Consensus dynamics for sensor networks

Sandro Zampieri - Universita' di Padova

Alessandro Chiuso, Luca Schenato, Gianluigi Pillonetto, Federica Garin, Simone Del Favero, Saverio Bolognani, Damiano Ravagnolo, Enrico Lovisari (Universita' di Padova)

Ruggero Carli, Francesco Bullo (Universita' di California Santa Barbara)

Paolo Frasca (CNR Roma)

Fabio Fagnani (Politecnico di Torino)

Giacomo Como (MIT)

Jean-Charles Delvenne (University of Louvain la Nueve)

Applications

1. Mobile multi-vehicles coordination (rendezvous, formation, ciclic pursuit, coverage, ...).
2. Distributed estimation and control for sensor/ actuator networks.
3. Sensor calibration for sensor networks (e.g. clock synchronization).
4. Load balancing for distributed computing systems.
5. Distributed optimization algorithms.

Scientific context

COOPERATION: Simple global behavior from local interactions

Flocking: collective animal behavior given by the motion of a large number of coordinated individuals

Scientific context

Social and economic networks: individual social and economic interactions produce global phenomena

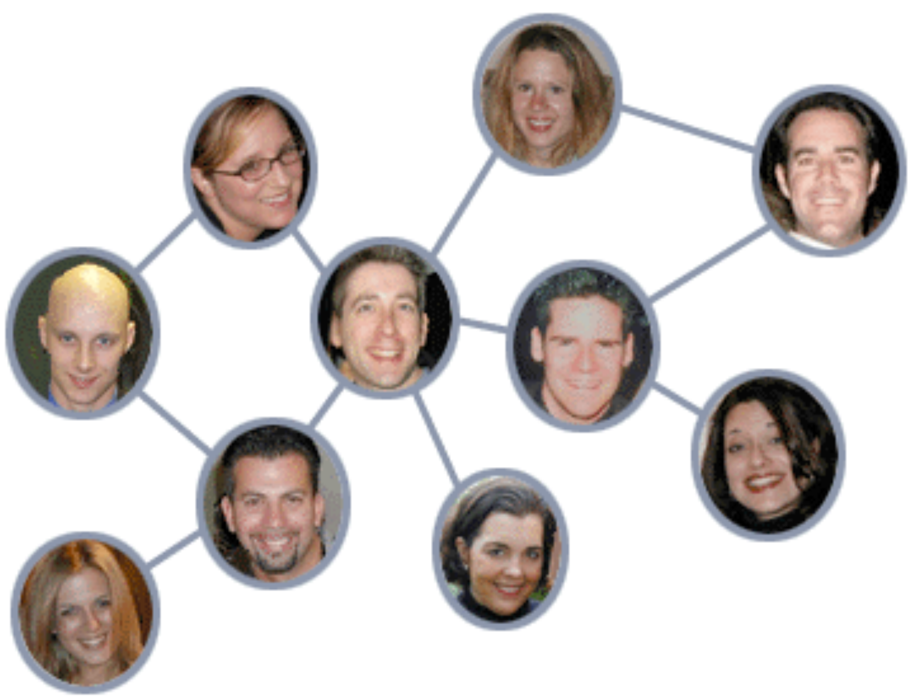

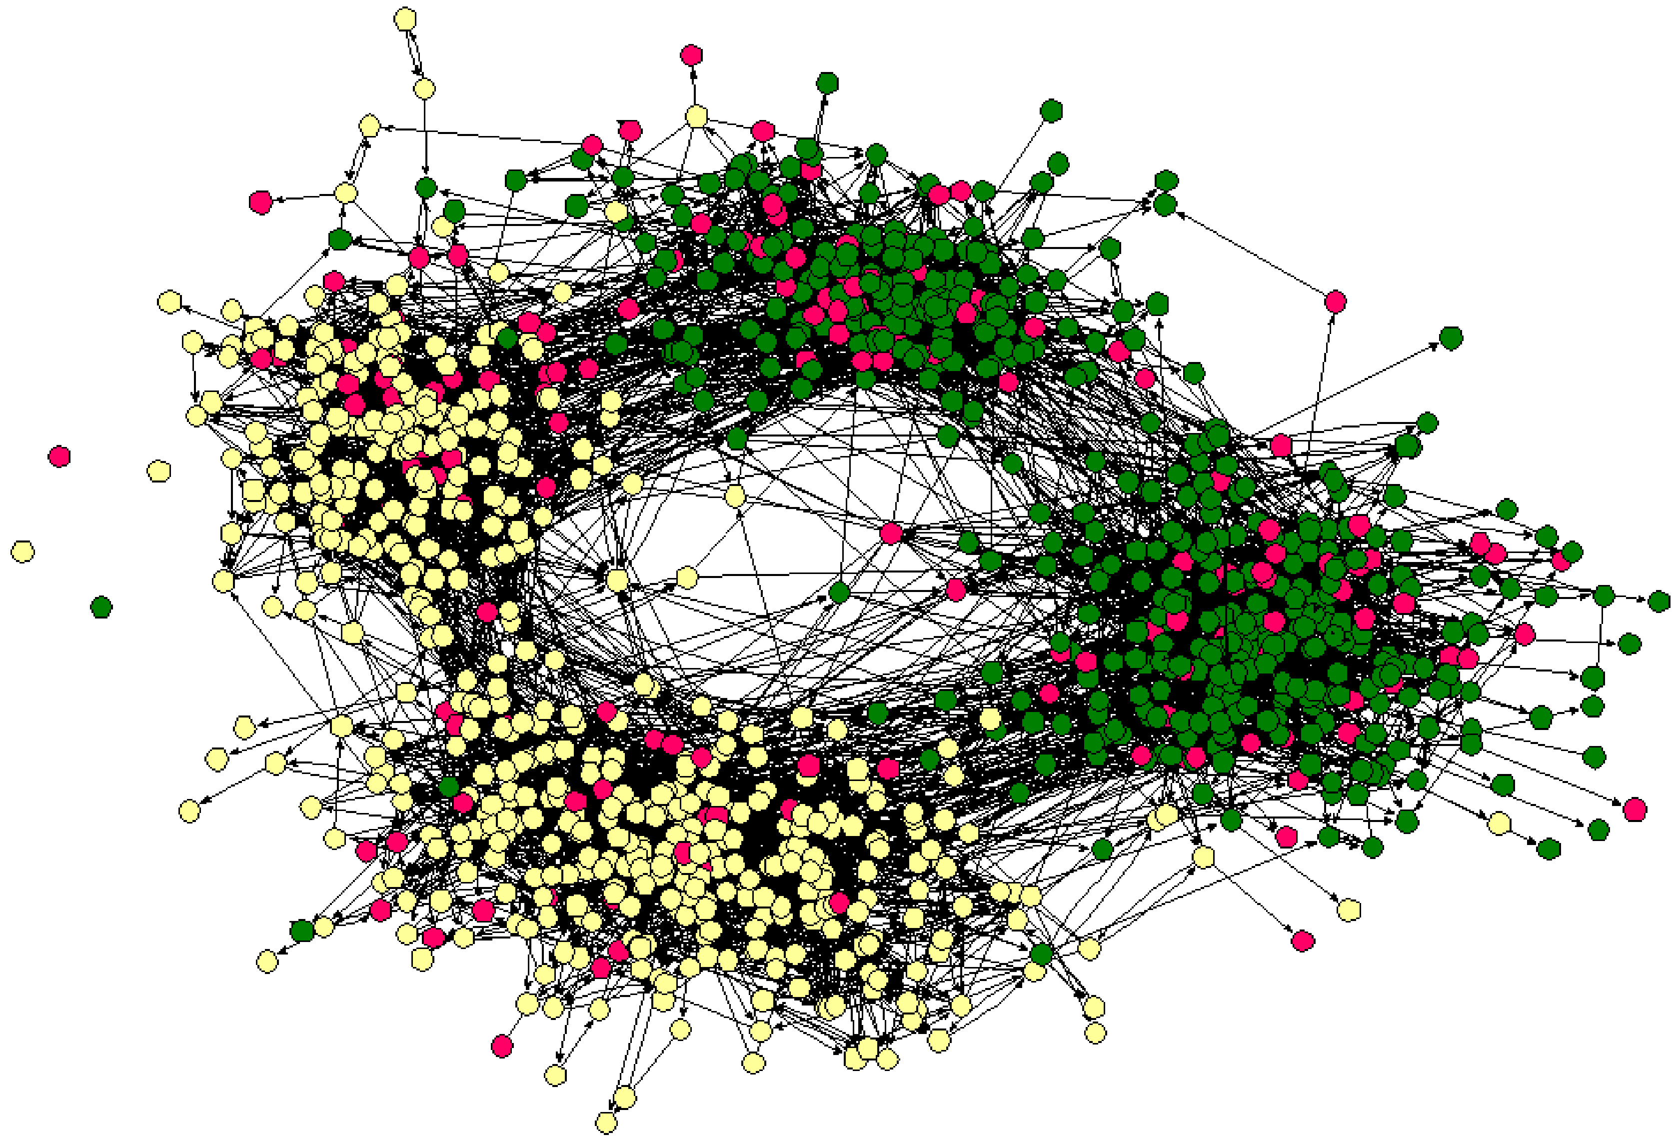

Graph describing friendship relations in an high school

Scientific context

Google page rank: from the complex internet web pages link connections to a global absolute web pages relevance evaluation

