Local optimization in cooperative agents networks

Dott.ssa Meritxell Vinyals
University of Verona

Seville, 30 de Junio de 2011

Outline

. Introduction and domains of interest

. Open problems and approaches

Outline

- Introduction and domains of interest
- Open problems and approaches

Motivating domains

Traffic light control

Many real-world problems can be modeled as a network of cooperative agents that have to coordinate their actions in order to optimize the system performance

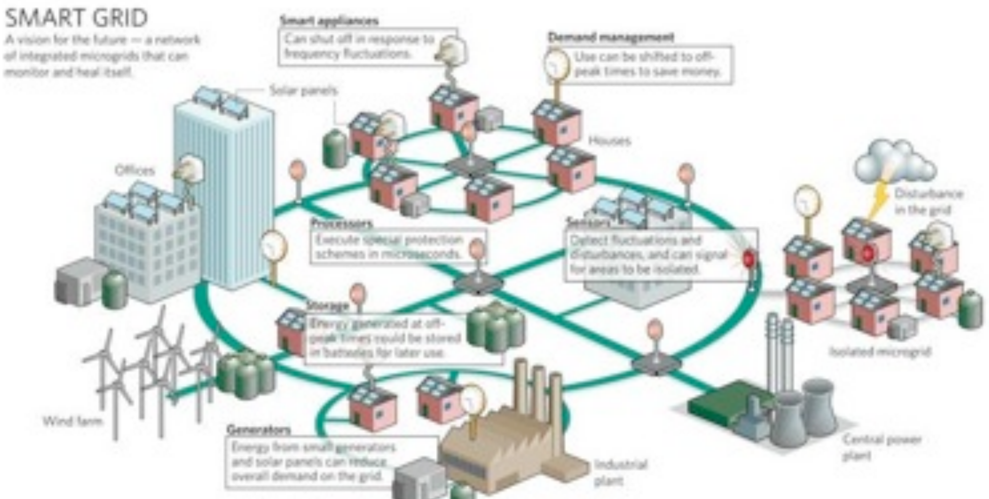

Smart (electricity) grid

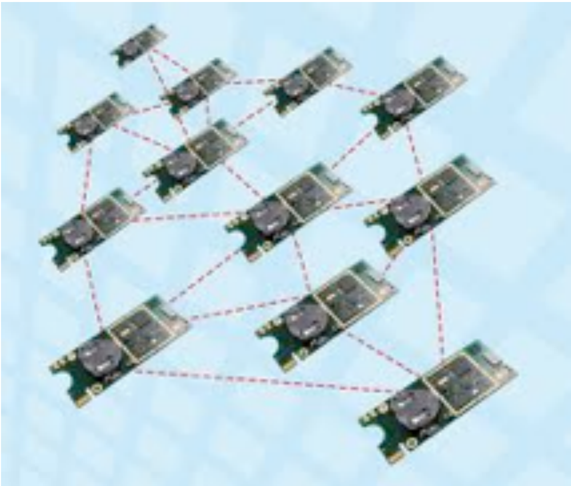

Energy-efficient sensor networks

Motivating domains

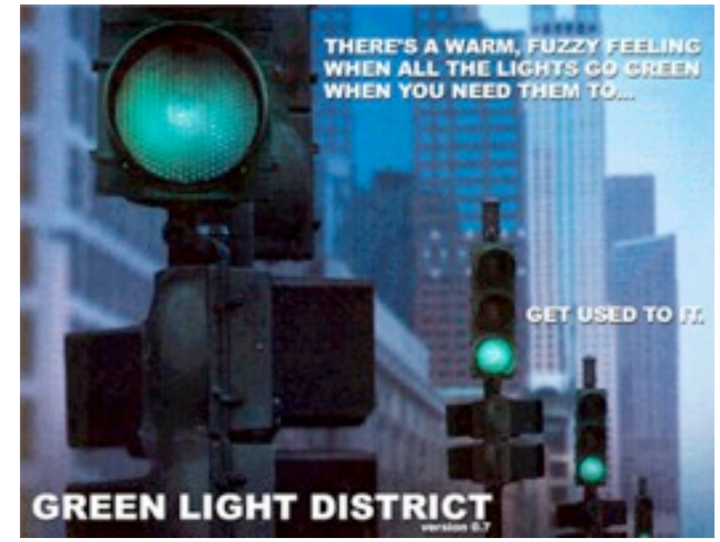

Traffic light control

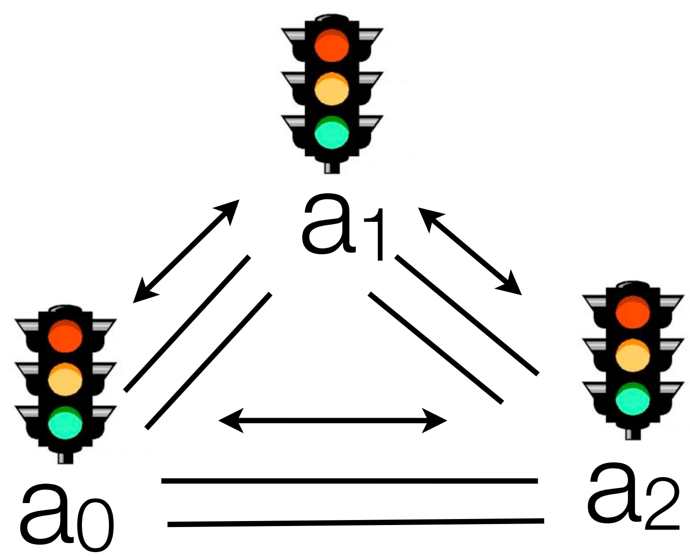

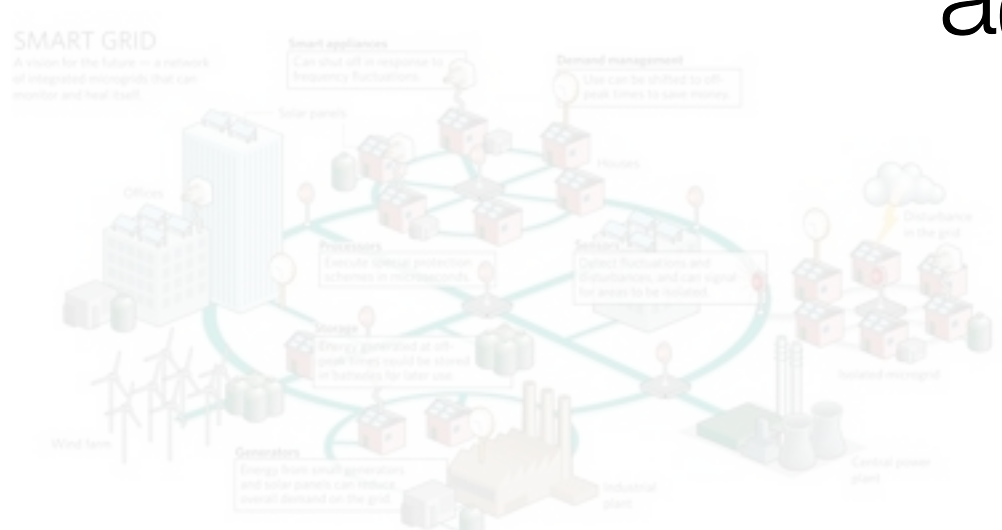

