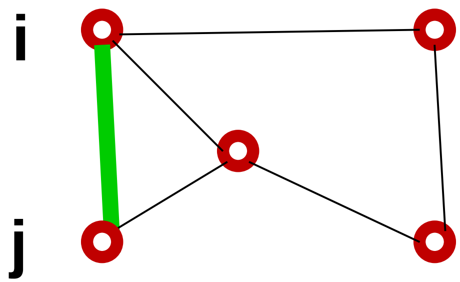

$$\begin{cases} x_i(k+1) = \dfrac{x_i(k) + x_j(k)}{2} \\ x_j(k+1) = \dfrac{x_i(k) + x_j(k)}{2} \\ x_{i'}(k+1) = x_{i'}(k) \end{cases}$$

Select 1 link randomly

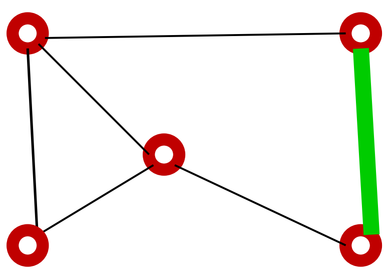

Is convergence ensured? What are the Metropolis weights at each step?

(...)

+ The union of the graphs that occur *infinitely often* is *jointly connected*

Images: created by the lecturers of the course

The last two results are important because…

- **Synchronous / asynchronous** multi-robot systems

  - **Impose synchronism?** (“rounds”)
  - What if a robot **fails** / leaves to perform other tasks?

Swarm of drones

https://www.youtube.com/watch?v=ZbkPJPg4kJY

Team of mobile manipulators

Image: Laboratories GmbH, CC BY-SA 3.0 <https://creativecommons.org/licenses/by-sa/3.0>, via Wikimedia Commons. https://commons.wikimedia.org/wiki/File:KUKA_omniRob.jpg

  - Which **level** of synchronization do these methods require?

Fixed graphs, classical consensus:

$$x_i(k+1) = x_i(k) + \alpha \sum_{j \,\epsilon\, N_i} (x_j(k) - x_i(k))$$

Gossip consensus:

$$x_i(k+1) = \frac{x_i(k) + x_j(k)}{2}$$

$$x_j(k+1) = \frac{x_i(k) + x_j(k)}{2}$$

$$x_{i'}(k+1) = x_{i'}(k)$$

  - How relevant is the **period** t=kT?

Exercise 2: ROS & multi-node

- **Synchronous / asynchronous** multi-robot systems

Topics: asynchronous communication

Services: **synchronous** communication

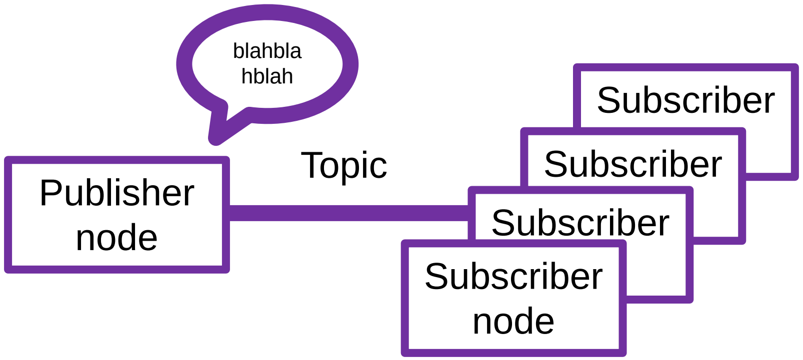

(Fixed graphs, **classical** consensus)

$$x_i(k+1) = x_i(k) + \alpha \sum_{j \in N_i} (x_j(k) - x_i(k))$$

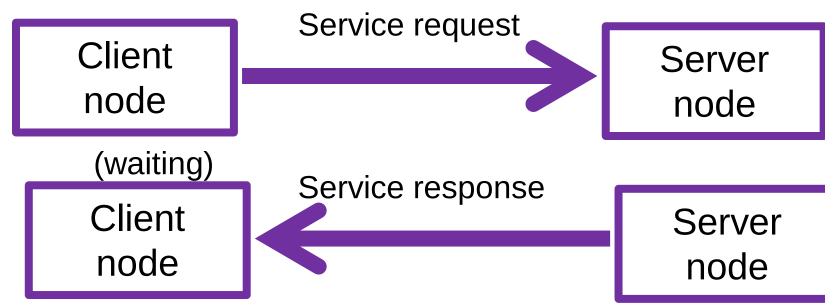

Gossip consensus

$$x_i(k+1) = \frac{x_i(k) + x_j(k)}{2}$$
$$x_j(k+1) = \frac{x_i(k) + x_j(k)}{2}$$
$$x_{i'}(k+1) = x_{i'}(k)$$