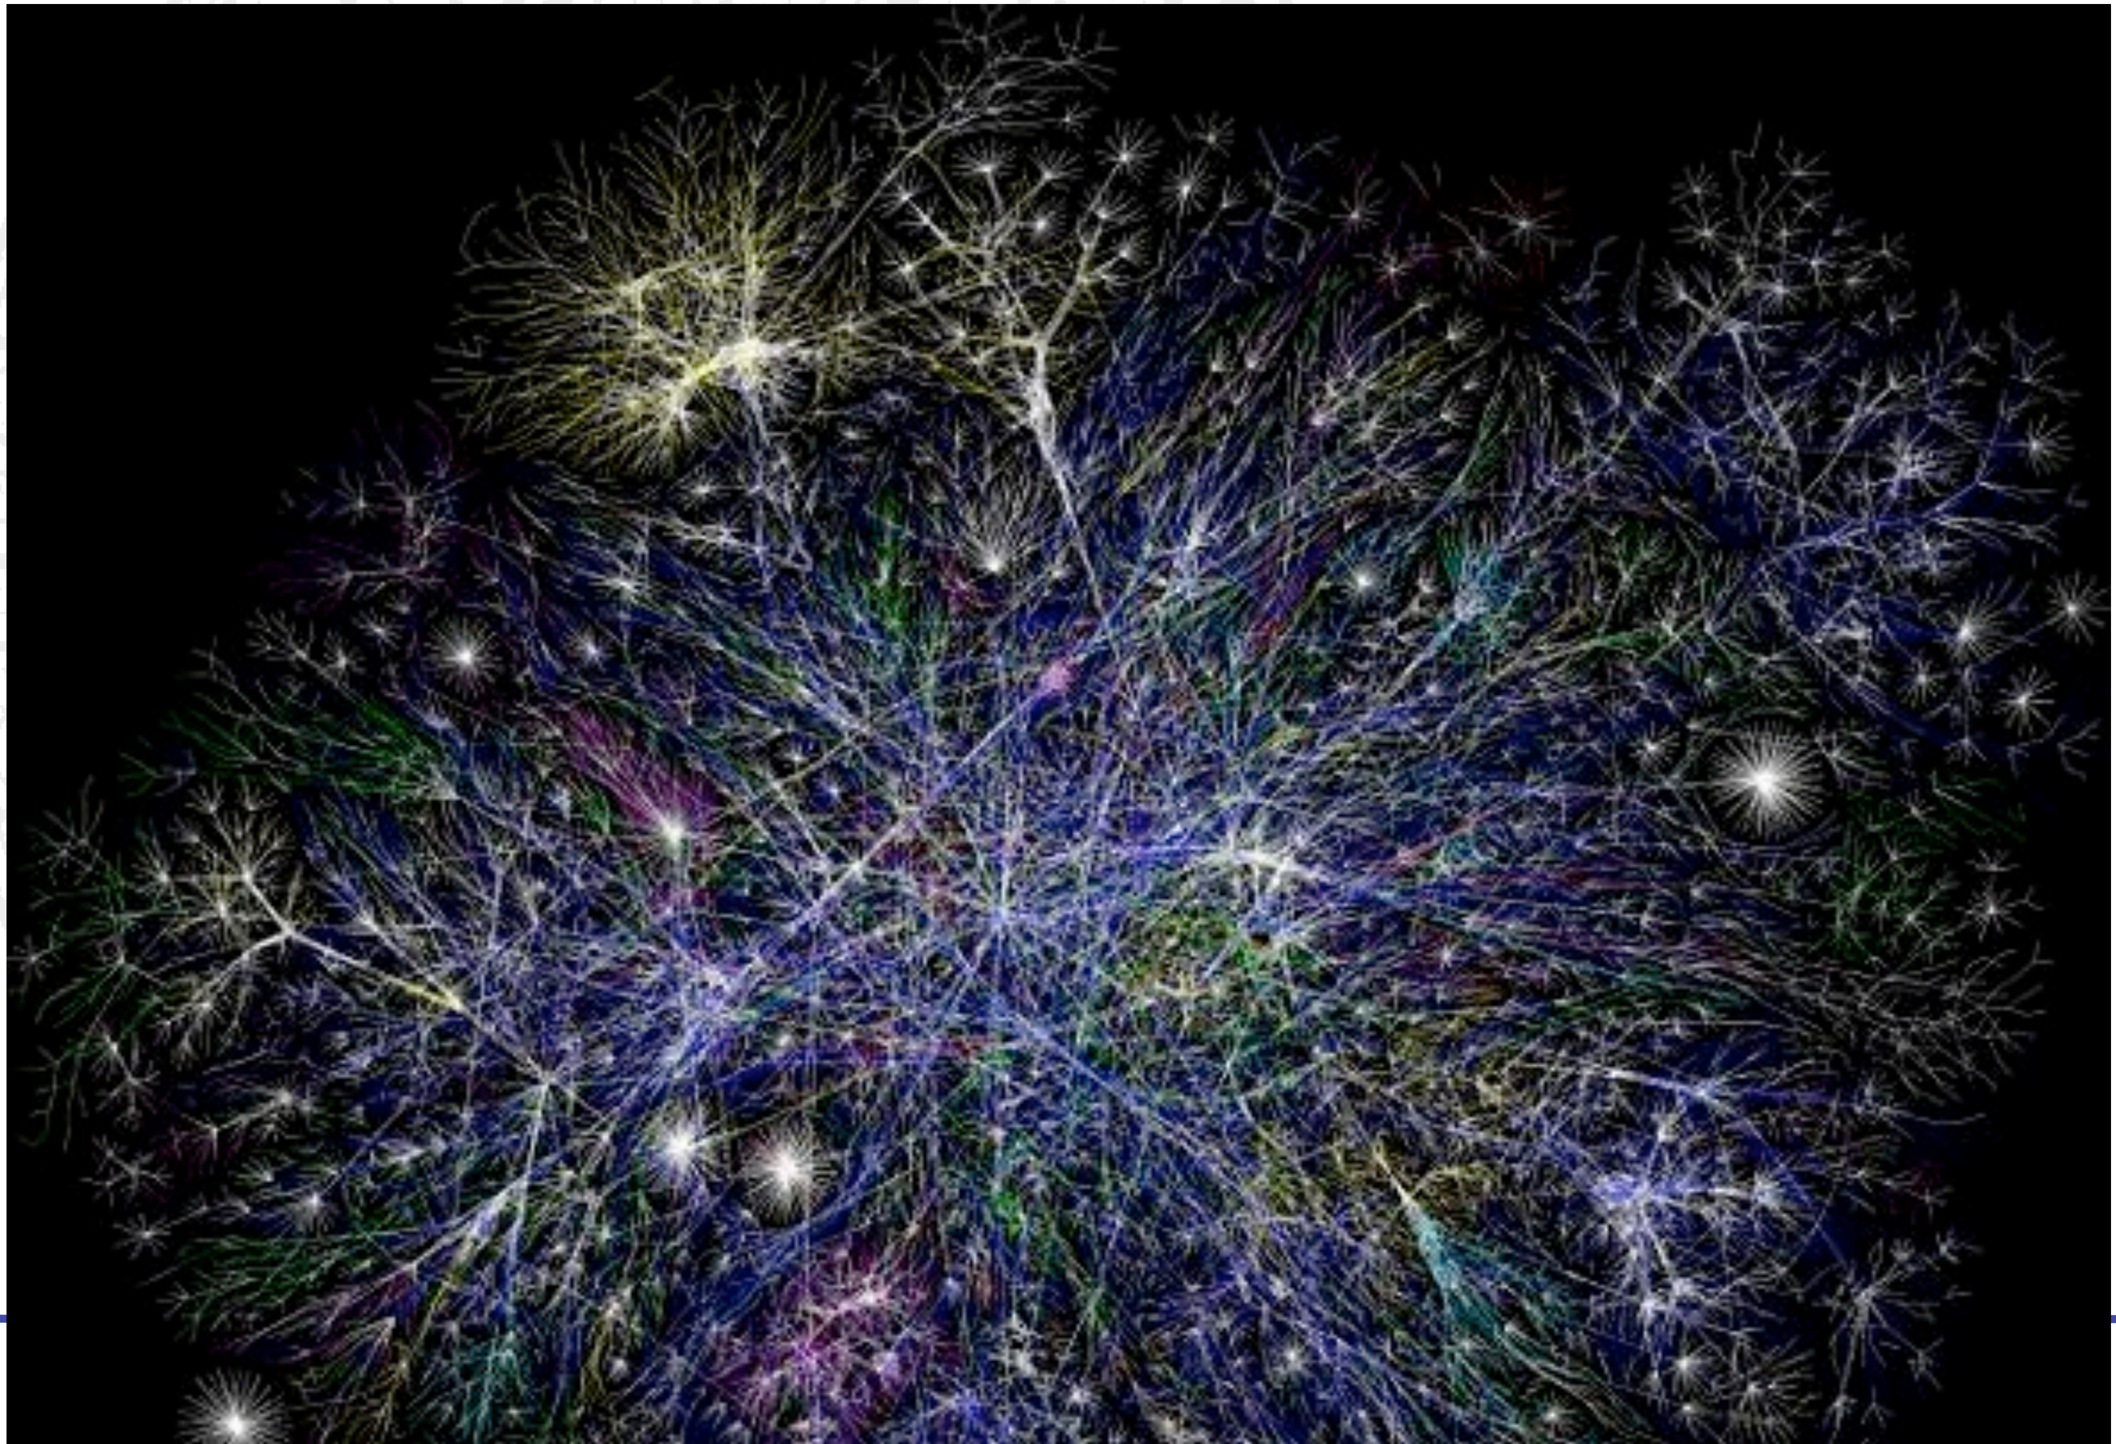

Scientific context

Complex biological systems: need for new instruments that allow to deal with complex interaction structures