Smart (electricity) grid

Energy-efficient
sensor networks

Motivating domains

Traffic light control

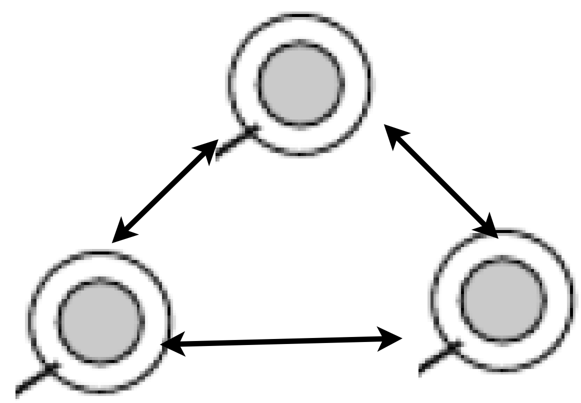

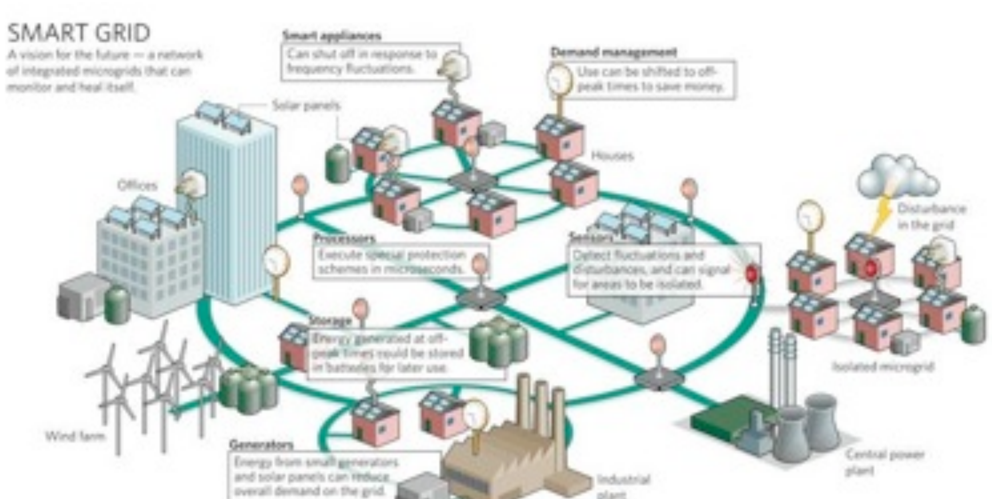

Smart (electricity) grid

Energy-efficient
sensor networks

Motivating domains

Traffic light control

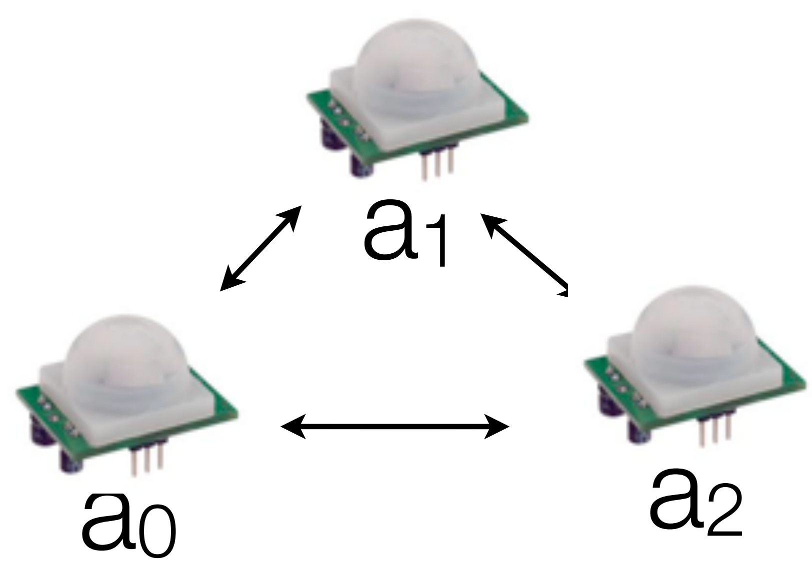

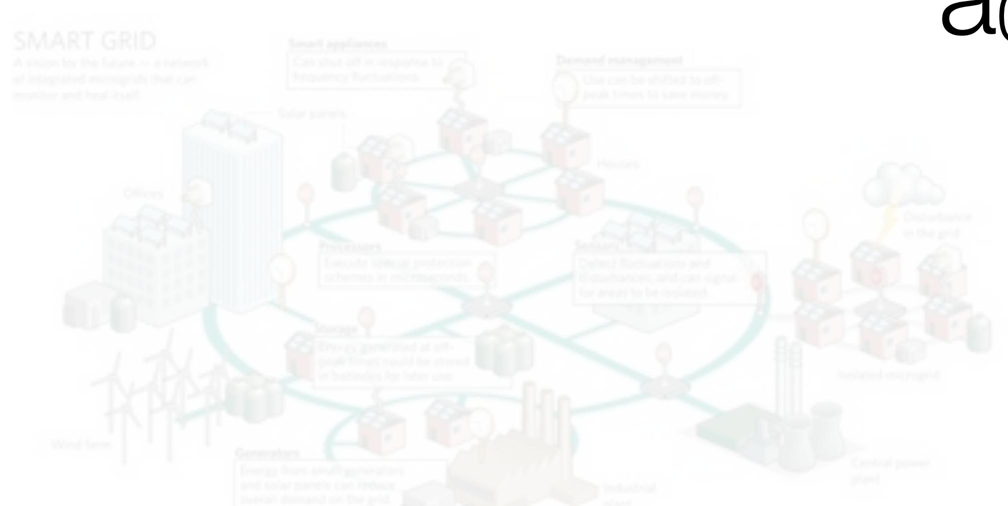

Smart (electricity) grid

Energy-efficient
sensor networks

Motivating domains

Traffic light control

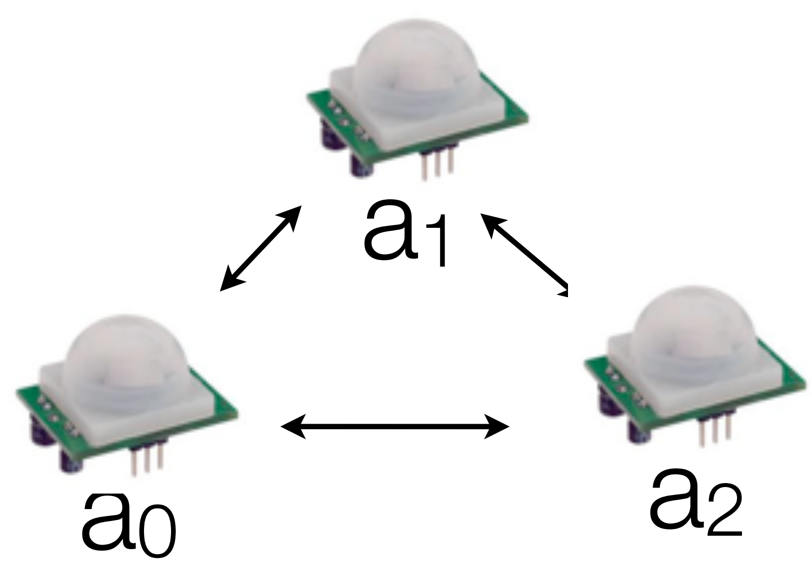

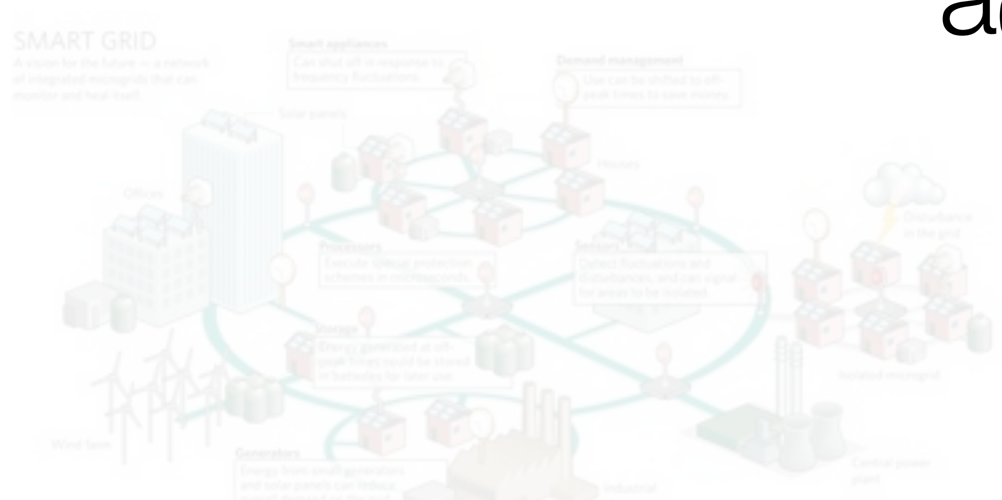

Smart (electricity) grid

Energy-efficient
sensor networks

Optimization with agents

Each agent can choose from a set of discrete actions

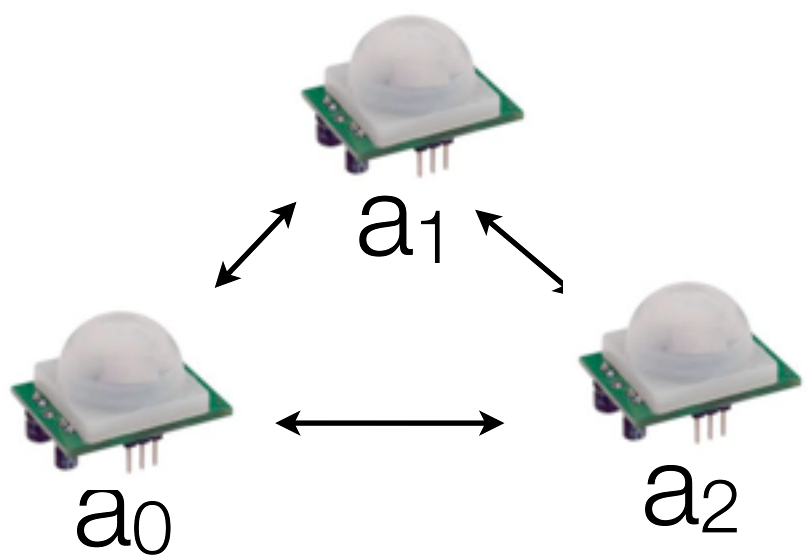

Each agent sensor has two actions:

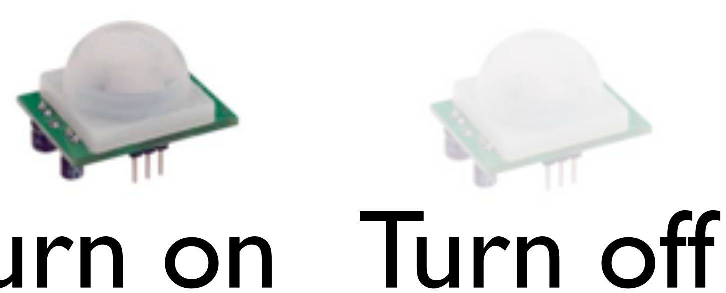

Optimization with agents

Agents report individual rewards for their actions

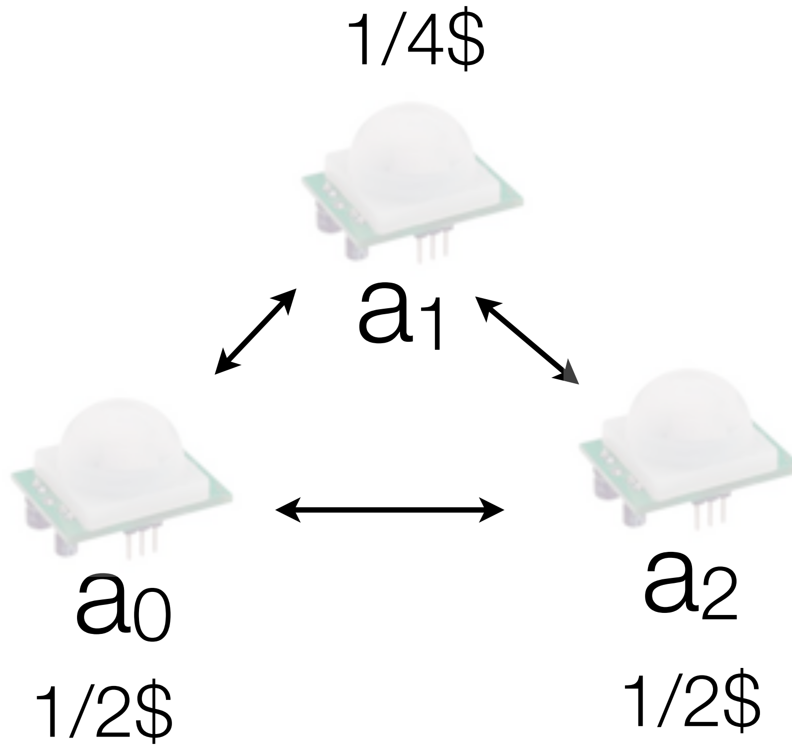

Agent sensors report a reward to turn off (e.g. which may vary depending on the remaining battery)

Optimization with agents

An edge stands for two agents that need to coordinate in order to receive a joint reward

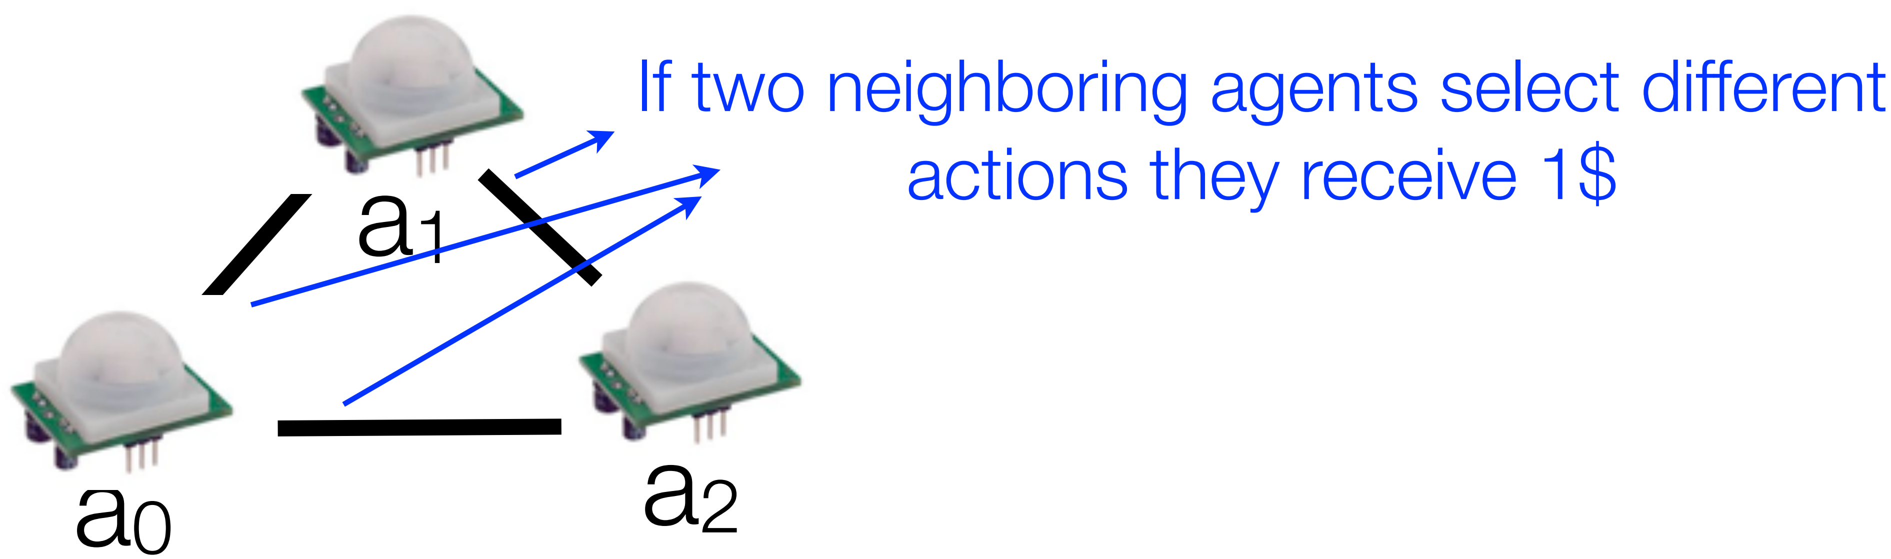

There is a reward if the region between two sensors is sampled by at least one sensor

Optimization with agents

An edge stands for two agents that need to coordinate in order to receive a joint reward

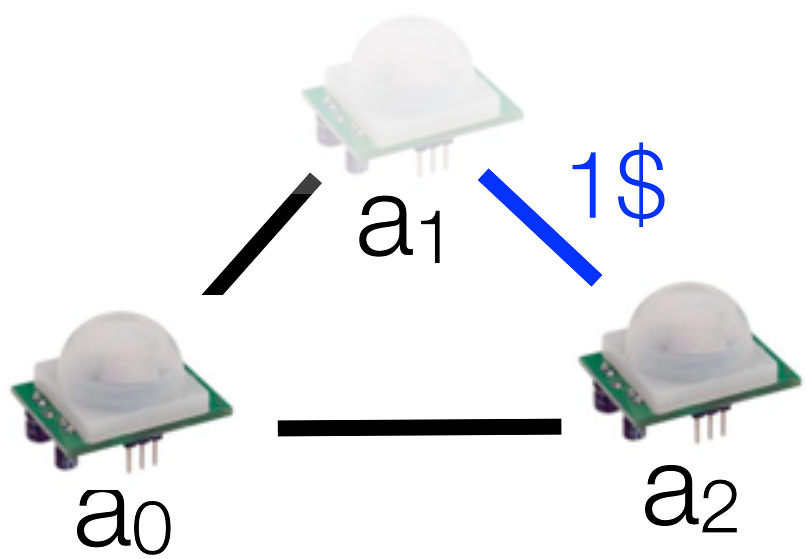

There is a reward if the region between two sensors is sampled by at least one sensor

Optimization with agents

An edge stands for two agents that need to coordinate in order to receive a joint reward

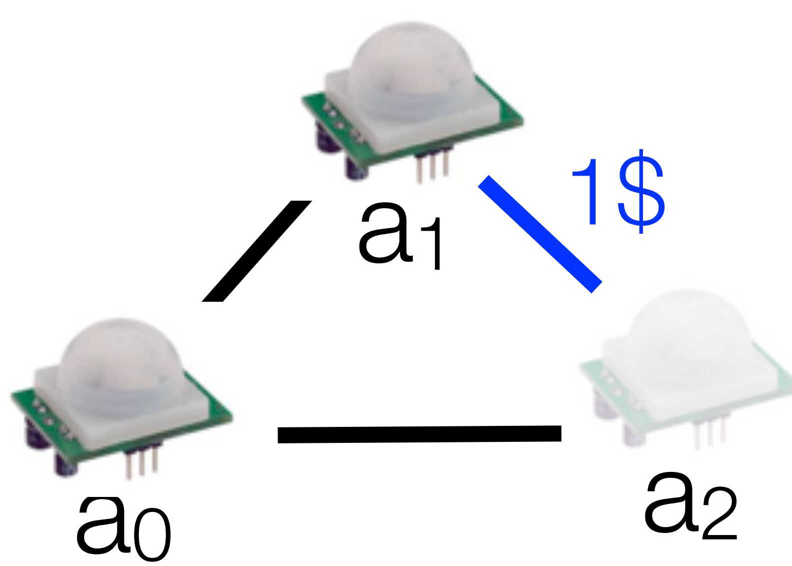

There is a reward if the region between two sensors is sampled by at least one sensor