Image turtlebot: Kuscu0, CC BY-SA 4.0 <https://creativecommons.org/licenses/by-sa/4.0>, via Wikimedia Commons. https://commons.wikimedia.org/wiki/File:TurtleBot3_Burger.jpg Remaining images: created by the lecturers of the course

Dynamic consensus

- To **track** the average of references that **change** along time $\boldsymbol{r_i}(\mathrm{k})$

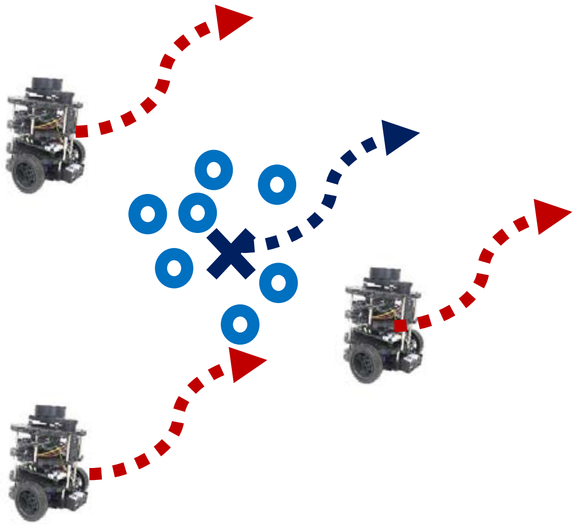

$$\boldsymbol{r_{avg}}(\mathrm{k}) = \frac{\boldsymbol{r_1}(\mathrm{k}) + \cdots + \boldsymbol{r_N}(\mathrm{k})}{N}$$

- Each i robot measures $\mathbf{r_i(k)}$. It initializes its state: $\boldsymbol{x_i}(0) = \boldsymbol{r_i}(0)$

$$\boldsymbol{x_i}(\mathrm{k}) = \boldsymbol{r_i}(\mathrm{k}) - \boldsymbol{r_i}(\mathrm{k}-1) +$$

Reference increment

$$+\boldsymbol{x_i}(\mathrm{k}-1) - \alpha \sum_{j\epsilon N_i} \left(\boldsymbol{x_i}(\mathrm{k}-1) - \boldsymbol{x_j}(\mathrm{k}-1) \right)$$

Consensus

Images: created by the lecturers of the course

Dynamic consensus

- **Convergence in undirected connected graphs**
- Conditions:
  - Parameter properly selected $\alpha \in \left(0, \frac{1}{d_{max}}\right)$
  - The reference increments are bounded

$$\mathbf{r_{max}} = \max_{k} \left\| (\mathbf{I} - \mathbf{1}\,\mathbf{1}^T/n)(\mathbf{r}(k) - \mathbf{r}(k+1)) \right\|_2 \text{ finite}$$

- Then, the states converge to a **neighborhood** of the average of the references $\boldsymbol{r_{avg}}(k)$

$$\lim_{k\to\infty} \left| x_i(k) - r_{avg}(k) \right| \le \frac{r_{max}}{\alpha \lambda_2}$$

(Framework)
Algebraic connectivity

- **Convergence in undirected connected graphs**

$$\alpha \in \left(0, \frac{1}{d_{max}}\right) \qquad \mathbf{r}_{\mathbf{max}} = \max_{k} \left\| (\mathbf{I} - \mathbf{1}\,\mathbf{1}^T/n)(\mathbf{r}(k) - \mathbf{r}(k+1)) \right\|_2 \ \text{finite}$$

- The states converge to a **neighborhood** of $r_{avg}(k)$

$$\lim_{k \to \infty} \left| x_i(k) - r_{avg}(k) \right| \le \frac{r_{max}}{\alpha\, \lambda_2}$$