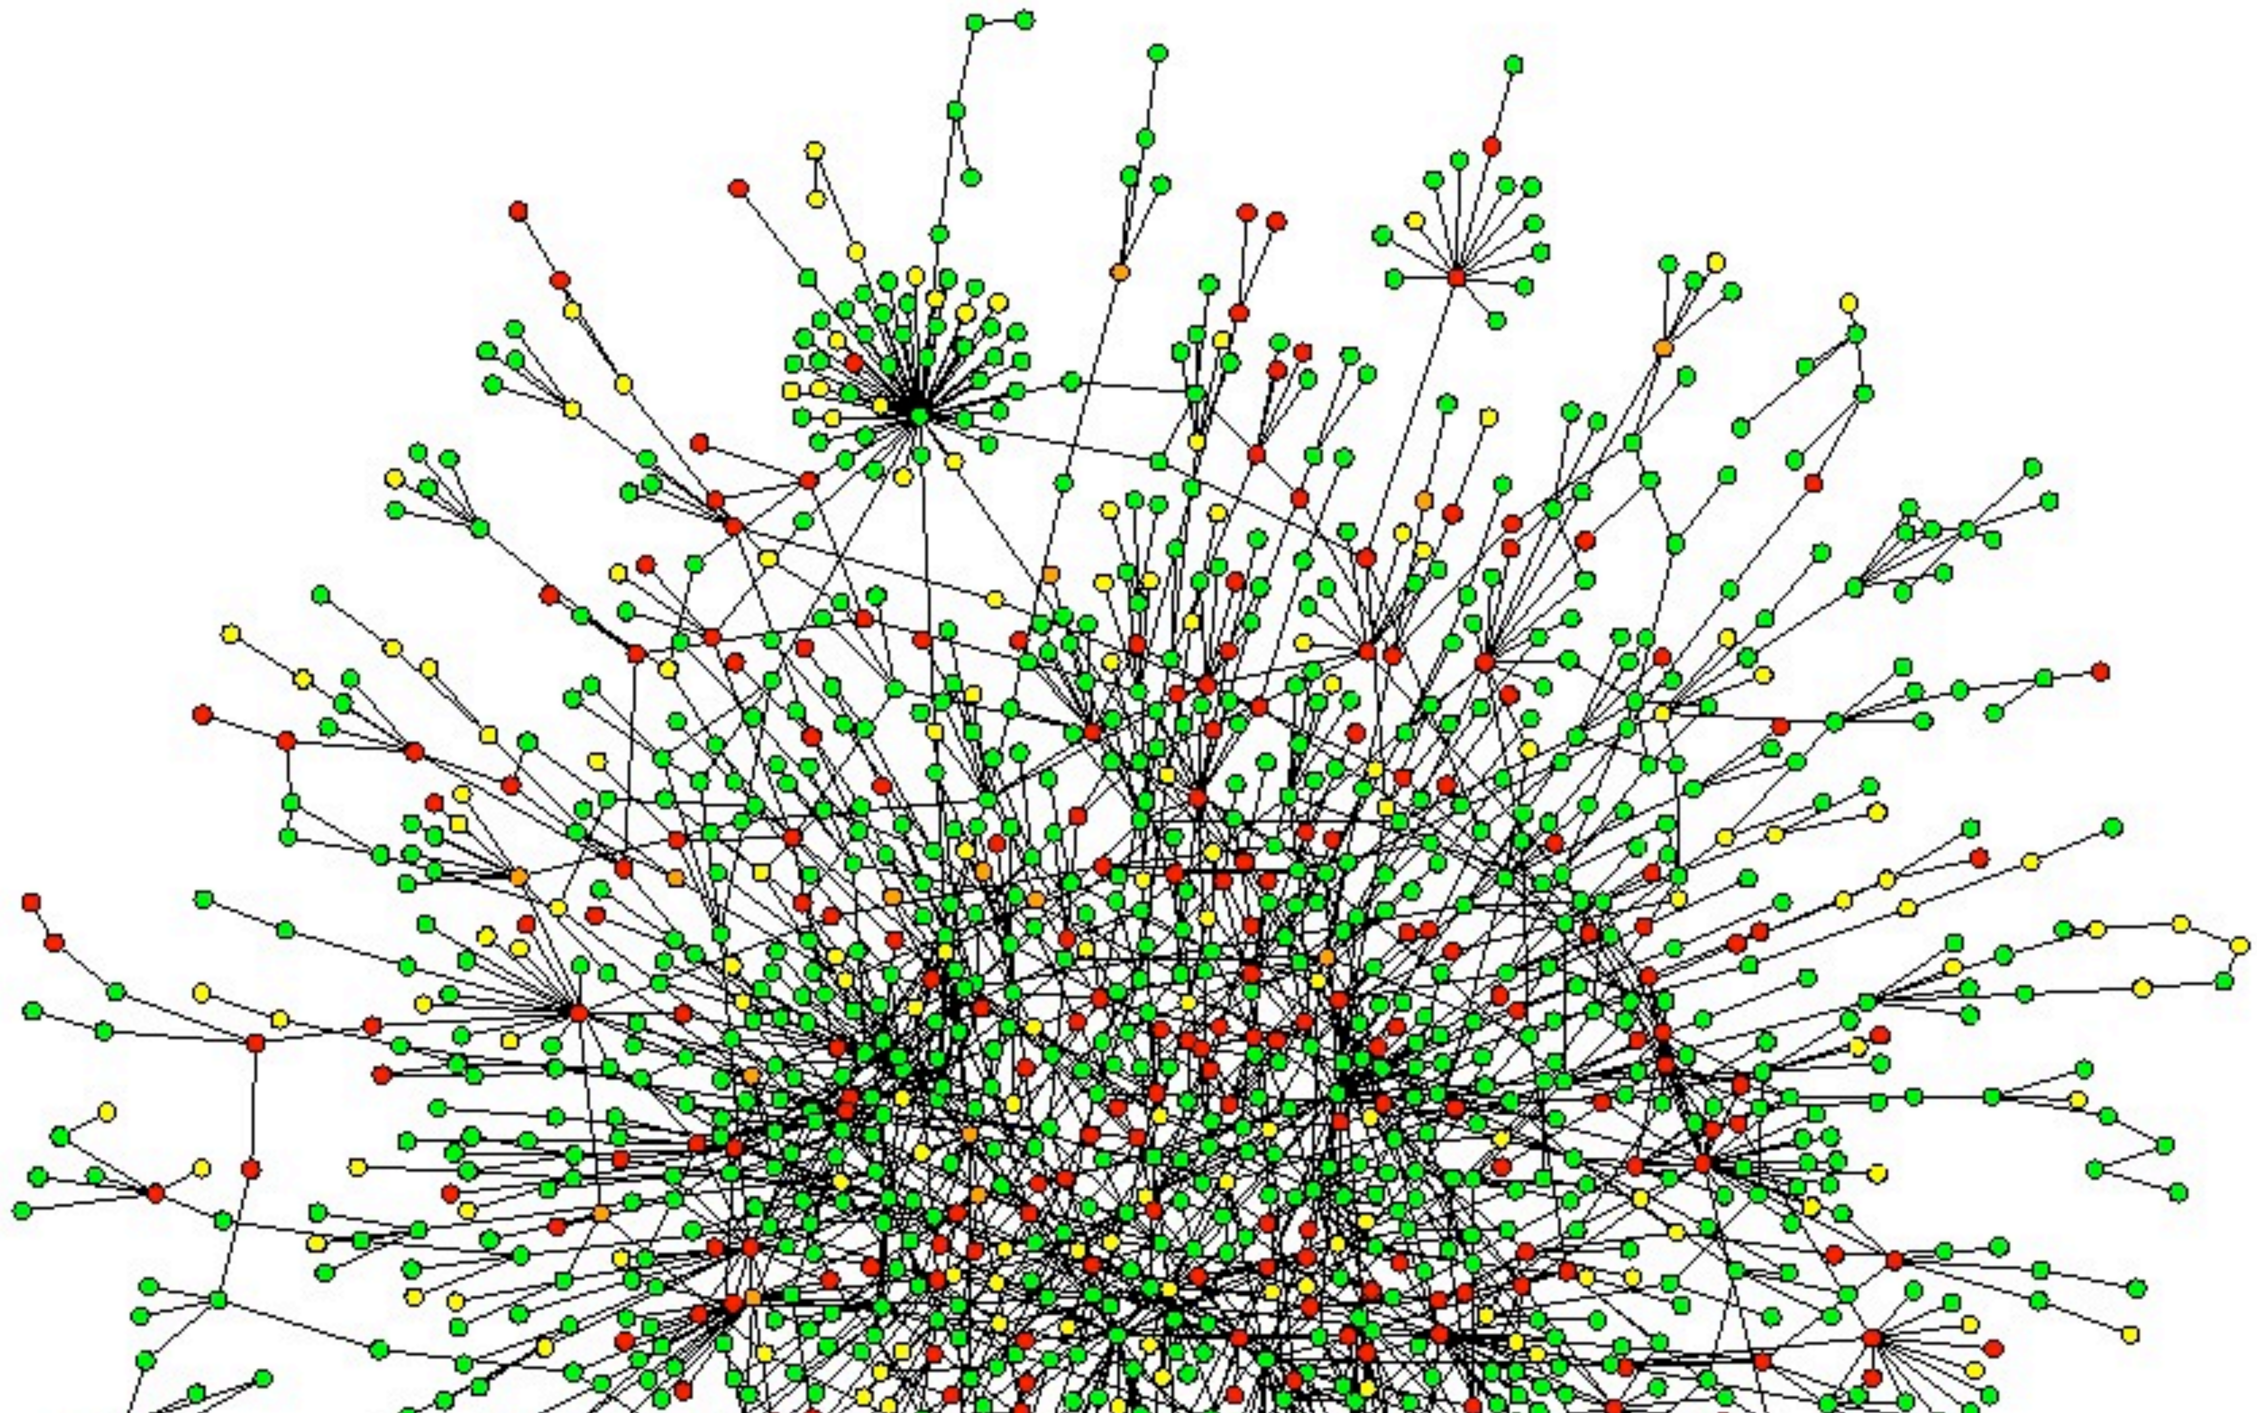

Protein interaction network

Distributed estimation for sensor networks

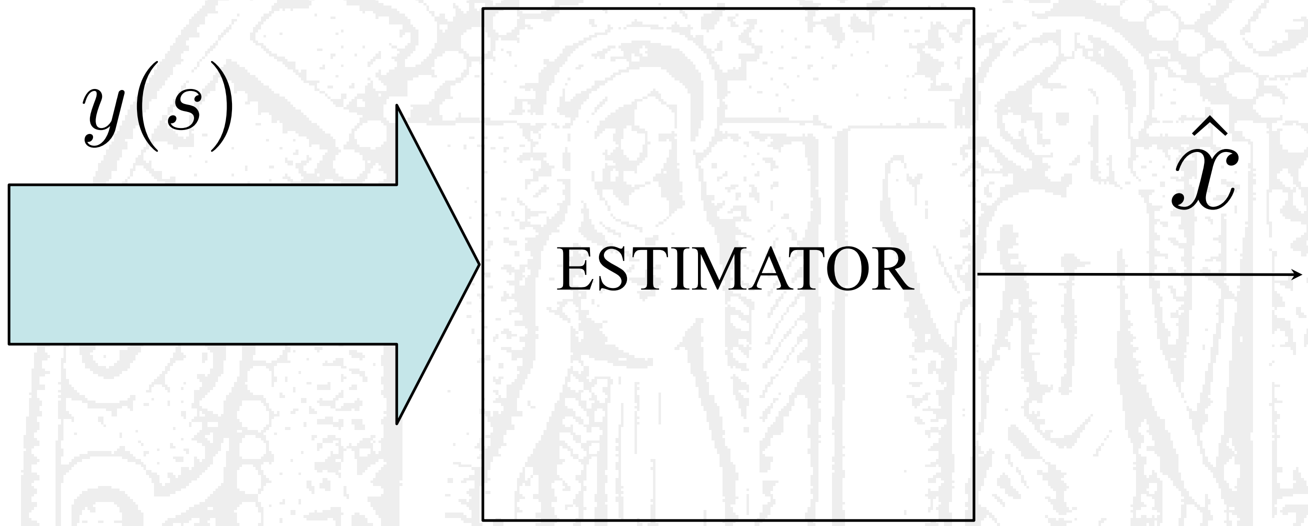

s = space variable
$y(s)$ = spatial data
$\hat{x}$ = data based decision

Distributed estimation for sensor networks

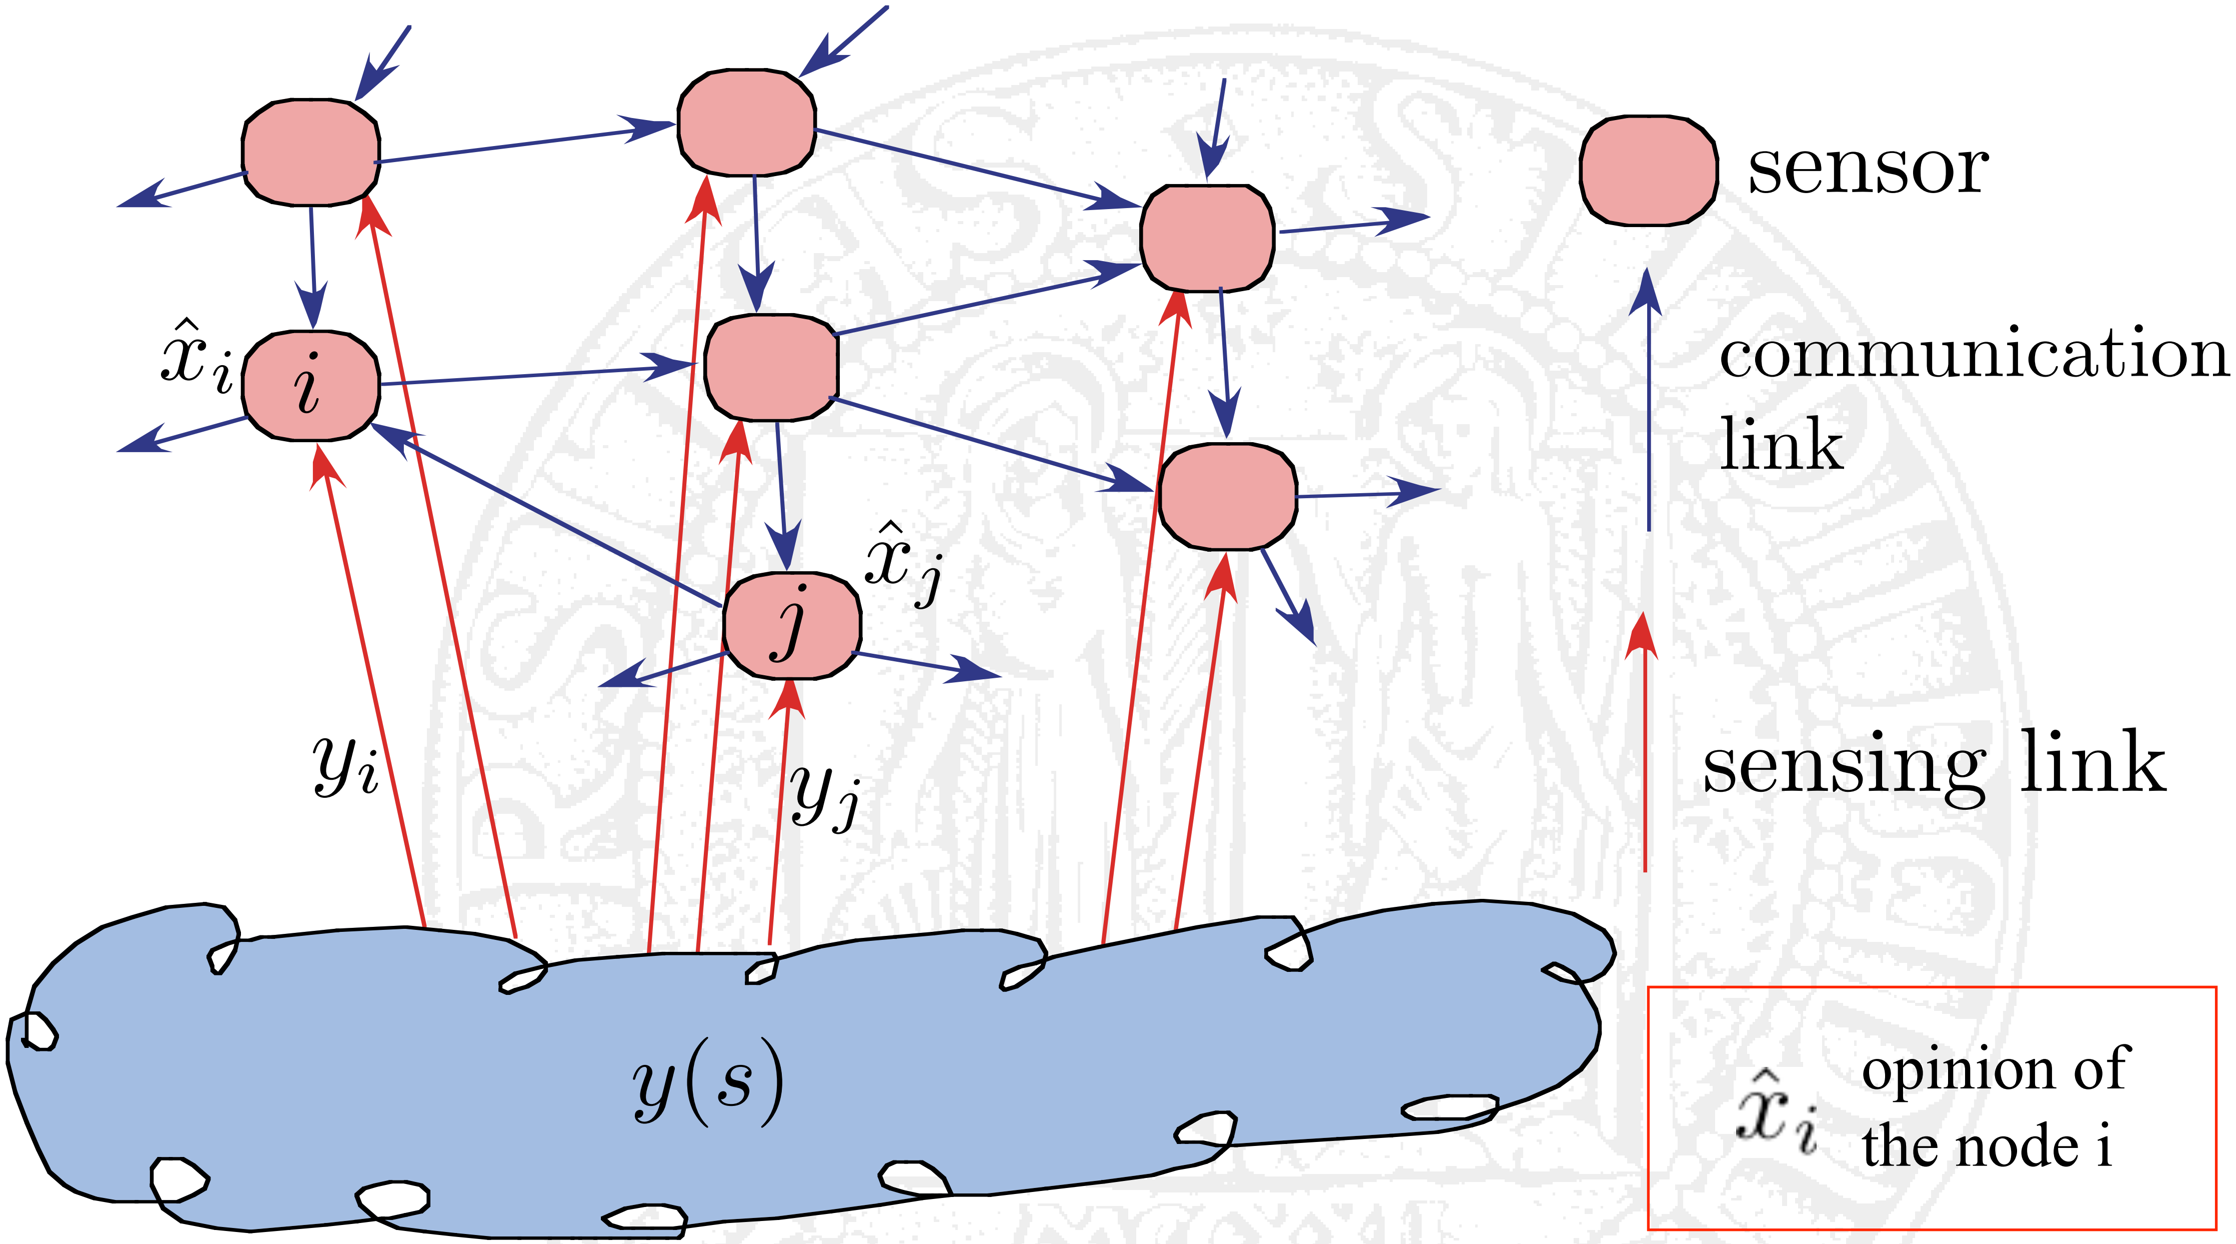

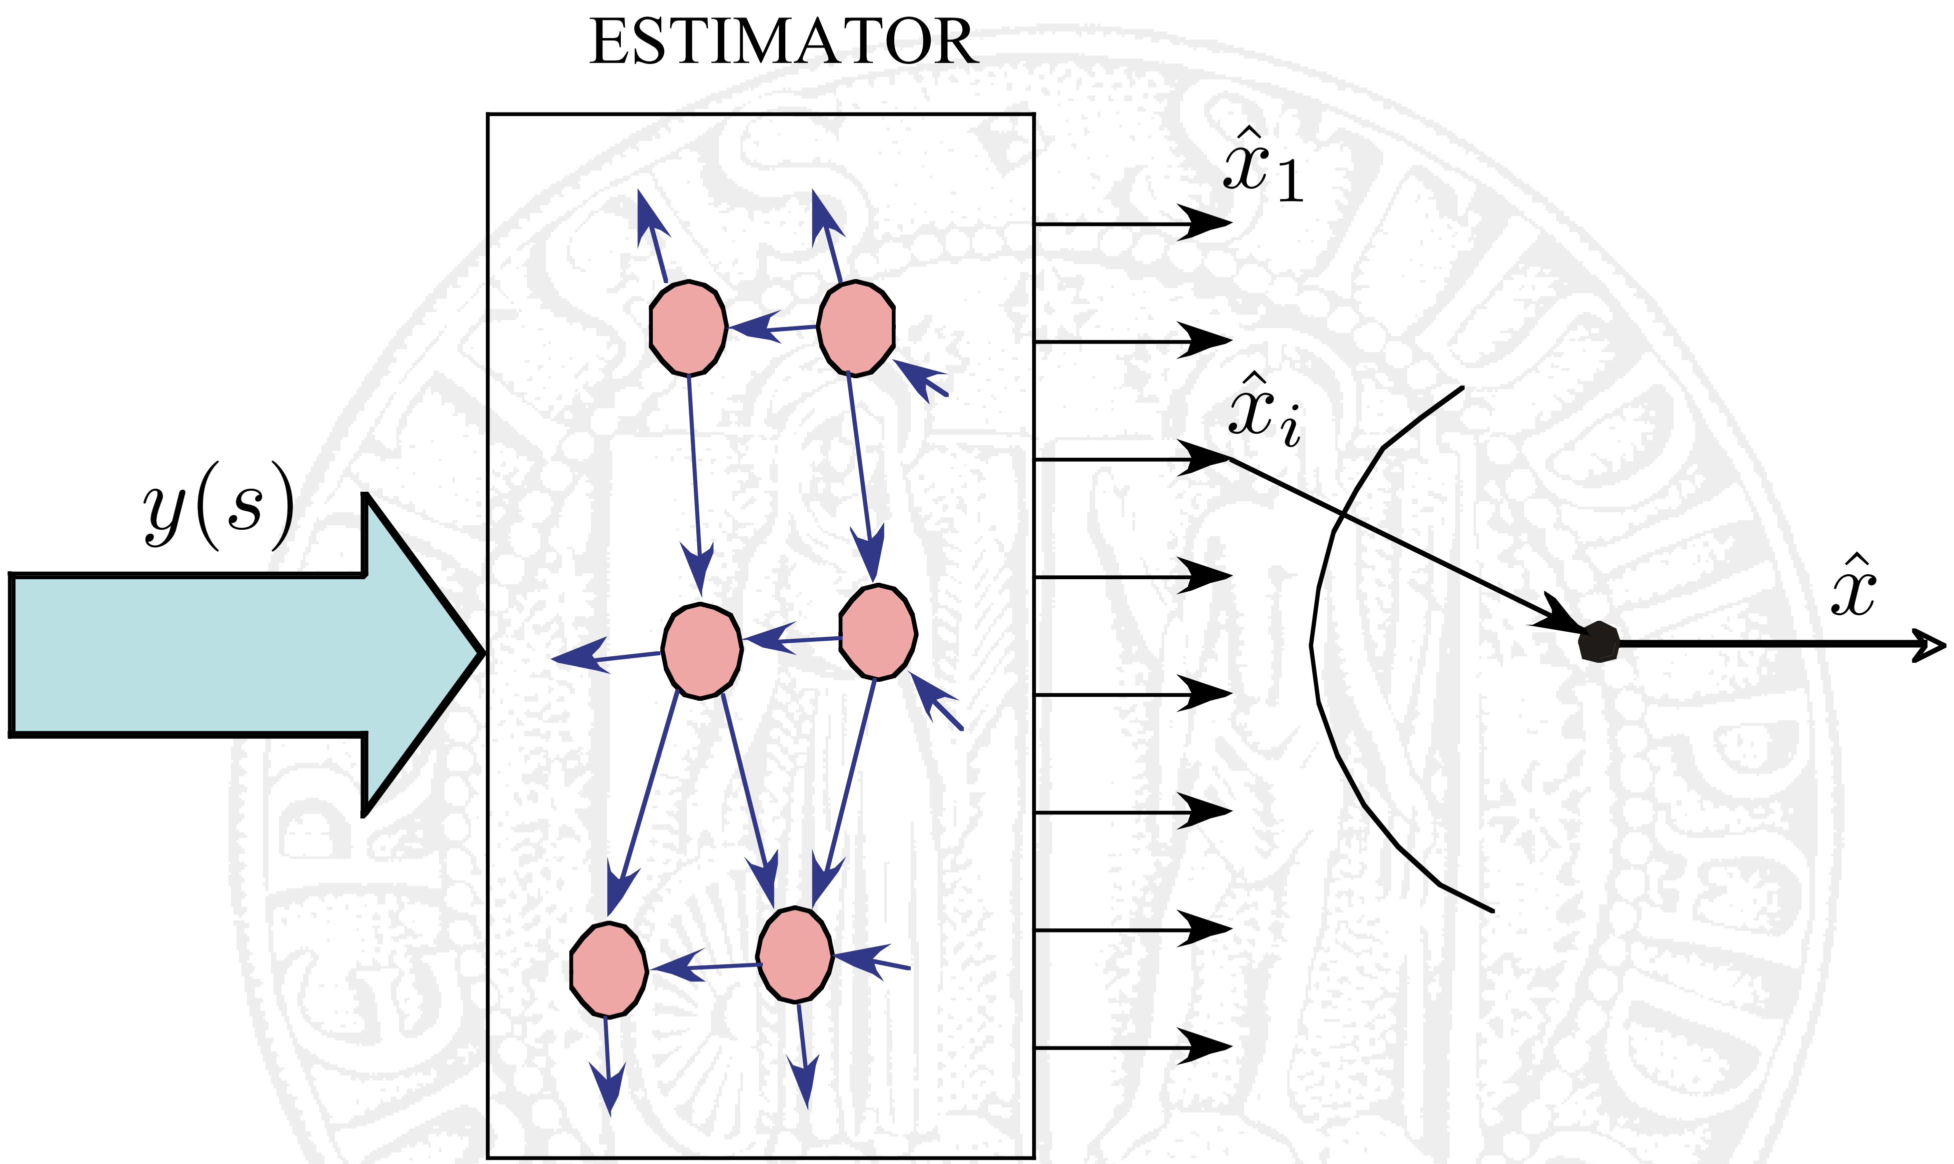

Advantages: intrinsic robustness and adaptivity due to redundancy

Distributed estimation for sensor networks

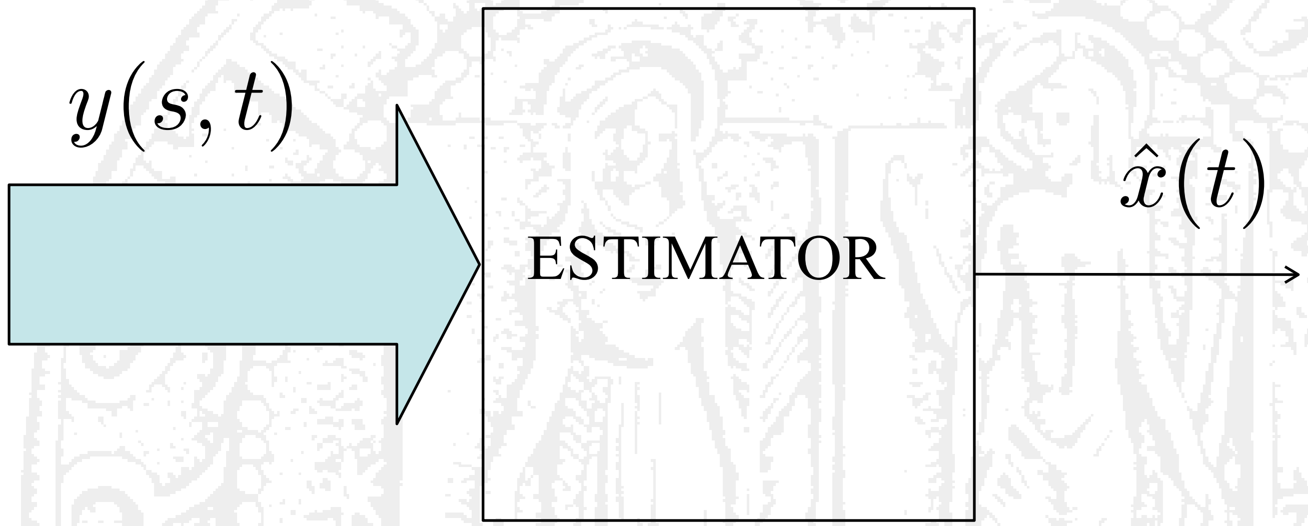

t = time
s = space variable
$y(s,t)$ = time-varying spatial data
$\hat{x}(t)$ = time-varying data base decision