Optimization with agents

The goal is to distributedly find a set of actions that maximize the overall reward

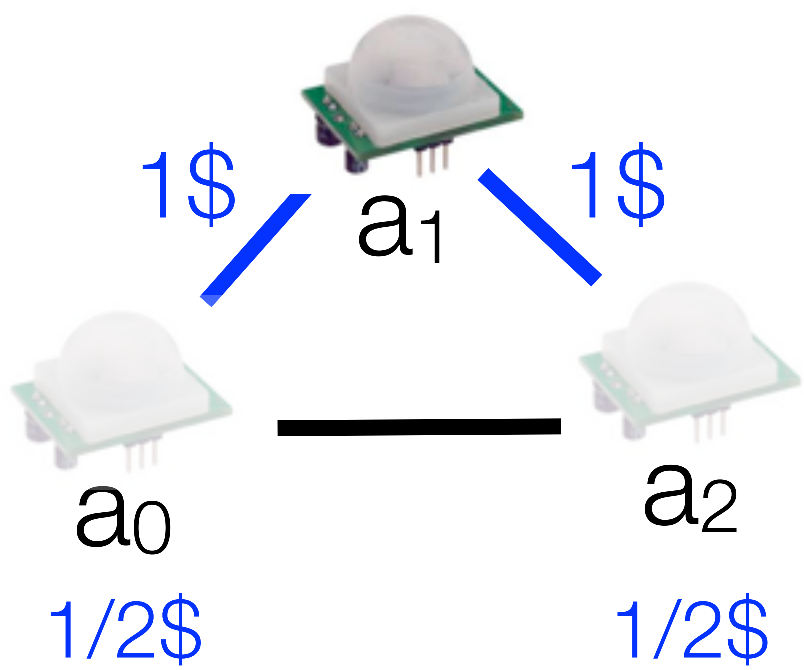

- Optimal configuration 3$

Motivating domains

Disaster management

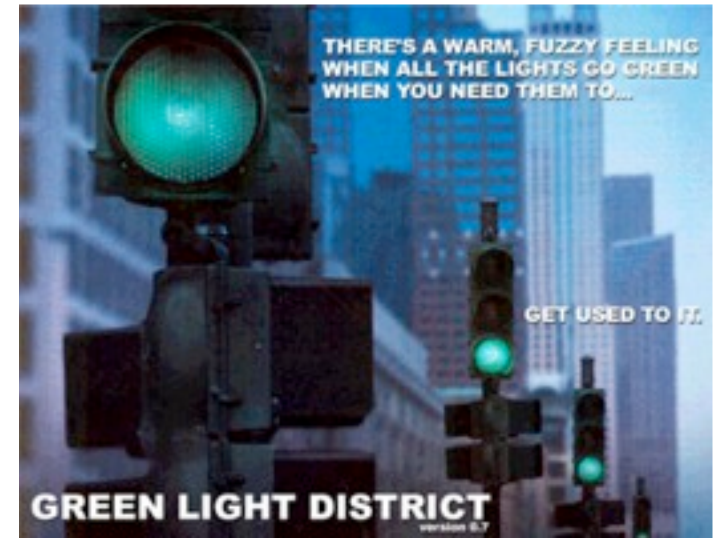

Traffic light control

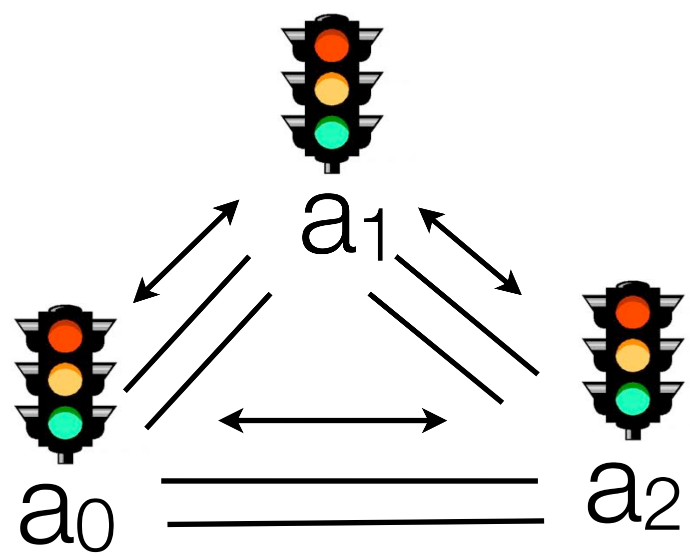

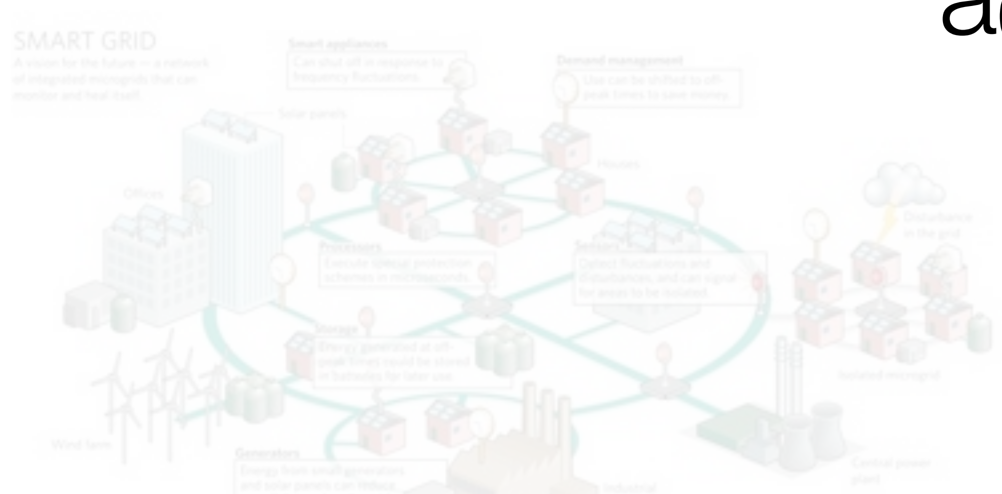

Smart (electricity) grid

Energy-efficient sensor networks

Traffic light control

Old times:
isolated traffic lights

Future generation:
social traffic lights

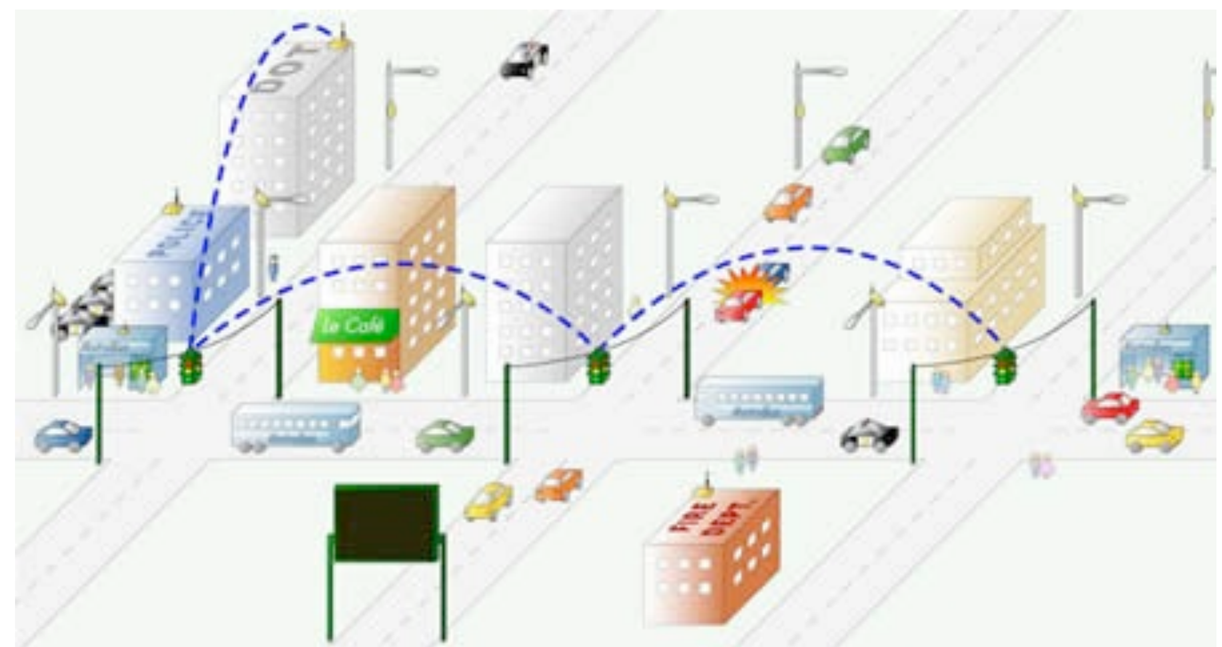

Coordinate traffic lights so that vehicles can traverse an arterial in one traffic direction, keeping a specified speed without stopping (green waves)