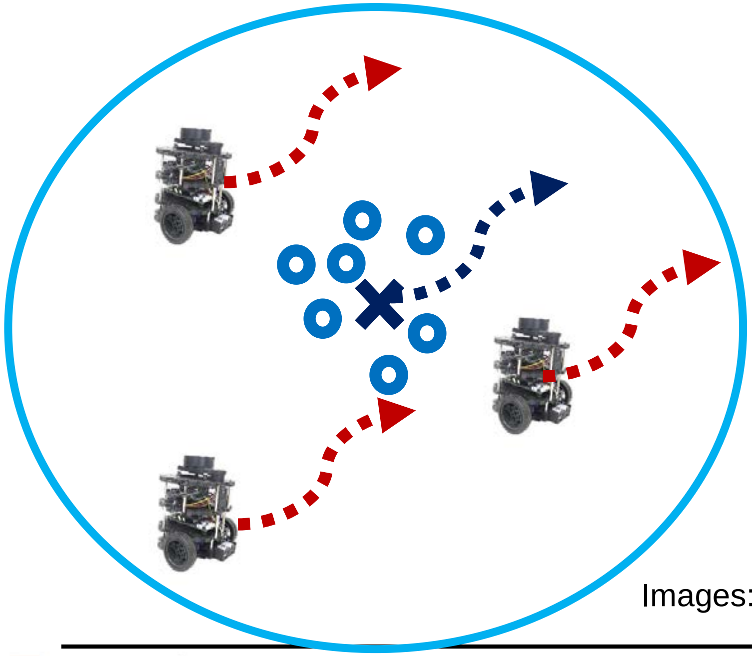

More communication links ?

λ_2 ↑ (Increase connectivity)

Neighborhood shrinks (more accurate)

Images: created by the lecturers of the course

- **Convergence in undirected connected graphs**

$$\alpha \in \left(0, \frac{1}{d_{max}}\right) \qquad \mathbf{r_{max}} = \max_{k} \left\| (\mathbf{I} - \mathbf{1}\,\mathbf{1}^T/n)(\mathbf{r}(k) - \mathbf{r}(k+1)) \right\|_2 \text{ finite}$$

- The states converge to a **neighborhood** of $r_{avg}(k)$

$$\lim_{k\to\infty} |x_i(k) - r_{avg}(k)| \le \frac{r_{max}}{\alpha\,\lambda_2}$$

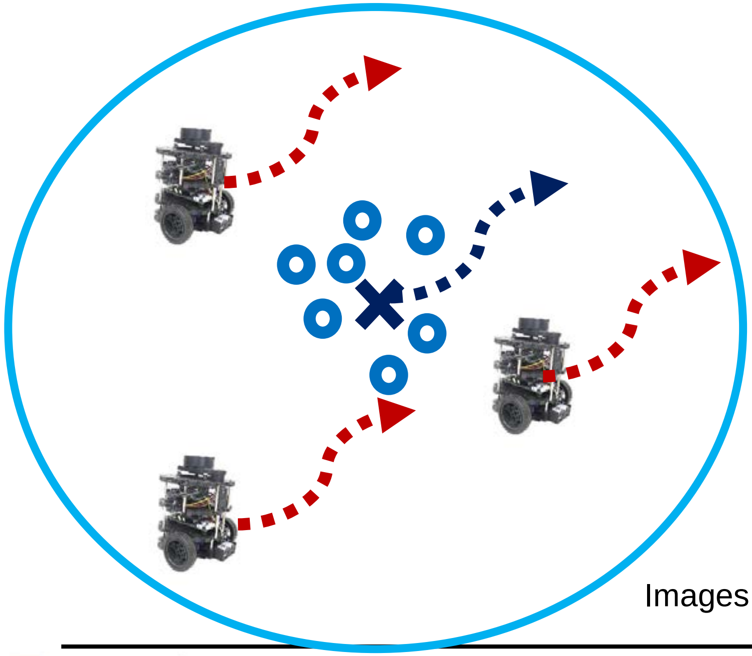

Any other alternatives to shrink the neighborhood?

α ↑ (there is a limit !)

Update more frequently (reference increments between k and k+1 smaller)

Images: created by the lecturers of the course

More results for other graphs

- **Strongly connected directed graphs**:

$$\mathbf{x}(\mathrm{k}+1) = \mathrm{W}\,\mathbf{x}(\mathrm{k}) \qquad \mathrm{W} = \mathbf{I} - \alpha\,\mathbf{L} \qquad \alpha \in \left(0, \frac{1}{d_{max}}\right)$$

- (Framework) A **directed** graph is said **strongly connected** if there exists a (directed) path joining any two nodes
- (Framework) $\mathrm{d}_{\mathrm{max}}$ is the maximum degree (max. number of neighbors)

- Convergence to consensus (**not** average!)