Distributed control

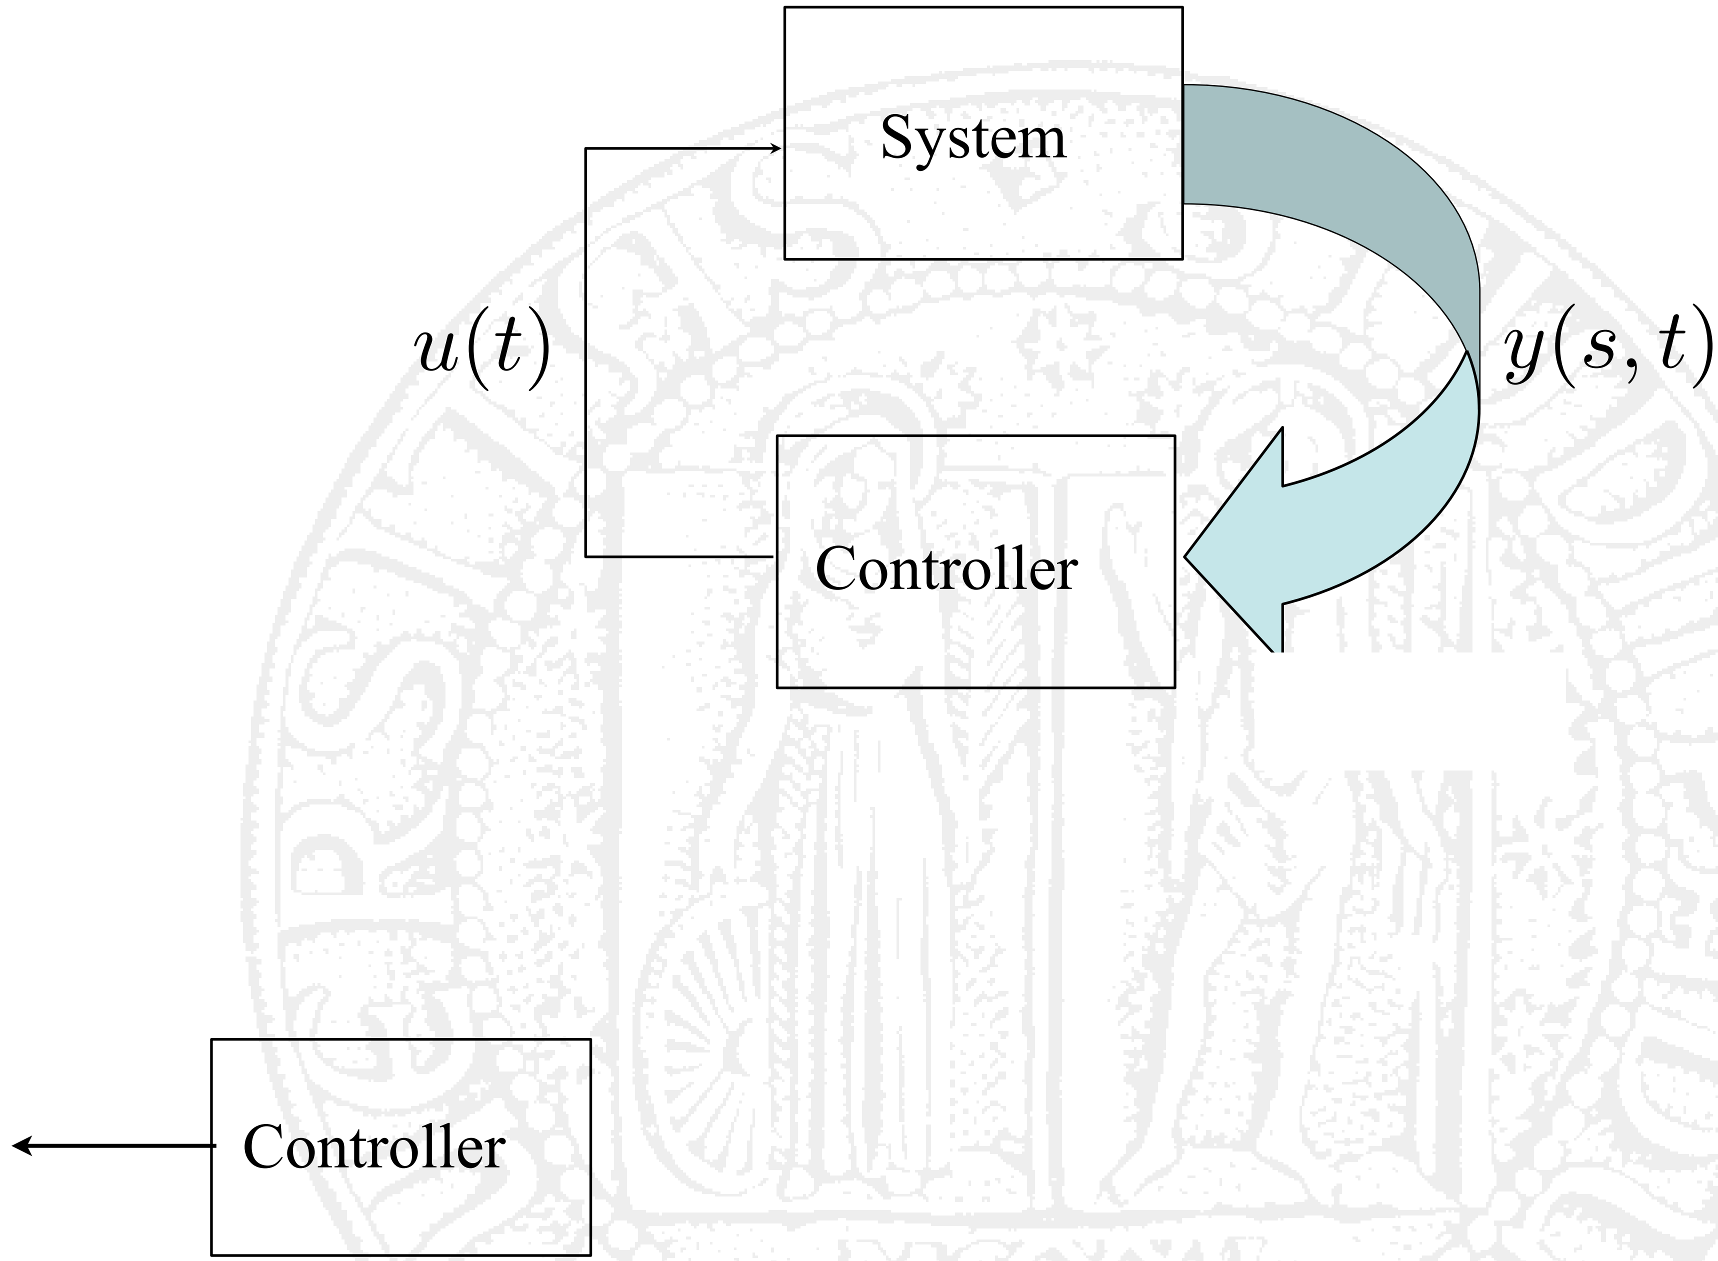

ESTIMATOR

Consensus algorithm

GOAL: each node has to obtain the average of the N values $y_1, \ldots, y_N$ where y_i is known only by the node i. This task has to be performed in a distributed way.

ALGORITHM: Each sensor produces at time t an estimate $x_i(t)$ of the average as follows

$$x_i(t+1) = \sum_{j=1}^{N} P_{ij} x_j(t) \qquad x_i(0) = y_i$$

COMMUNICATION: $x_j(t)$ needs to be transmitted from the node i to the node j iff

$$P_{ij} \neq 0$$

Consensus algorithm

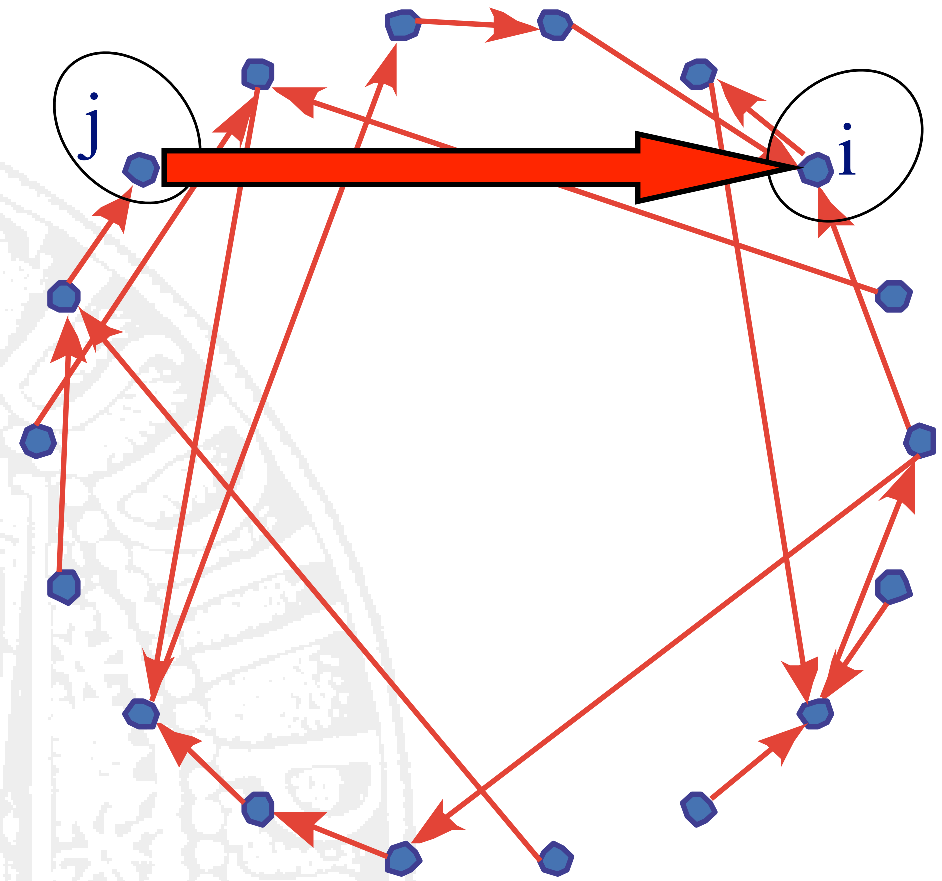

$$x(t+1) = Px(t)$$
$$x(0) = y$$

If the graph associated with $\mathcal{G}_P$ associated with P is strongly connected, then all estimates converge to the same value (consensus)

$$x_i(t) \longrightarrow \sum_{j=1}^{N} \mu_j x_j(0)$$

where the weights μ_j are nonnegative and sum to one.

Consensus algorithm

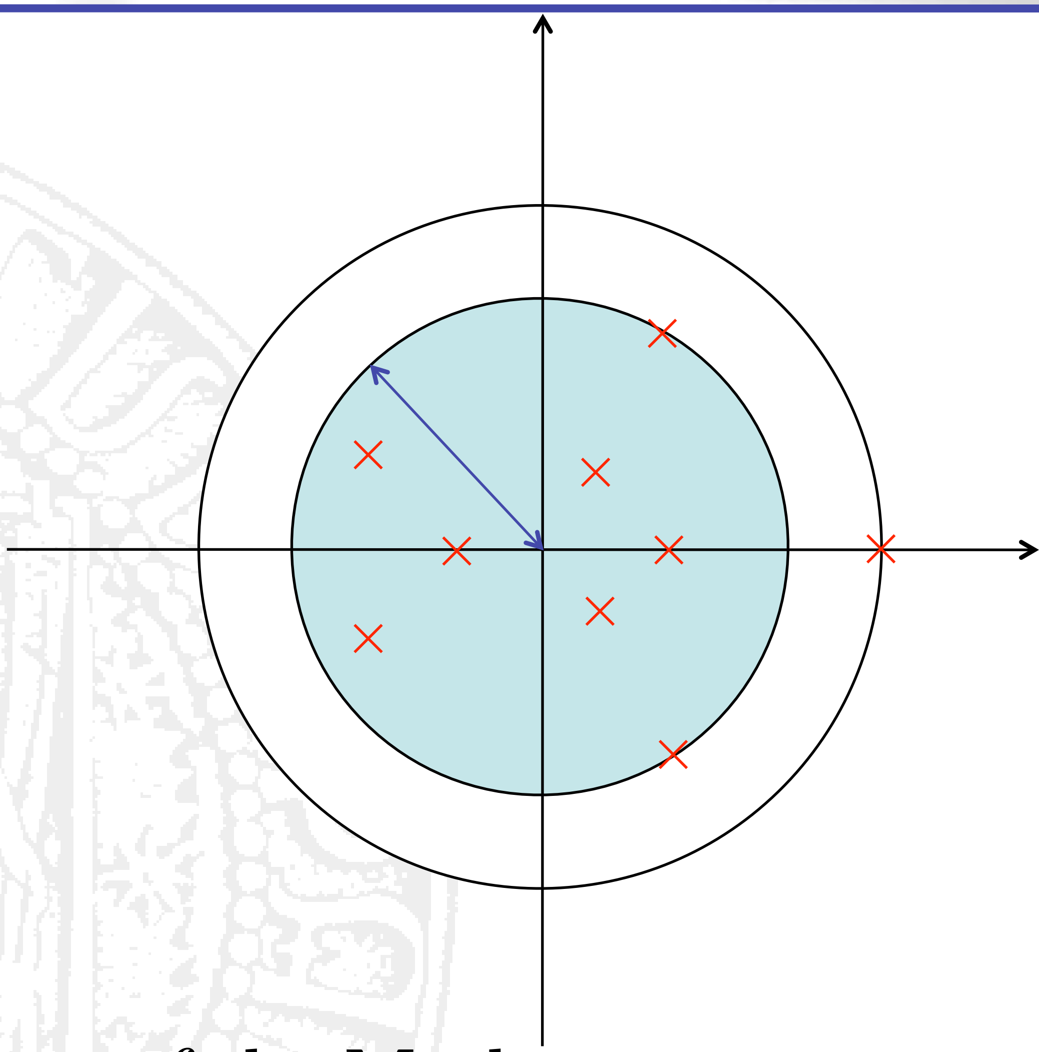

PERFORMANCE INDICES

1. The difference between μ_j and $1/N$.

2. Speed of convergence of $x_i(t)$ to $x_i(\infty)$

MARKOV CHAINS THEORY

1. The vector $(\mu_1, \ldots, \mu_N)$ is the invariant measure of the Markov chain. Therefore $\mu_j = 1/N$ if and only if P is doubly stochastic.