Traffic light control

[R. Junges and A. L. C. Bazzan, 2010] uses a multi-agent system approach in which:

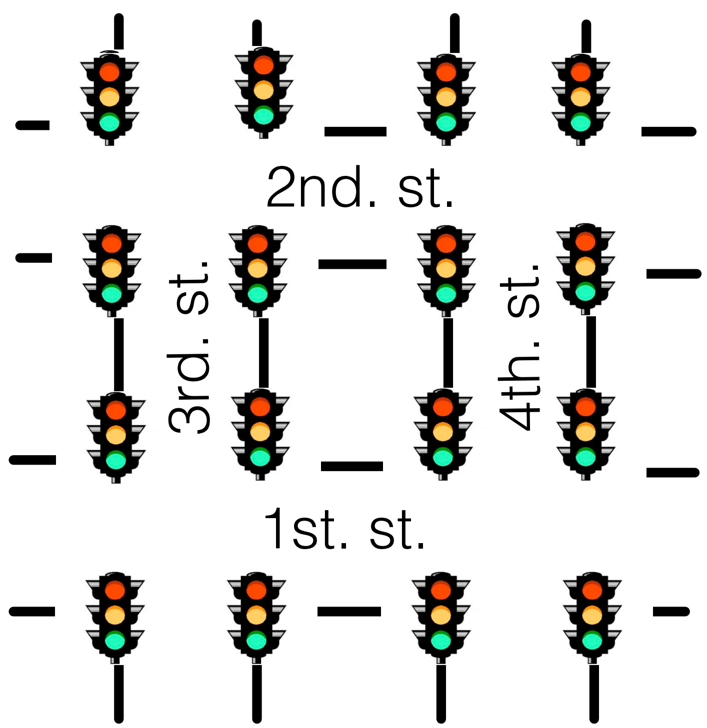

Traffic light control

[R. Junges and A. L. C. Bazzan, 2010] uses a multi-agent system approach in which:

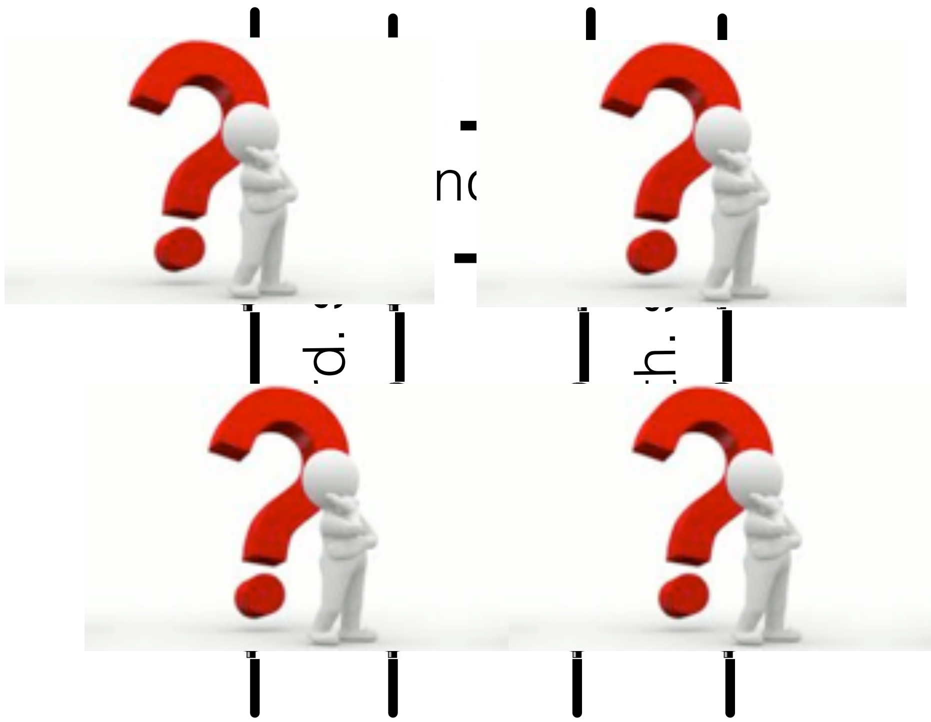

Each agent is in charge of a crossing

Traffic light control

[R. Junges and A. L. C. Bazzan, 2010] uses a multi-agent system approach in which:

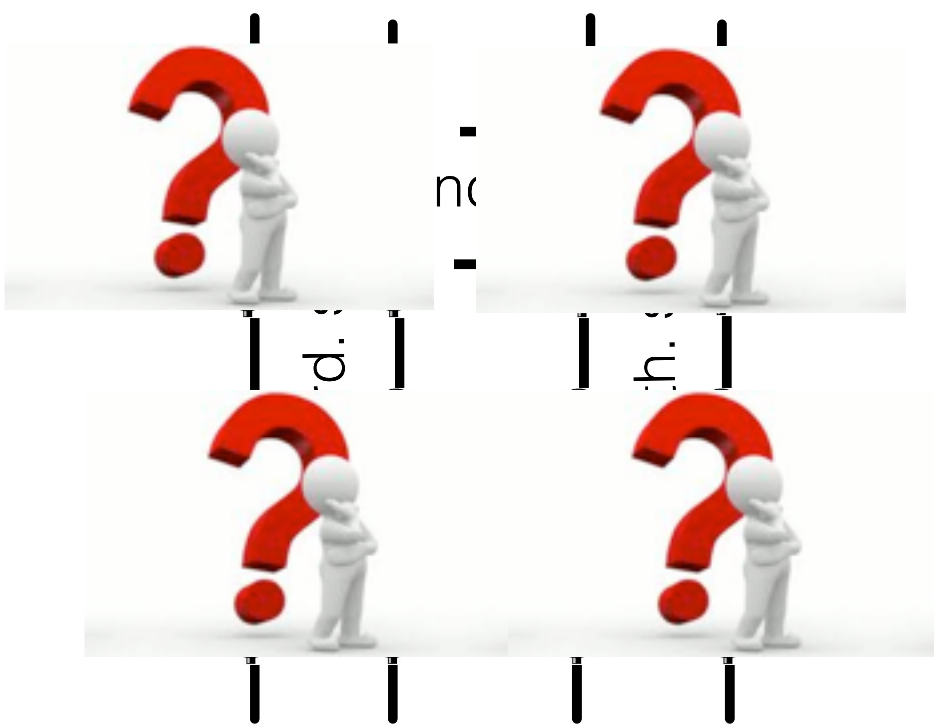

The decision of an agent is composed of a set of signals plans for the traffic lights in the crossing

$\oplus$ $\otimes$

Traffic light control

[R. Junges and A. L. C. Bazzan, 2010] uses a multi-agent system approach in which:

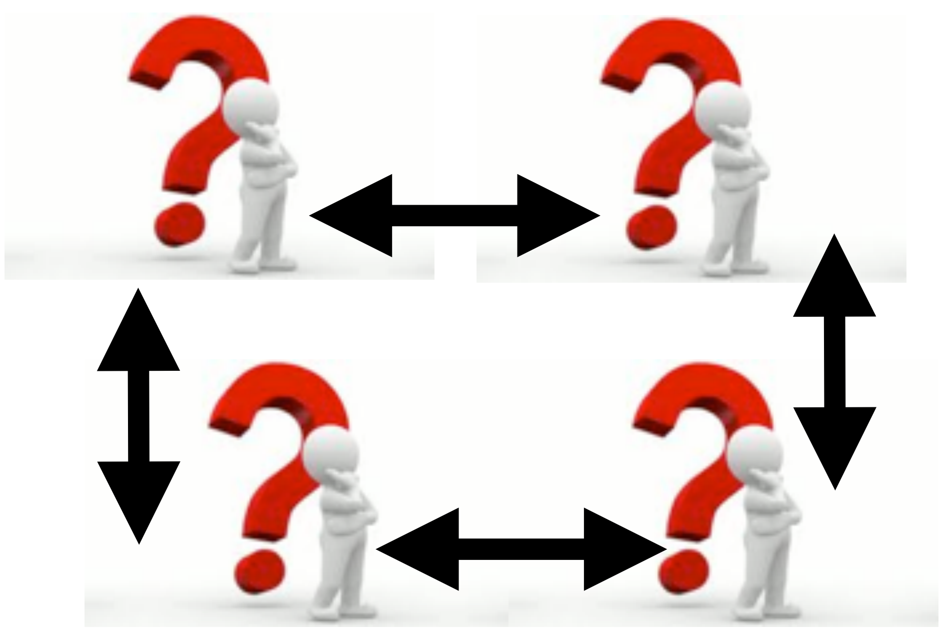

Two agents in two adjacent crossings need to coordinate their plans

Traffic light control

[R. Junges and A. L. C. Bazzan, 2010] uses a multi-agent system approach in which:

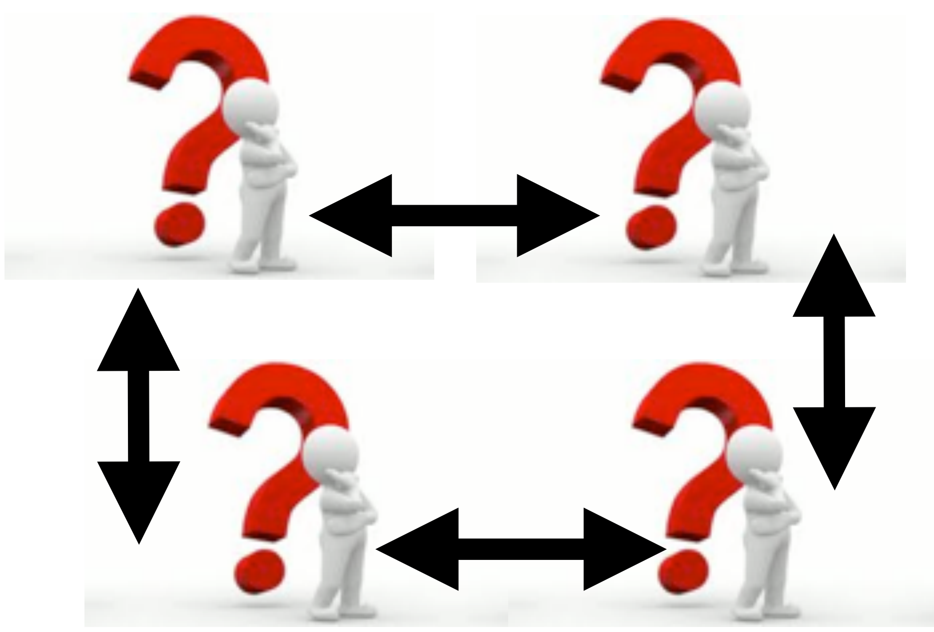

The reward to execute two plans in neighboring crossings is in function of:

(I) the degree in which these two plans synchronize

(II) the volume of vehicles in that direction

Traffic light control

[R. Junges and A. L. C. Bazzan, 2010] uses a multi-agent system approach in which:

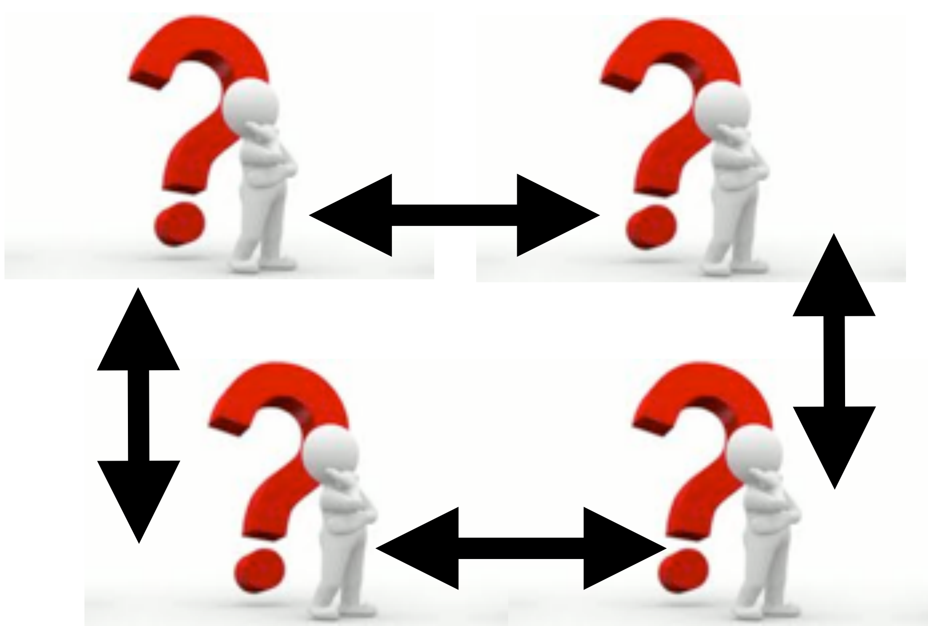

	⊕	⊗
⊕	2$	1$
⊗	2$	0$

Smart grid

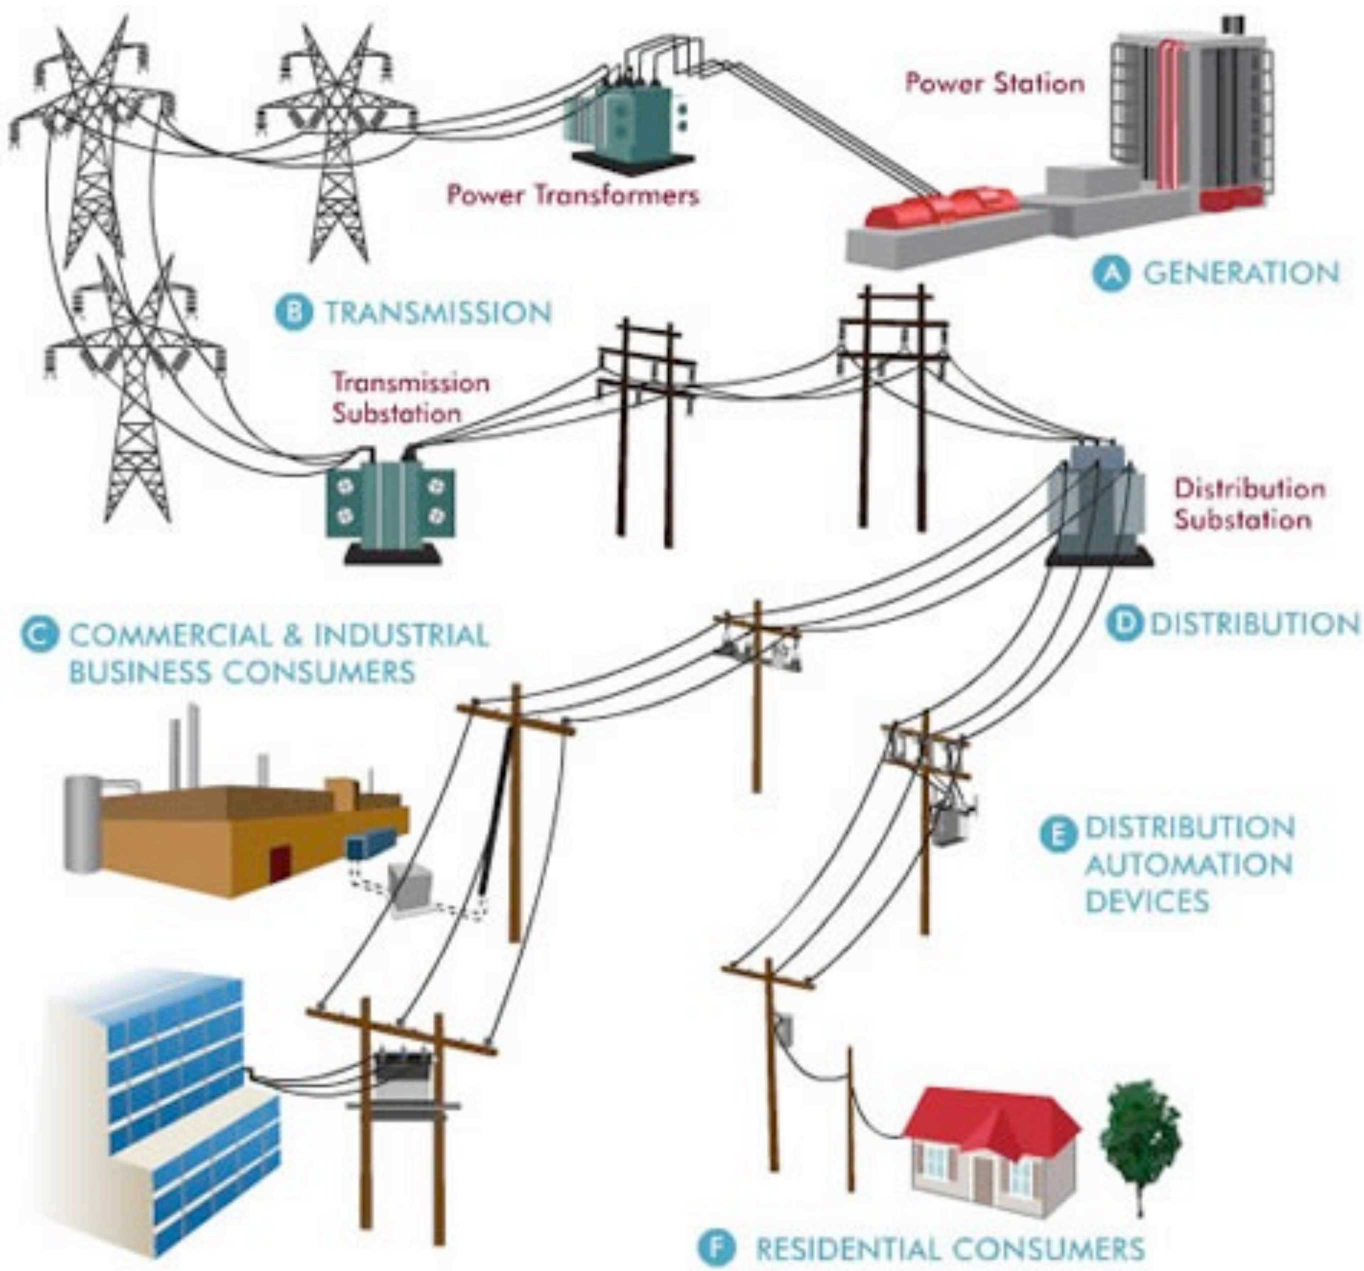

. The current hierarchical, centrally-controlled grid is obsolete
. Problems on scalability, efficiency and integration of green energies
. Most of the decisions about the operation of a power system are made in a centralized fashion