$$\lim_{k\to\infty} \mathbf{x}(\mathrm{k}) = \frac{\mathbf{1}\boldsymbol{w}^T}{\boldsymbol{w}^T\mathbf{1}}\mathbf{x}(0) = \frac{\left(w_1x_1(0) + w_2x_2(0) + \cdots + w_Nx_N(0)\right)}{w_1 + w_2 + \cdots + w_N}\mathbf{1} = \begin{bmatrix}\text{consen.}\\ \vdots \\ \text{consen.}\end{bmatrix}$$

where **w** is the left eigenvector of W associated with the eigenvalue 1

$$\boldsymbol{w}^T W = \boldsymbol{w}^T$$

Key idea: consensus (not necessarily the average)

More results for other graphs

- **Weakly connected directed graphs with leaders and followers**
- Leader robots: $\boldsymbol{x_i}(\mathrm{k}+1) = \boldsymbol{x_i}(\mathrm{k})$
- Follower robots:

Seen in Exercise 1 (colab & python)

$$\boldsymbol{x_i}(\mathrm{k}+1) = \boldsymbol{x_i}(\mathrm{k}) - \alpha \sum_{j \in N_i} \left(\boldsymbol{x_i}(\mathrm{k}) - \boldsymbol{x_j}(\mathrm{k})\right)$$

  - (Framework) N_i are the neighbors of robot i
  - Here, neighbors include leaders and followers

- One leader: asymptotic convergence of followers to the leader state
- Several leaders (containment control): asymptotic convergence of followers to the *convex hull* of the leader states (minimal convex set containing all the leader states, formally:)

$$\text{ConvexHull}(z_1, \cdots, z_N) = \{\textstyle\sum_{i=1}^{N} a_i z_i \,|a_i > 0, a_1 + \cdots + a_N = 1\}$$

Main ideas in this lecture ?

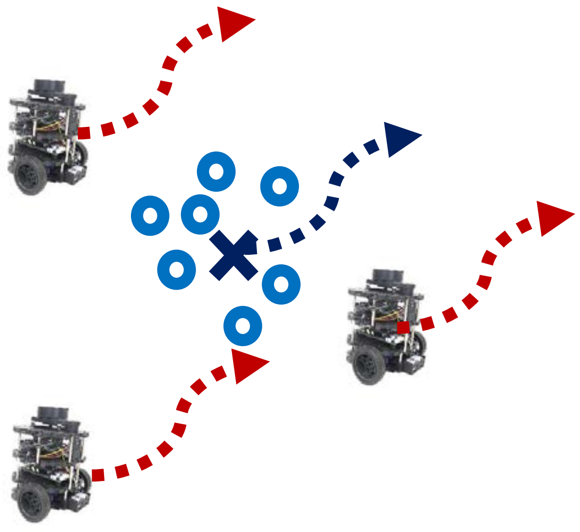

- Main results (and sketch of the proof) on convergence of the consensus protocol
- Different communication schemes
- Gossip: only requires synchronizing 2 robots
- Dynamic consensus (tracking the average)
- Some results for other graphs

- (…) Tons of additional theoretical / applied results in the literature !

Next session (the last session of this Part)

- Lab: multi-node or multi-robot (gossip, asynchronous)
  - Rendezvous
  - Deployment on a line
- Lab environment: Colab / Python / SmallWorl2D / ROS / ROS turtlesim / ROS & Gazebo (have your setup ready)