2. The convergence is exponential with rate given by the second largest eigenvalue ρ of P. The number $1-\rho$ is called the spectral gap of P.

Example: vehicles formation

Assume we have N vehicles moving on the plane. Each vehicle has coordinates $z_i(t) = (x_i(t), y_i(t))^T$. The goal is the rendezvous of the vehicles in one point of the plane (can be generalized to formation reaching).

Solution:

$$z_i(t+1) = \sum_{i=1}^{N} P_{ij} z_j(t)$$

The vehicles will reach asymptotically the centroid on the initial positions.

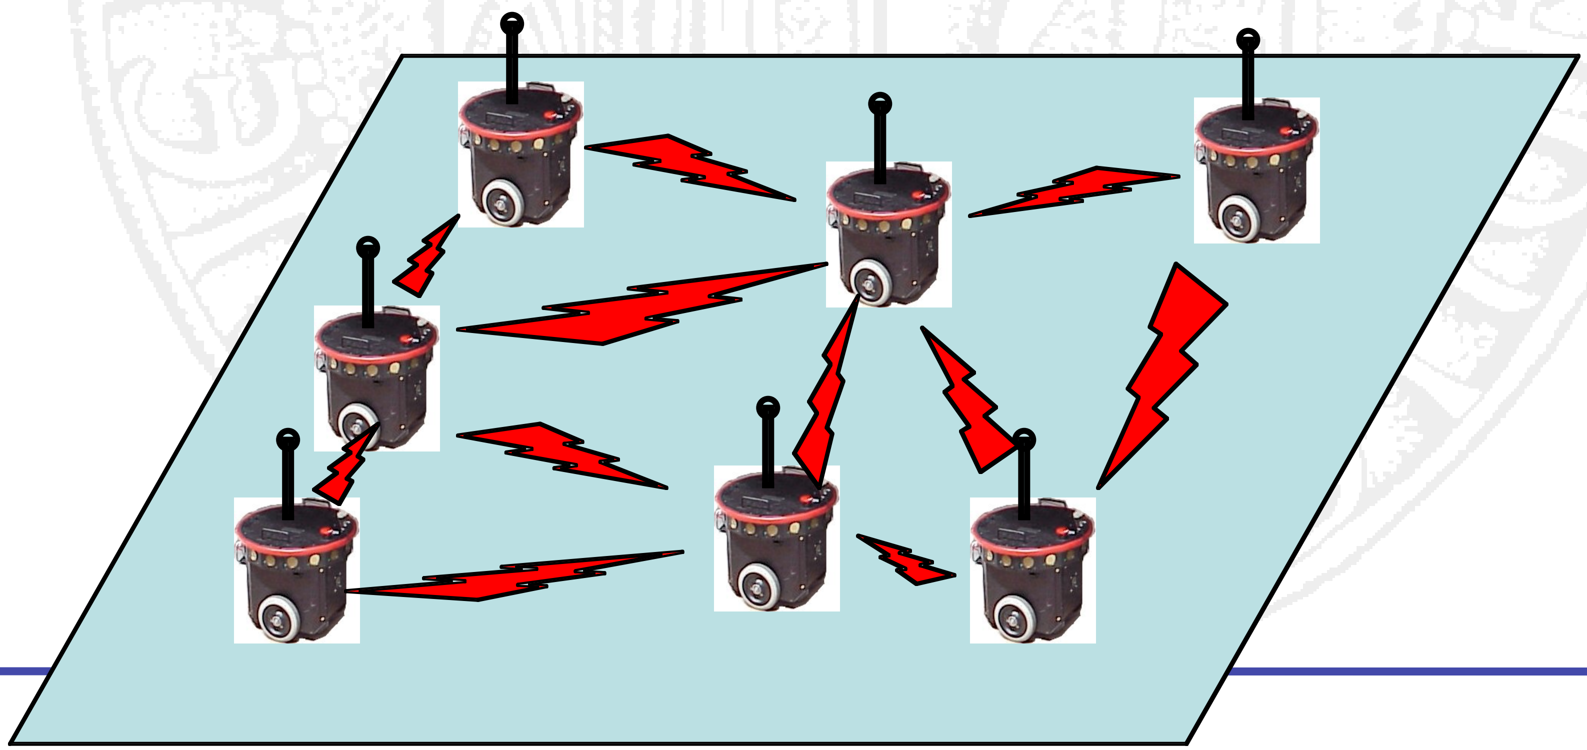

Example: distributed estimation

Assume that N sensors have to estimate a quantity $x \in \mathbf{R}$ from their noisy measurements. The result of the measure of the sensor i is

$$y_i = x + v_i$$

where v_i are independent noises of zero mean and variance r. The best estimate of x from the measurements is

$$\hat{x} := \frac{1}{N} \sum_i y_i$$

Example: distributed least square

Assume that each sensor i measures two variables x_i, y_i and that the relation between these needs to be estimated. The relation is modeled by a finite dimensional function space

$$f(x) = \sum_{i=1}^{n} \theta_i f_i(x)$$

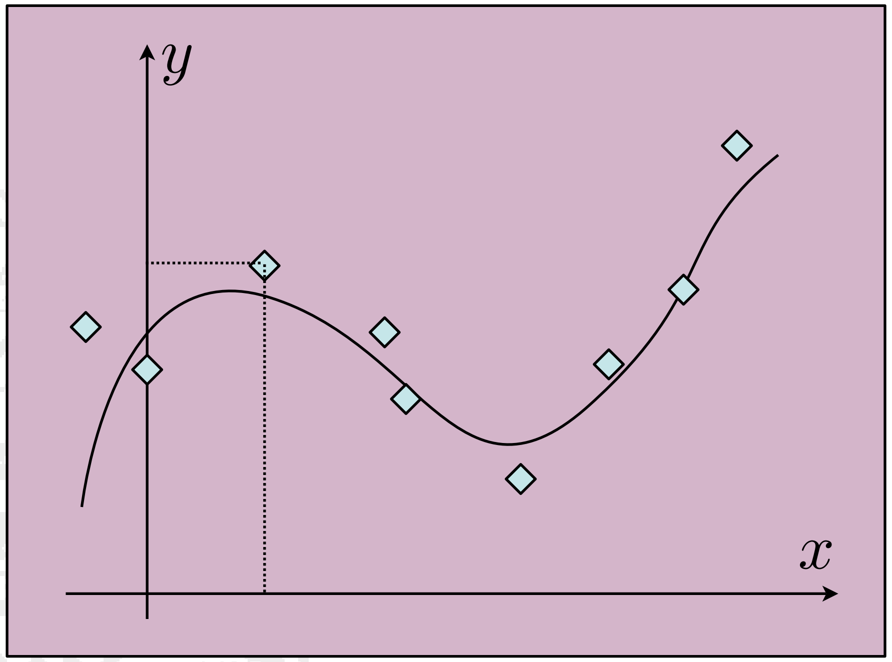

where the functions $f_i(x)$ form the basis of the function space. We need to estimate the coefficients θ_i. We can write

$$f(x) = F^T(x)\Theta$$

where

$$F^T(x) = [f_1(x) \ \cdots \ f_n(x)] \qquad \Theta = [\theta_1 \ \cdots \ \theta_n]^T$$

PROBLEM: Determine

$$\hat{\Theta} := \text{argmin}_{\Theta} \sum_{j=1}^{N} (y_i - F^T(x_i)\Theta)^2$$

Example: distributed least square

According the theory of least square optimization we have that

$$\hat{\Theta} = \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i) \right)^{-1} \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i \right)$$

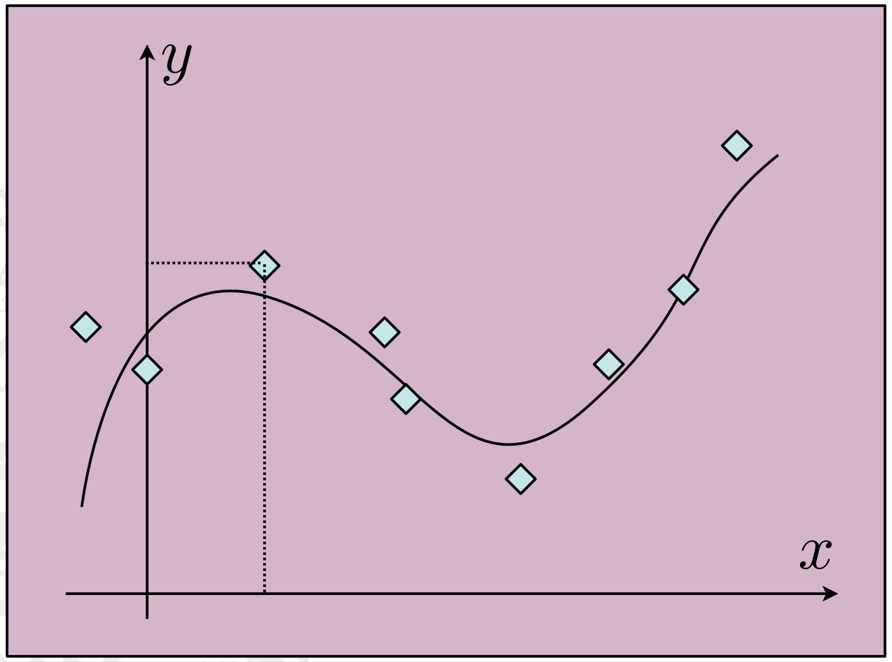

SOLUTION

$$M_i(0) = F(x_i) F^T(x_i) \in \mathbb{R}^{n \times n} \xrightarrow{\text{average consensus}} M_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i)$$

$$v_i(0) = F(x_i) y_i \in \mathbb{R}^{n} \xrightarrow{\text{average consensus}} v_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i$$

Example: distributed least square

According the theory of least square optimization we have that

$$\hat{\Theta} = \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i) \right)^{-1} \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i \right)$$

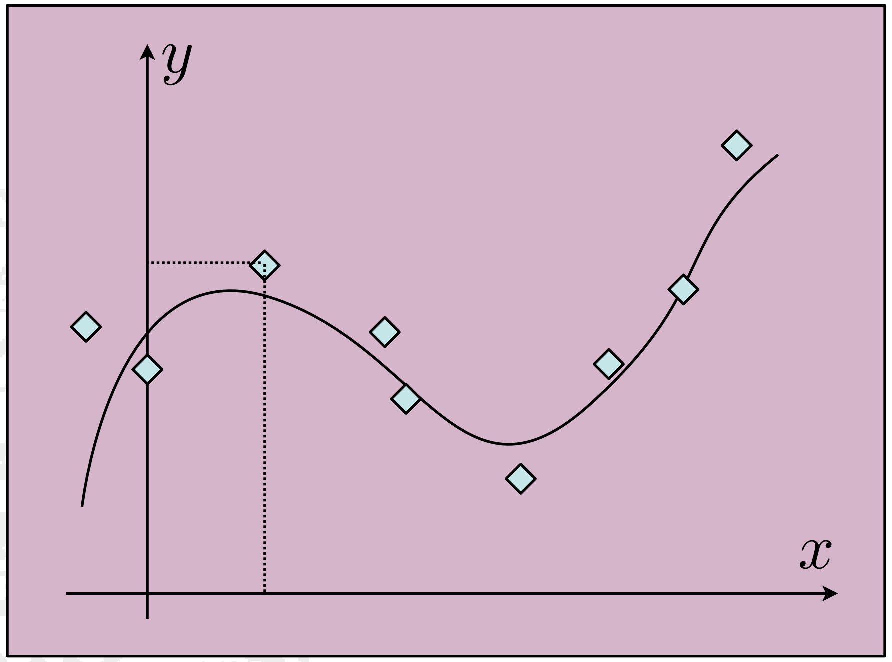

SOLUTION

$$M_i(0) = F(x_i) F^T(x_i) \in \mathbb{R}^{n \times n} \xrightarrow{\text{average consensus}} M_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i)$$

Initial knowledge of the node i

$$v_i(0) = F(x_i) y_i \in \mathbb{R}^n \xrightarrow{\text{average consensus}} v_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i$$

Example: distributed least square

According the theory of least square optimization we have that

$$\hat{\Theta} = \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i) \right)^{-1} \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i \right)$$

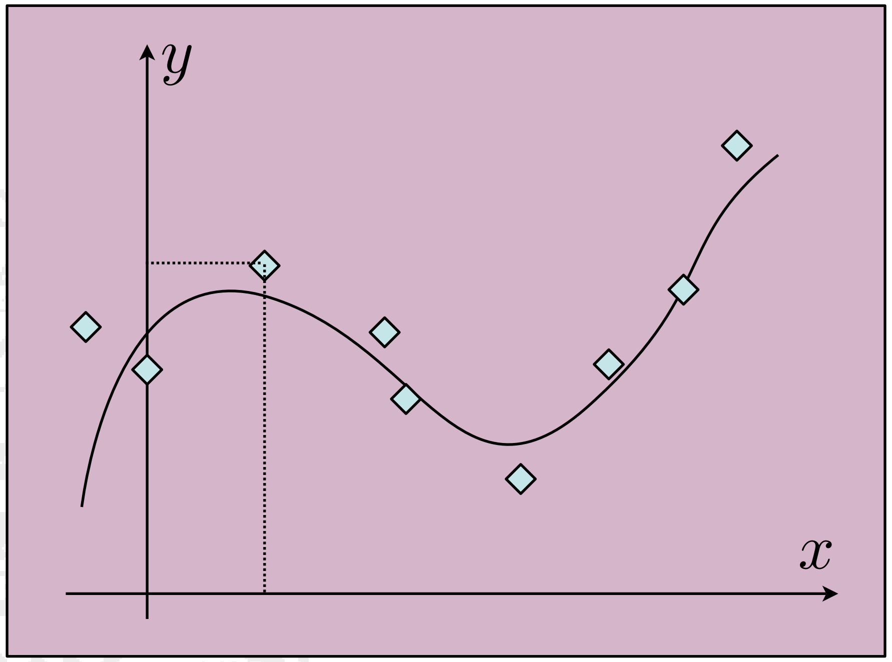

SOLUTION

$$M_i(0) = F(x_i)F^T(x_i) \in \mathbb{R}^{n \times n} \xrightarrow{\text{average consensus}} M_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i)$$

Initial knowledge of the node i

Final knowledge of the node i

$$v_i(0) = F(x_i) y_i \in \mathbb{R}^n \xrightarrow{\text{average consensus}} v_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i$$

Example: distributed least square

According the theory of least square optimization we have that

$$\hat{\Theta} = \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i) \right)^{-1} \left(\frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i \right)$$