Smart grid

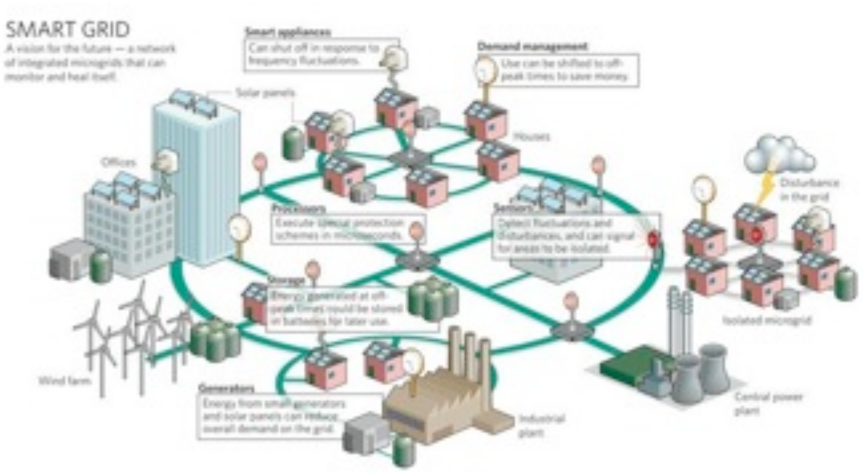

. Centralized control is replaced with decentralized control:

. efficiency and scalability

. complex control mechanisms needed

. Introduction of intelligence at all levels, especially at lower levels, to provide timely and accurate control responses

Smart grid

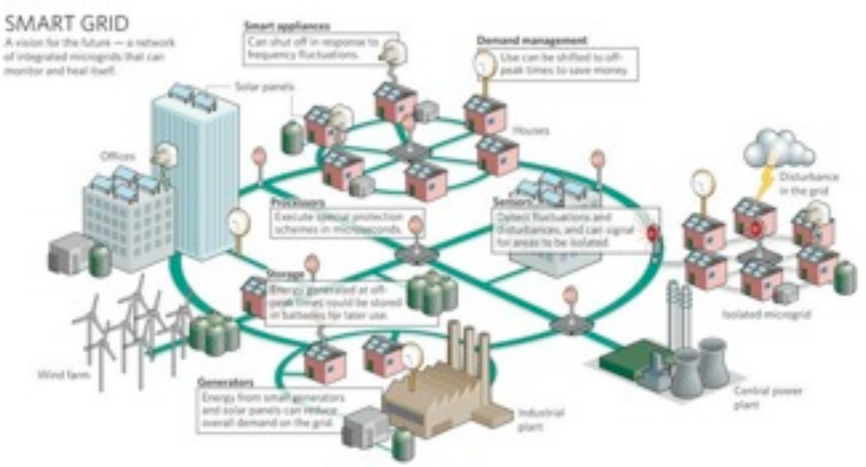

. Home/neighborhood level

. Distribution level

. Transmission level

Smart grid

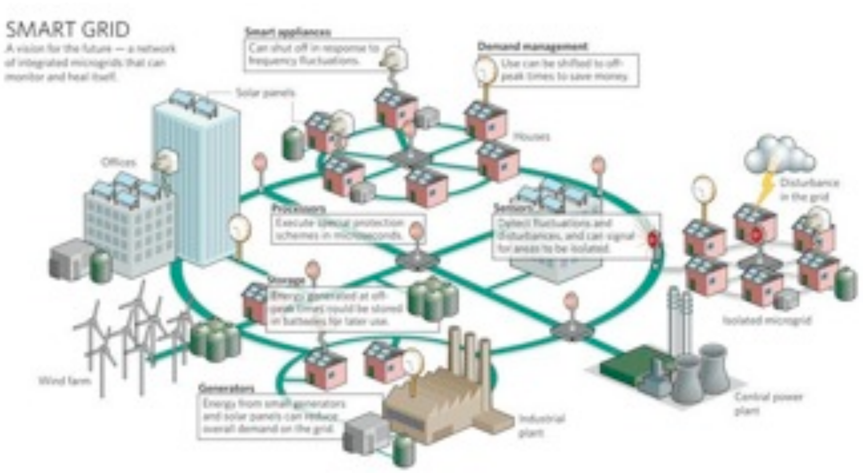

. Home/neighborhood level

. Coordinate home appliances to reduce the peak

. Create coalitions of energy profiles to reduce the peak

Smart grid

[Petcu & Faltings, 2008] . Transmission level

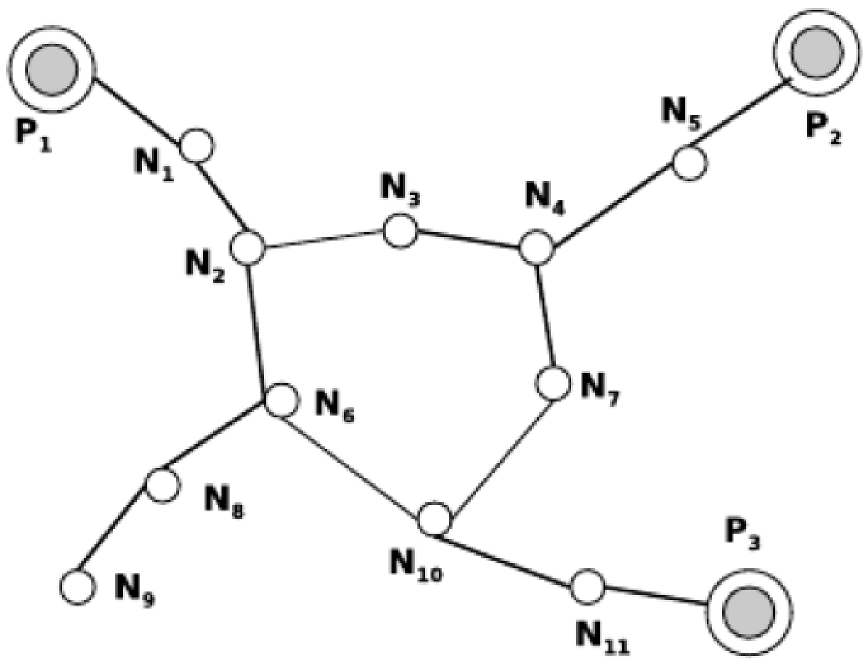

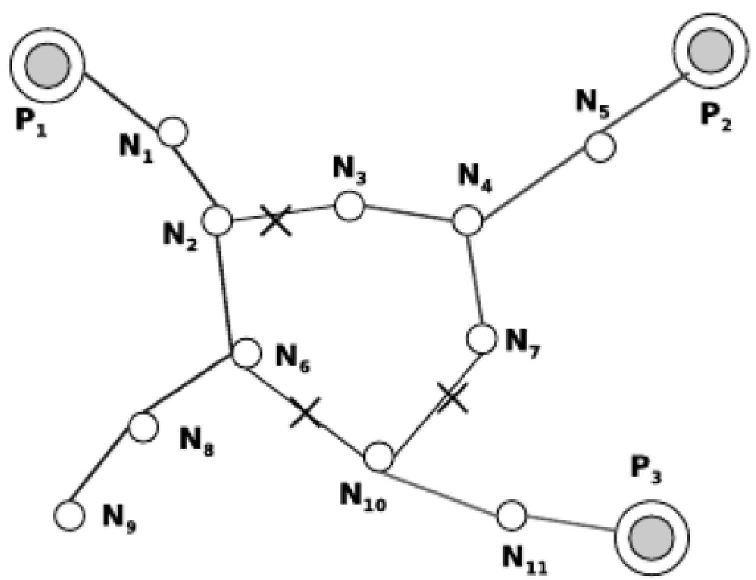

How sinks configure the network by enabling transmission lines such that are:

(a) cycle free; and

(b) the amount of power lost is minimized.