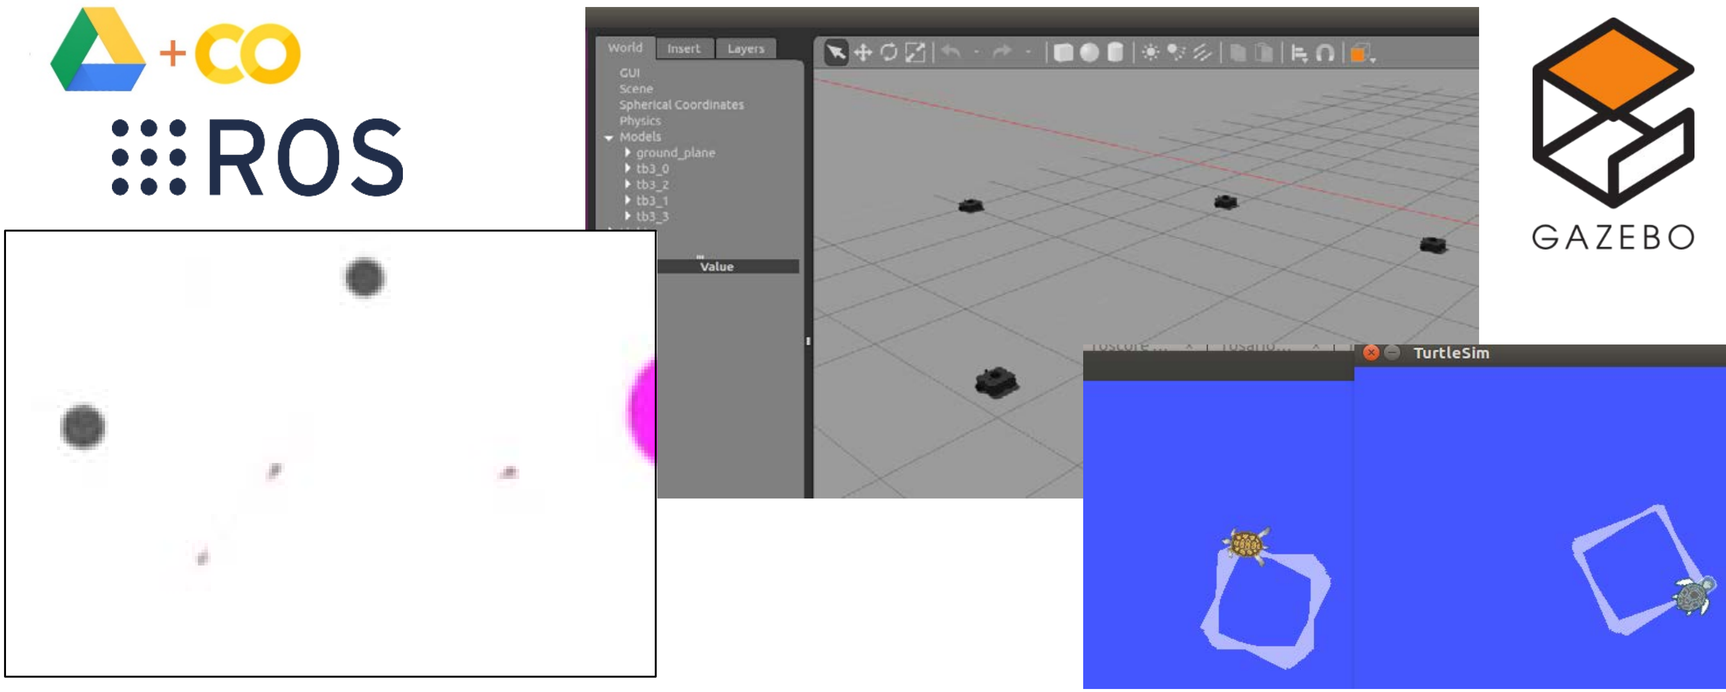

Images: results obtained by the lecturers of the course after running simulation software

Bibliography

- Xiao, L., Boyd, S., & Lall, S. (2005, April). A scheme for robust distributed sensor fusion based on average consensus. In IPSN 2005. Fourth International Symposium on Information Processing in Sensor Networks, 2005. (pp. 63-70). IEEE.
- Reza Olfati-Saber, J. Alex Fax, and Richard M. Murray. Consensus and Cooperation in Networked Multi-Agent Systems. Proceedings of the IEEE (95)1:215-233, 2007.
- L. Elsner, I. Koltracht, and M. Neumann, “On the convergence of asynchronous paracontractions with application to tomographic reconstruction from incomplete data,” Linear Algebra Appl., vol. 130, pp. 65–82, 1990.
- Cortés, J. Global and robust formation-shape stabilization of relative sensing networks. Automatica 45 (2009): 2754-2762.
- Franceschelli M., Gasparri A. Gossip-Based Centroid and Common Reference Frame Estimation in Multiagent Systems. IEEE TRANSACTIONS ON ROBOTICS 30 (2): 524-531 , 2014.

Bibliography

- Kan, Z., John M. Shea, J. M, Dixon, W. E. Leader–follower containment control over directed random graphs. Automatica 66 (2016) 56–62.
- Kia, S. S., Van Scoy, B., Cortes, J., Freeman, R. A., Lynch, K. M., & Martinez, S. (2019). Tutorial on dynamic average consensus: The problem, its applications, and the algorithms. IEEE Control Systems Magazine, 39(3), 40-72.
- Kia, S. S., Cortés, J and Martínez, S. Dynamic average consensus under limited control authority and privacy requirements. Int. J. Robust Nonlinear Control 2015; 25:1941–1966.