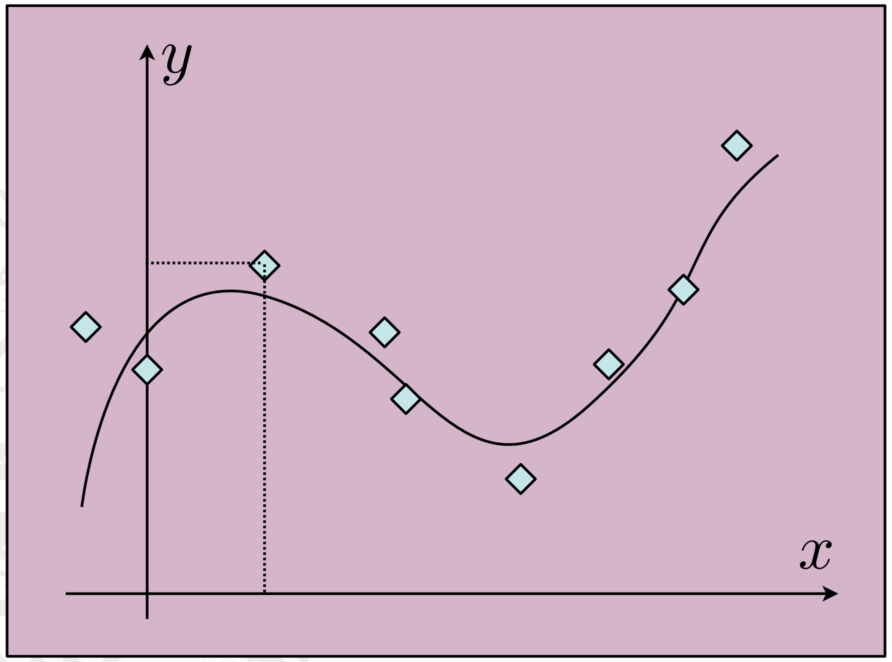

SOLUTION

average consensus

$$M_i(0) = F(x_i) F^T(x_i) \in \mathbb{R}^{n \times n} \longrightarrow M_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) F^T(x_i)$$

Initial knowledge of the node i

Final knowledge of the node i

average consensus

$$v_i(0) = F(x_i) y_i \in \mathbb{R}^n \longrightarrow v_i(\infty) = \frac{1}{N} \sum_{i=1}^{N} F(x_i) y_i$$

Example: distributed decision making

We have a binary random variable x such with prior

$$P(x=0) = P(x=1) = 1/2$$

N sensors can estimate x though a binary random variable y_i which are conditional independent and with conditional probabilities

$$P(y_i = 1|x=0) = P(y_i = 0|x=1) = e_i$$

$$P(y_i = 0|x=0) = P(y_i = 1|x=1) = 1 - e_i$$

It can be seen that the normalized log-likelihood function is

$$\mathcal{L}(y_1, \ldots, y_N) = \frac{1}{N} \log \frac{P(0|y_1, \cdots, y_N)}{P(1|y_1, \cdots, y_N)} = \frac{1}{N} \sum_i (1-2y_i) \log \frac{1-e_i}{e_i}$$

$$\hat{x} = 0 \iff \mathcal{L}(y_1, \ldots, y_N) > 0$$

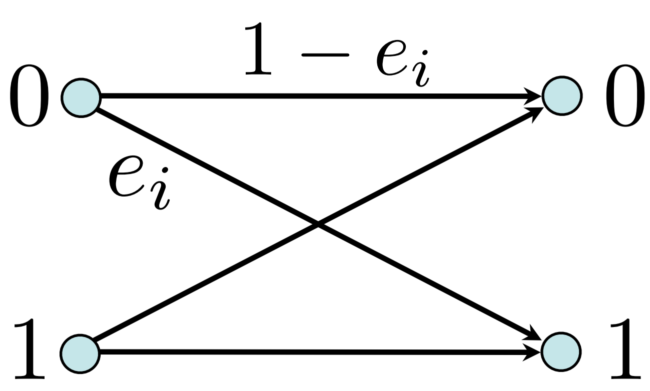

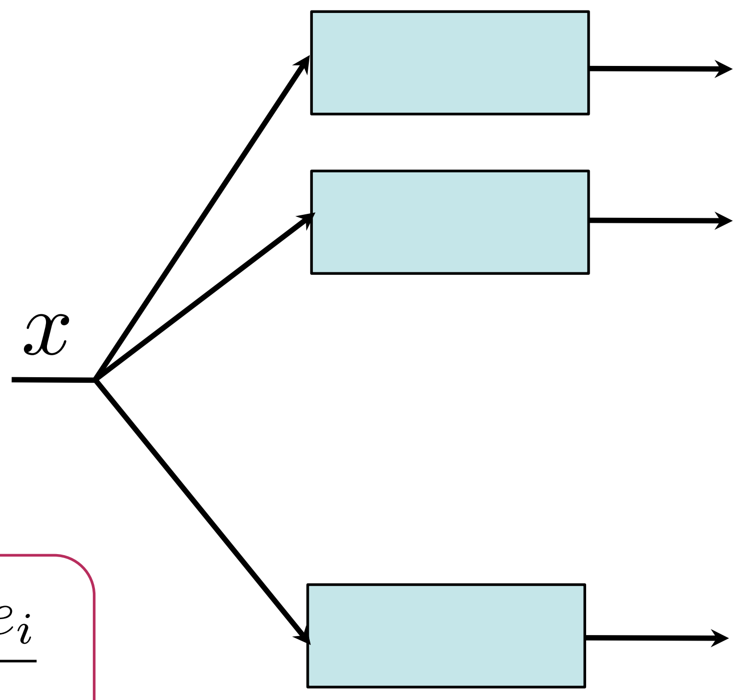

Cayley graphs and Cayley matrices

Let $\mathbf{G}$ be a group with N elements. A matrix P is called a Cayley matrix with respect to $\mathbf{G}$ iff

$$P_{i+l,j+l} = P_{i,j} \qquad \forall i, j, l \in \mathbf{G}$$

A grapf if a Cayley graph iff its adjacency matrix is Cayley.

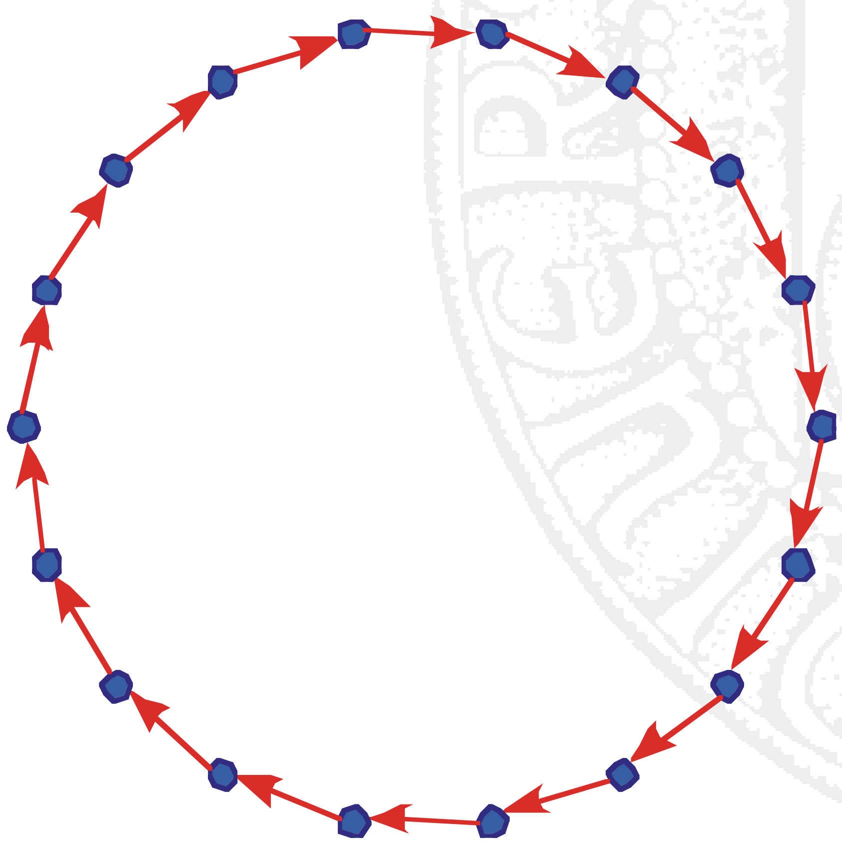

$\mathbf{G} = \mathbf{Z}_N$

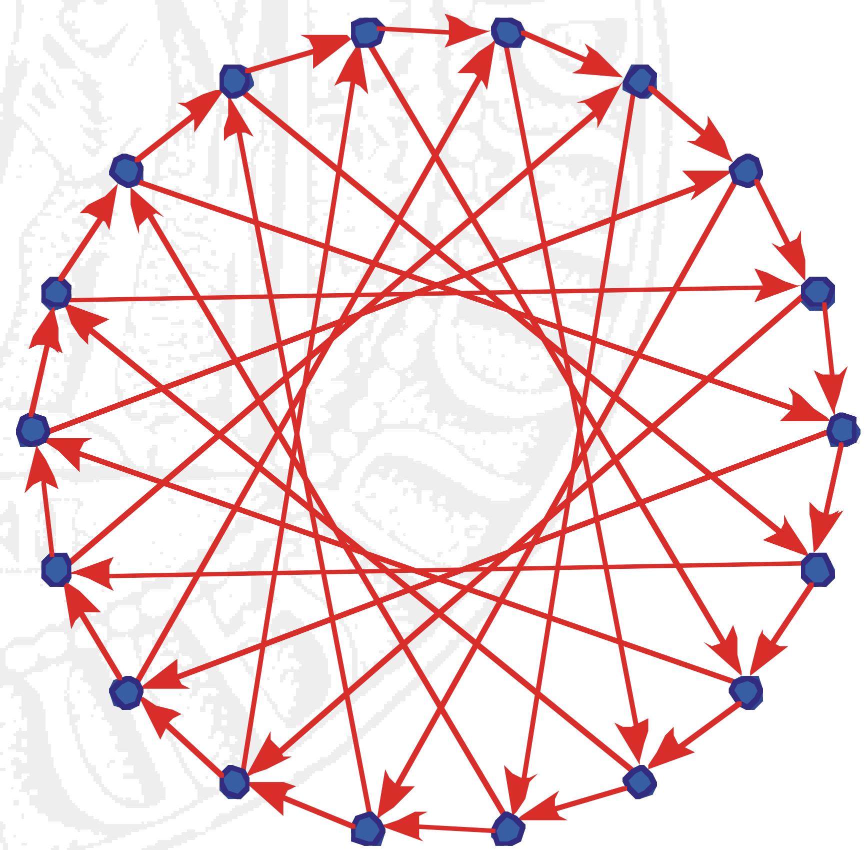

Cayley graphs and Cayley matrices

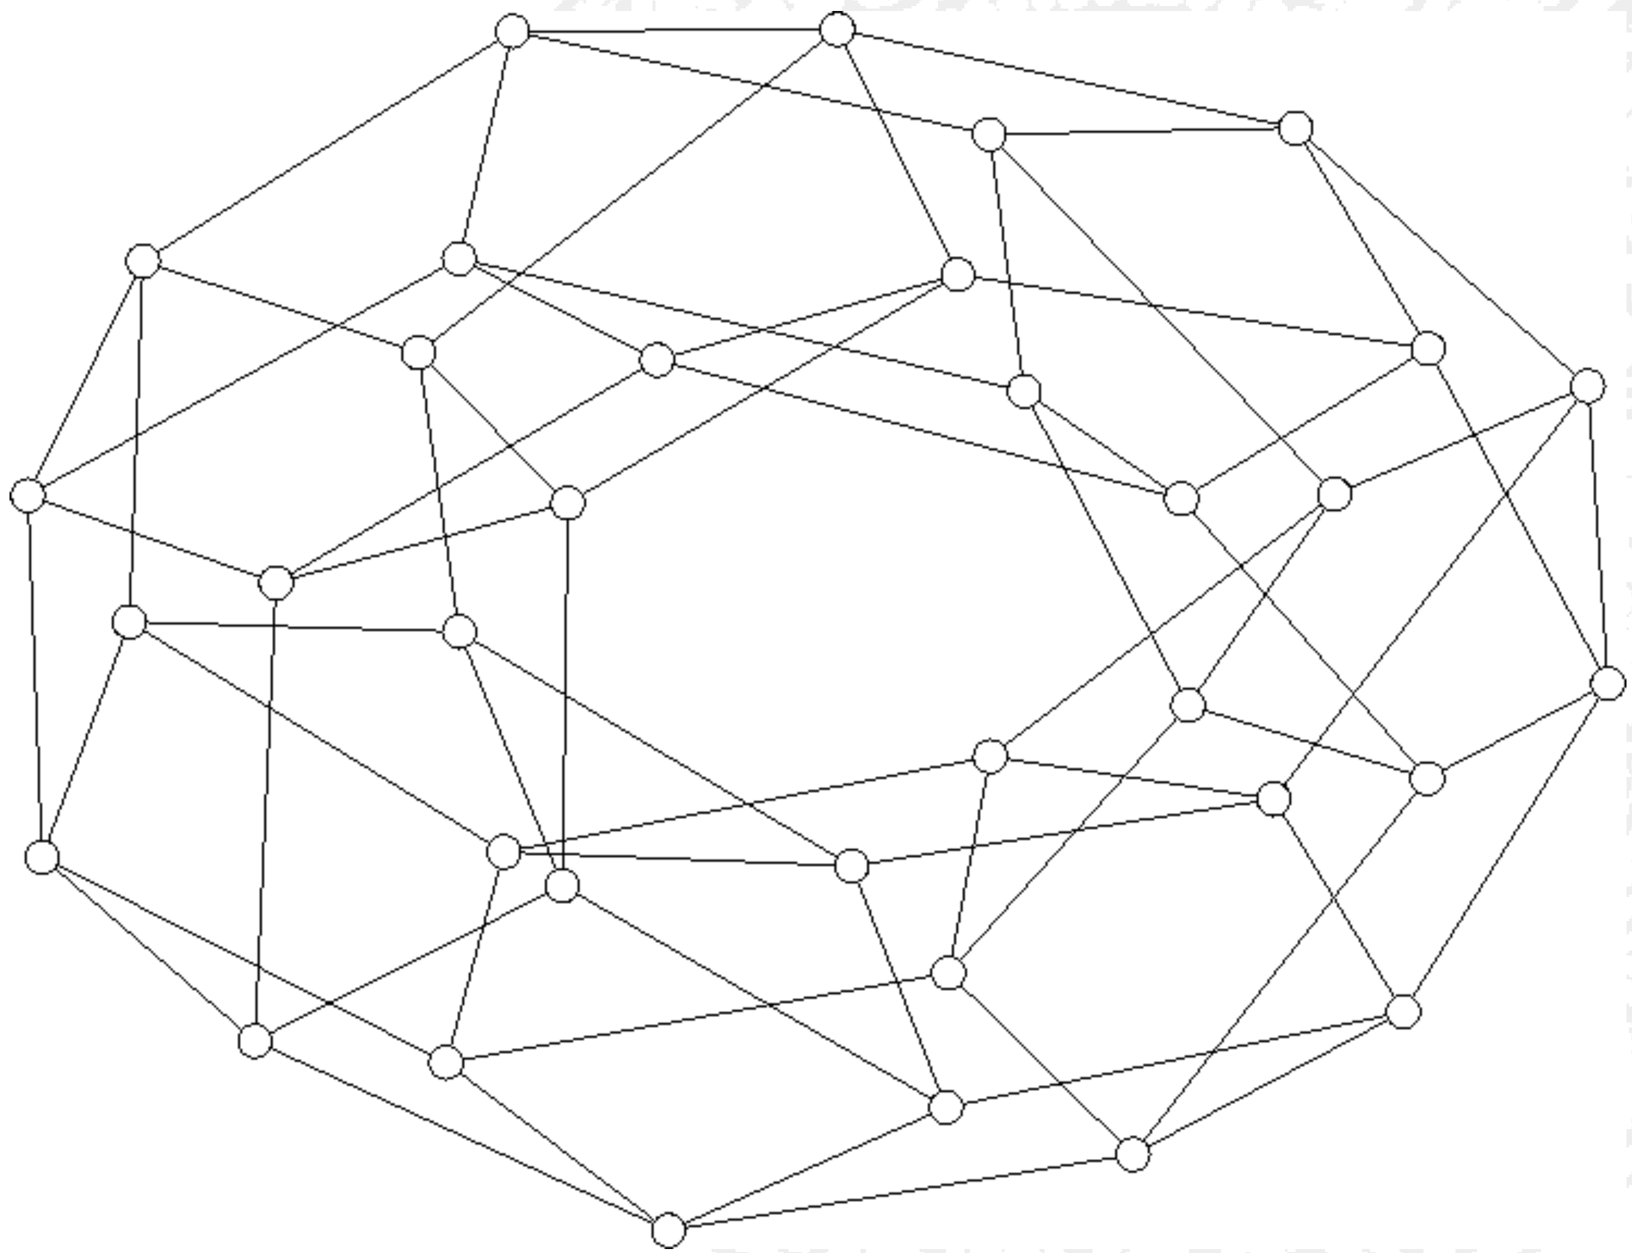

Communication graphs with symmetries

Communication graphs with symmetries

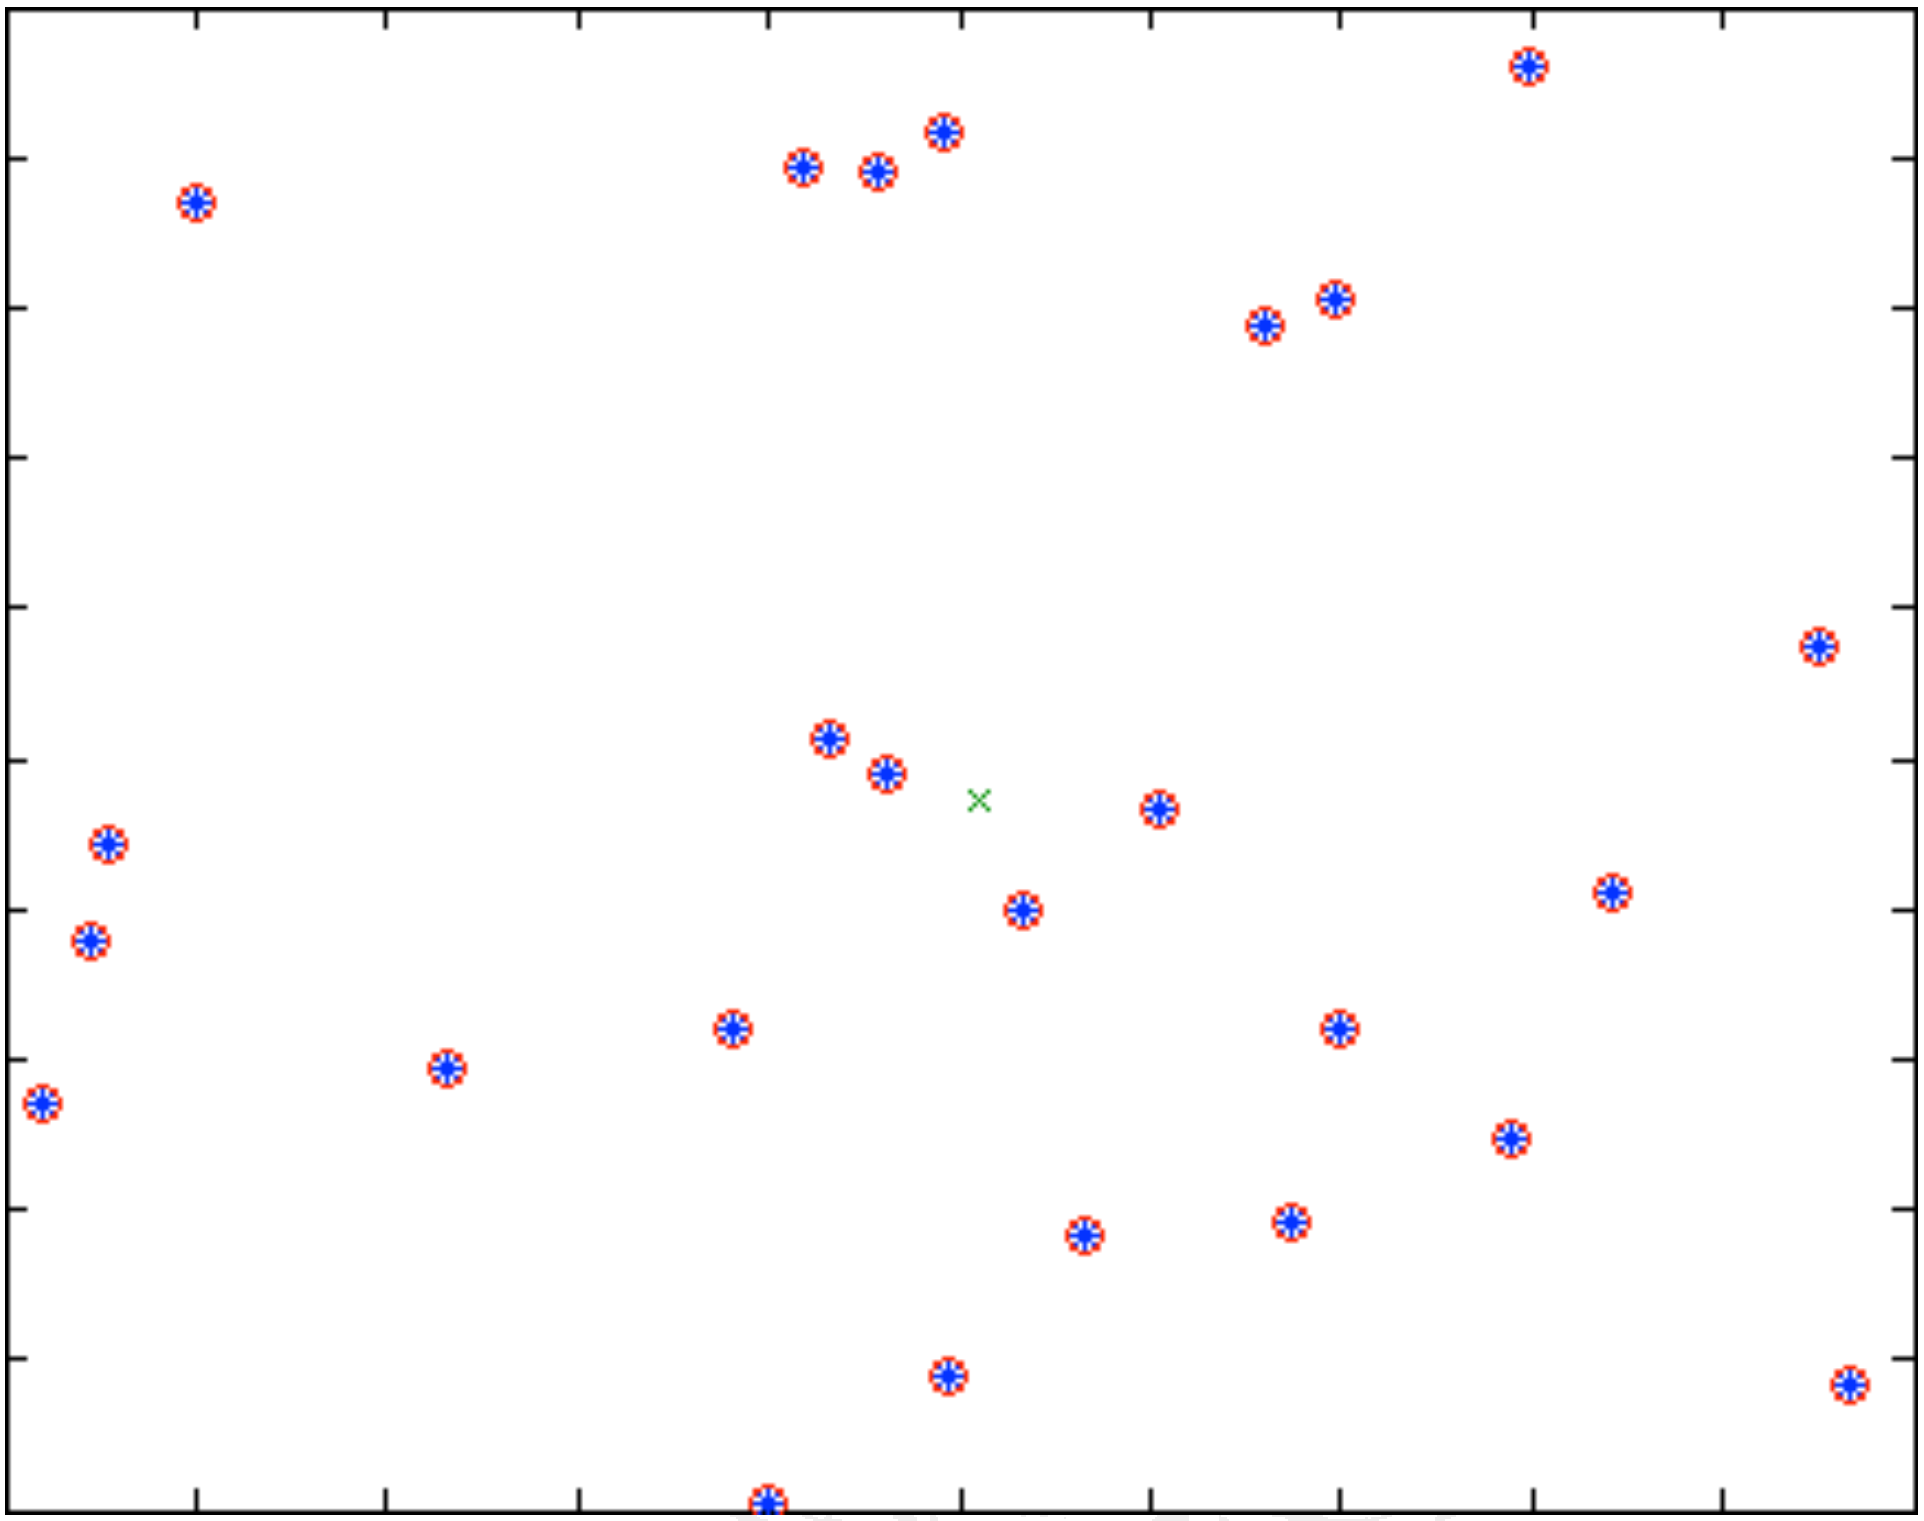

Communication graphs with symmetries

Communication graphs with symmetries

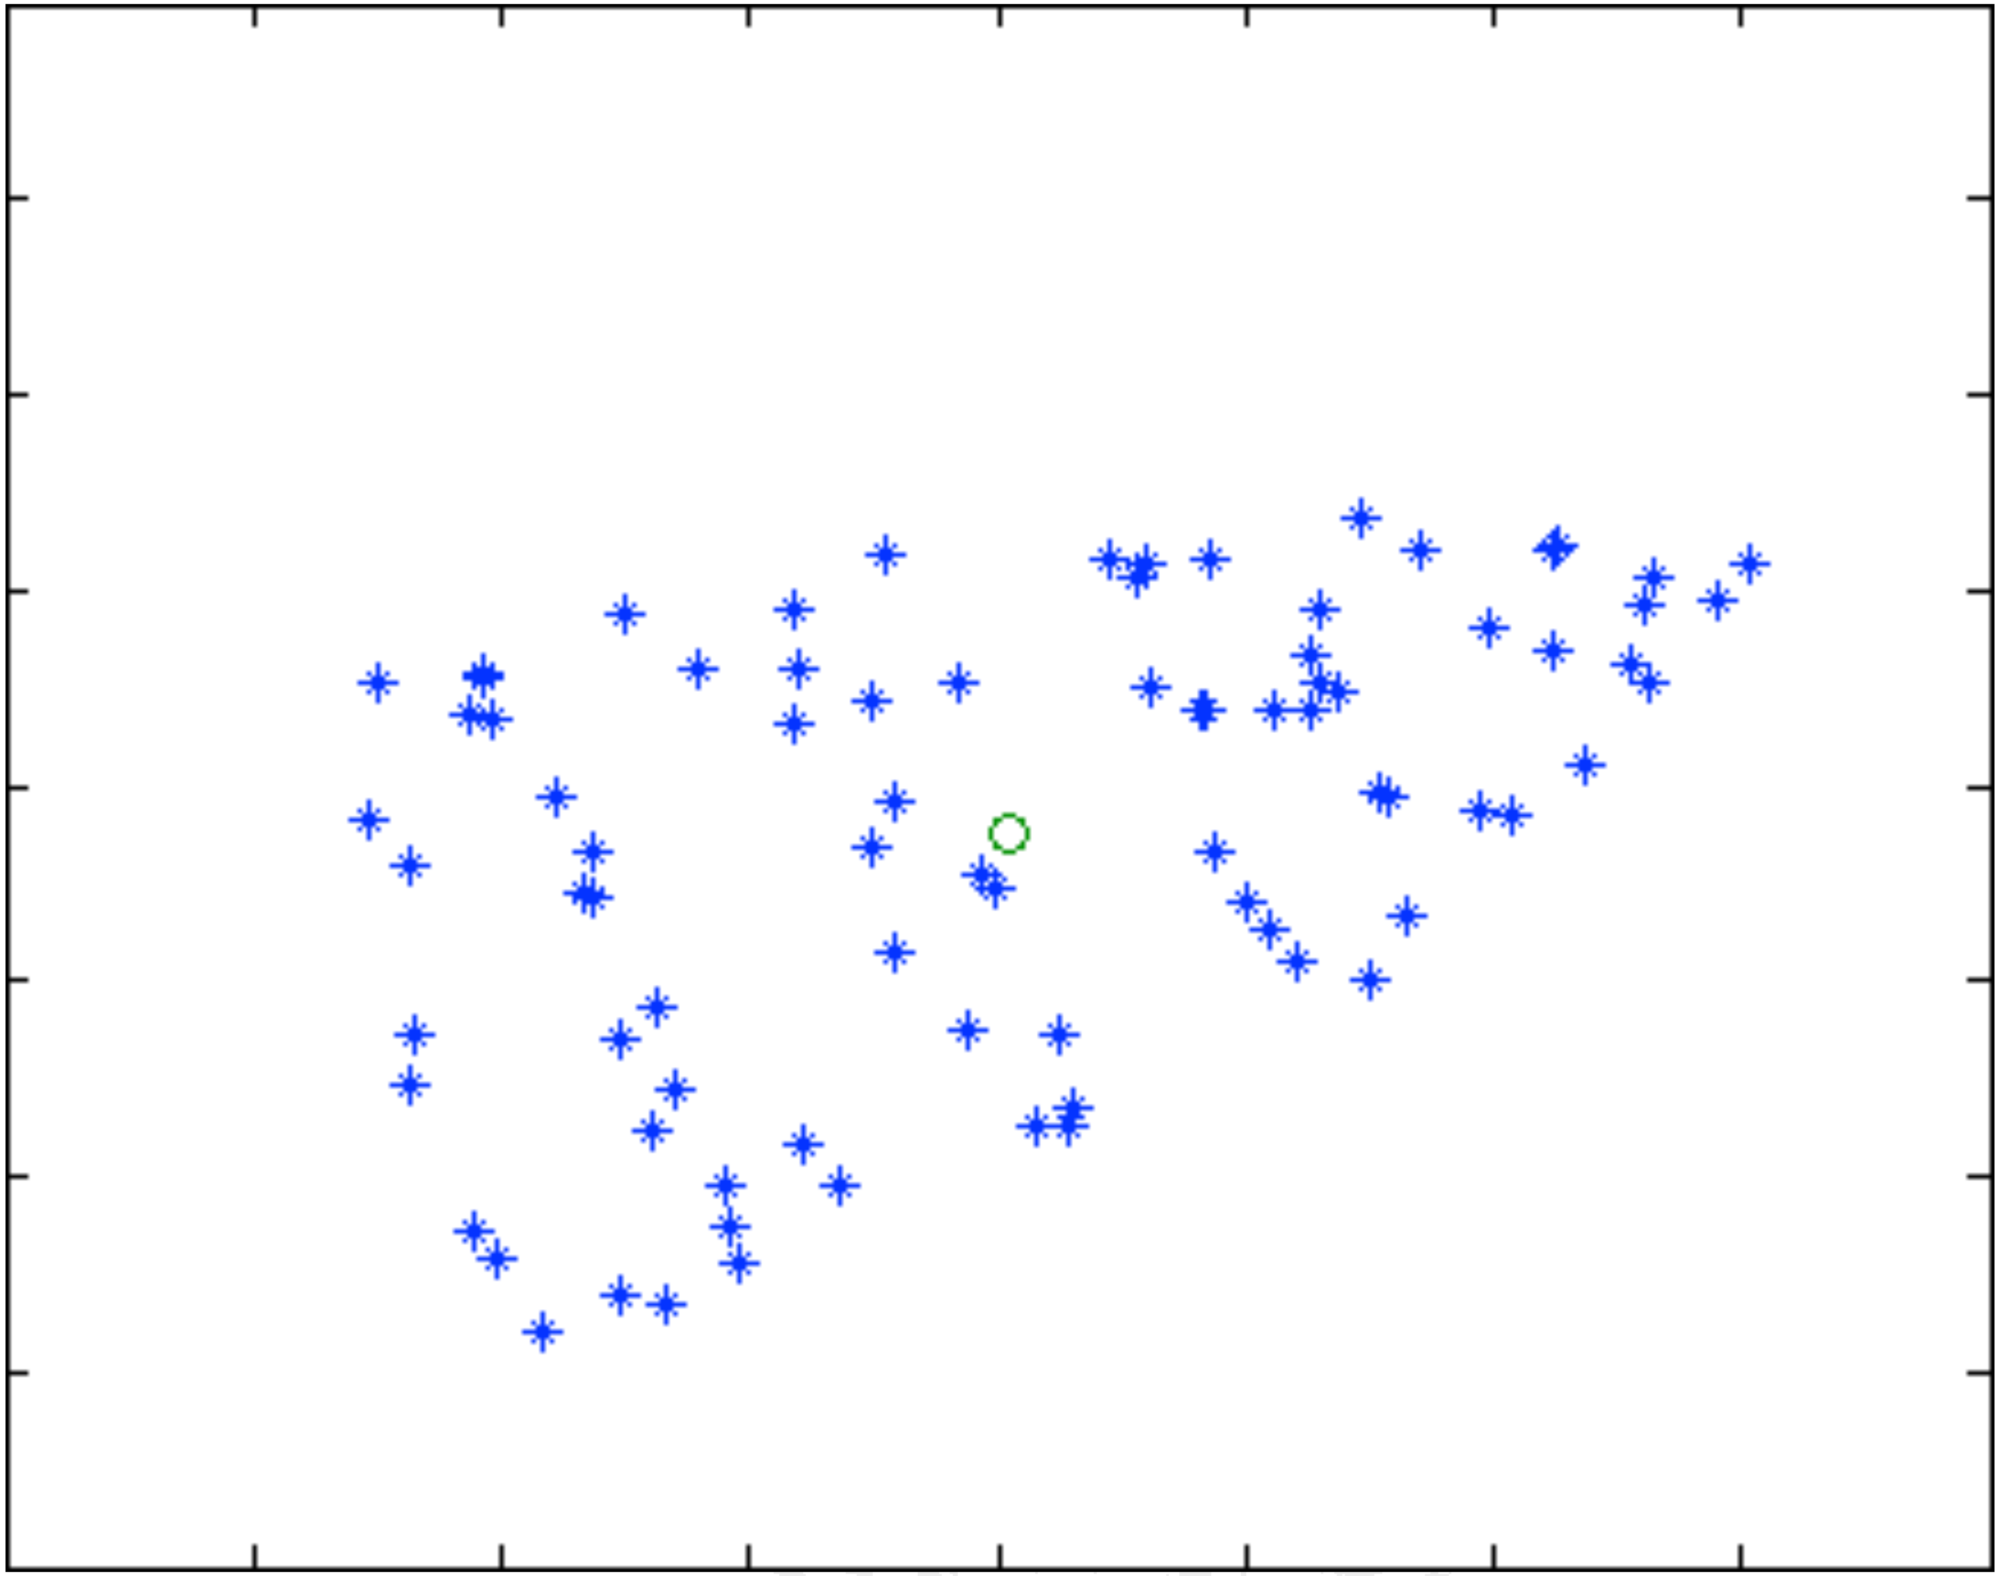

Time varying network topology

We have now randomly time-varying stochastic matrices $P(t)$.
We obtain a randomly switching linear system

$$x(t+1) = P(t)x(t)$$

PROBABILISTIC CONSENSUS

$$x_i(t) \to c \qquad \text{almost surely}$$

where c is a random variable depending on $x(0)$.

THEOREM (Cogburn 1987) Assume that $P(t)$ are i.i.d. Then we have probabilistic consensus iff $\overline{P} = \mathbb{E}[P(t)]$ (which is always stochastic) yields deterministic consensus.

Time varying network topology

We have now randomly time-varying stochastic matrices $P(t)$.
We obtain a randomly switching linear system

$$x(t+1) = P(t)x(t)$$

PROBABILISTIC CONSENSUS

$$x_i(t) \to c \qquad \text{almost surely}$$

where c is a random variable depending on $x(0)$.

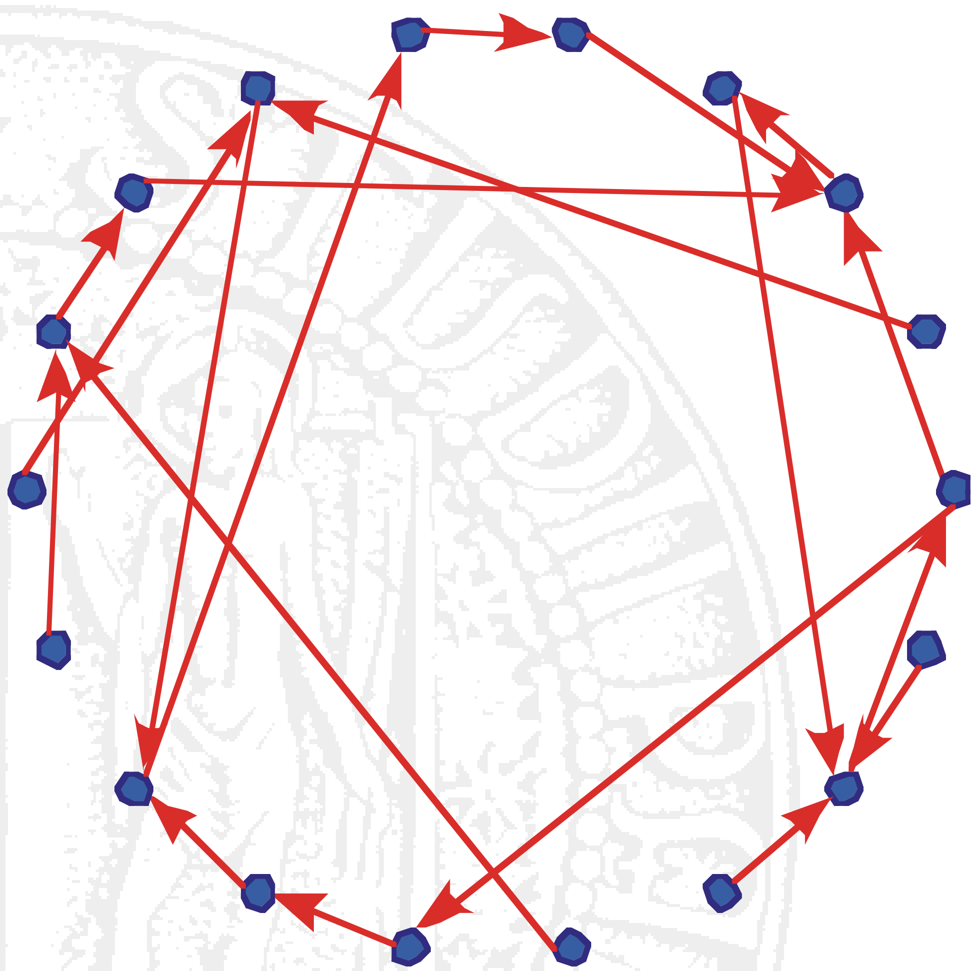

THEOREM (Cogburn 1987) Assume that $P(t)$ are i.i.d. Then we have probabilistic consensus iff $\overline{P} = \mathbb{E}[P(t)]$ (which is always stochastic) yields deterministic consensus.

Time varying network topology

Time varying network topology

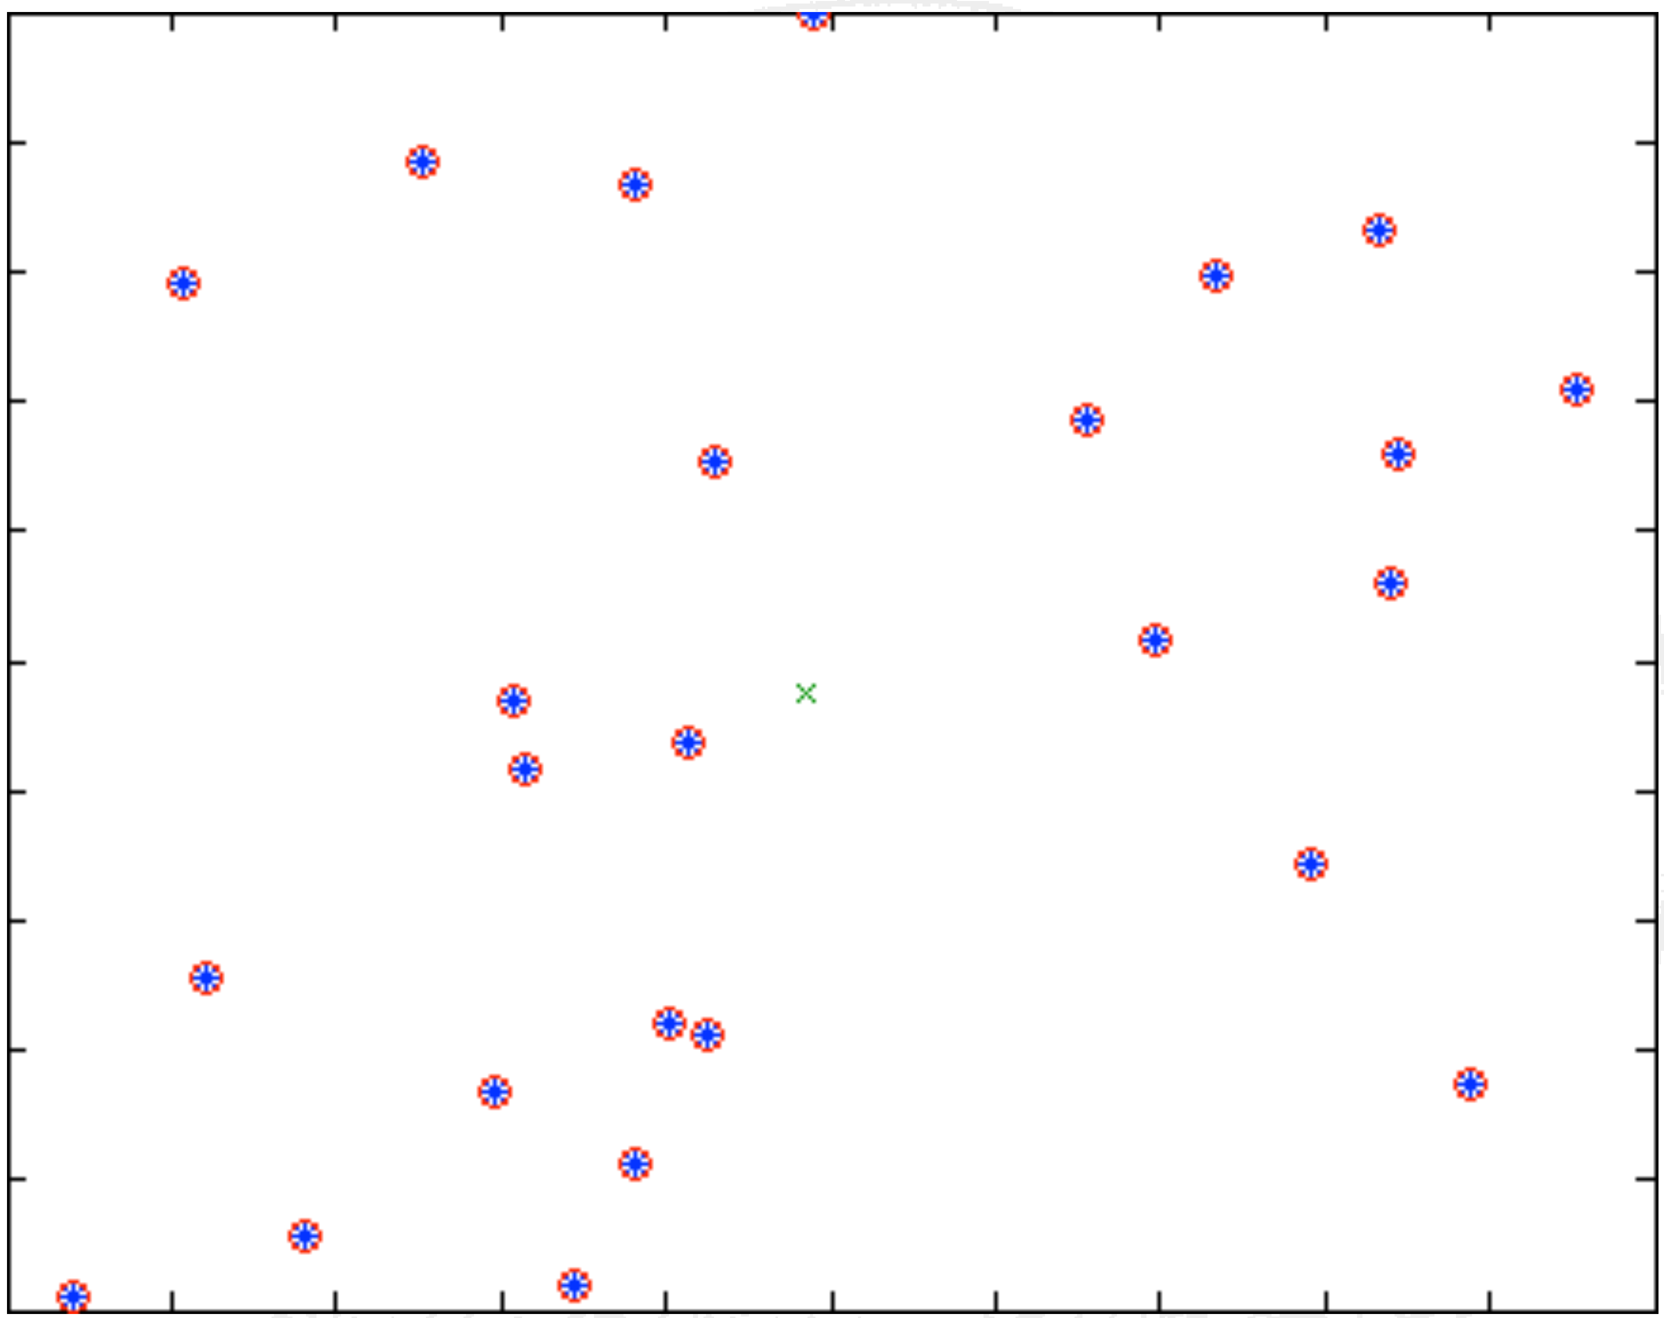

Grazie dell'attenzione