Outline

- Introduction and domains of interest
- Open problems and approaches

Maxims for researchers

First takes all the credit, second gets nothing

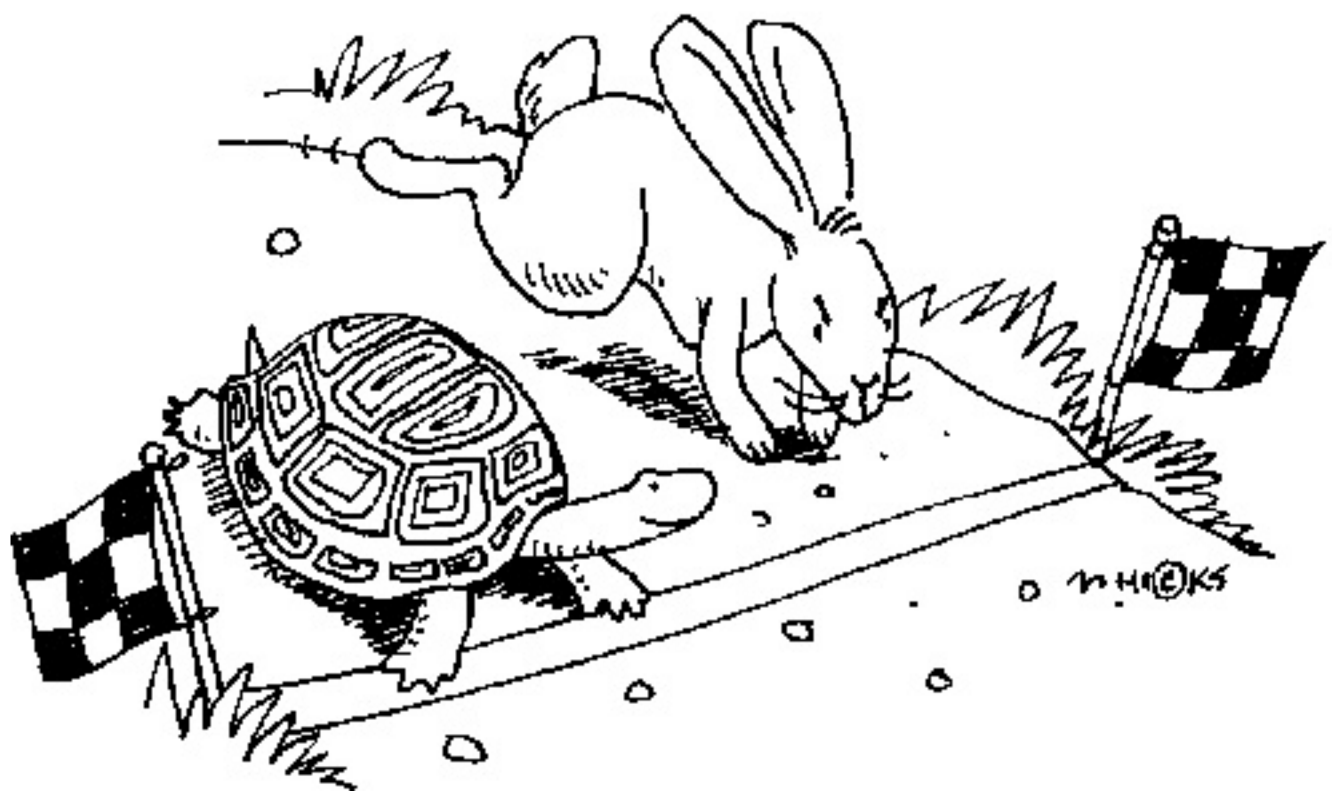

Maxims for researchers

Either you are the first or you are the best in the crowd

The moral

Identify open problems, preferably with few contributions

Open problems

Trade-off quality vs cost

Dynamism

Uncertainty

Hierarchical optimization

Non-cooperative agents

Open problems

Trade-off quality vs cost

Dynamism

Uncertainty

Hierarchical optimization

Non-cooperative agents

Optimality: the idealistic (but usually impractical) term

Researchers have proposed optimal algorithms that aim to minimize the communication/computation needed by agents to find their optimal actions

DPOP [A. Petcu & B. Faltings, 2005] ADOPT [Modi et al., 2005]
OptAPO [R. Mailler & V. Lesser, 2004] Action-GDL [M. Vinyals et al., 2010]

All of them have an exponential cost (either in size/number of messages/computation)

Optimality: the idealistic (but usually impractical) term

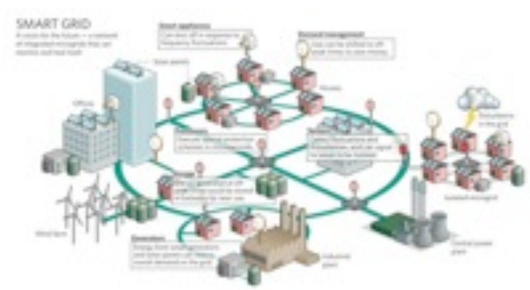

Smart grid

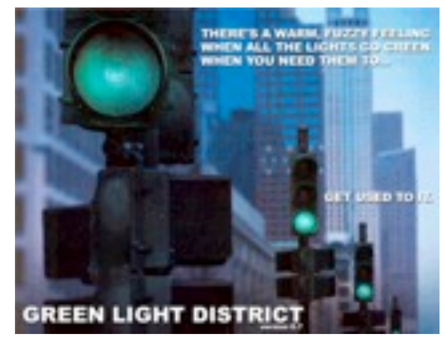

Traffic light control

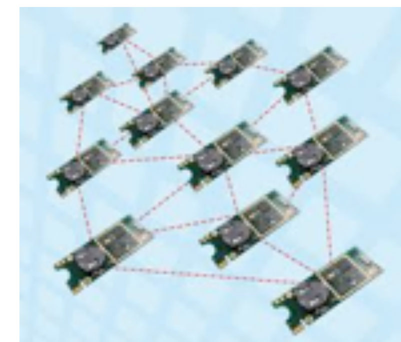

Energy-efficient sensor networks

In many domains the price of optimality is simply not affordable

Time

Scale

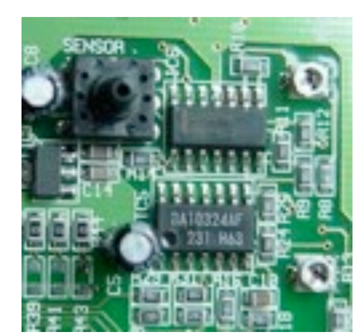

Communication/Computation

Optimality: the idealistic (but usually impractical) term

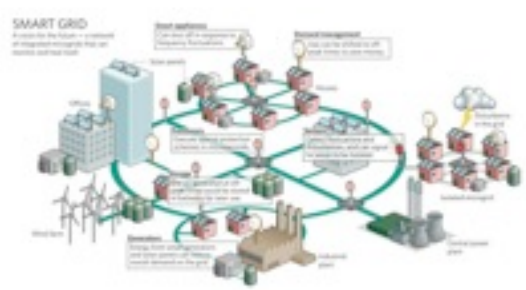

Smart grid

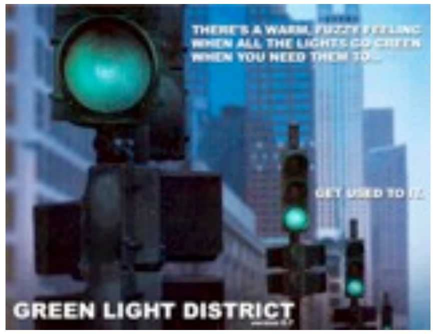

Traffic light control

Energy-efficient sensor networks

In many domains the price of optimality is simply not affordable

Time

Scale

Communication/Computation

Optimality: the idealistic (but usually impractical) term

Smart grid

Traffic light control

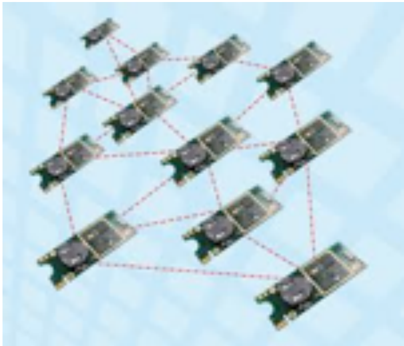

Energy-efficient sensor networks

In many domains the price of optimality is simply not affordable

Time

Scale

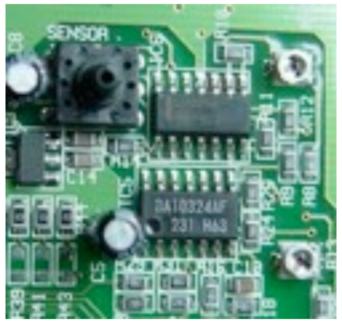

Communication/Computation

Suboptimality: low-cost at not guarantees

Researchers have also proposed suboptimal algorithms:

. Return fast good solutions in average

. Small amount of communication/computation per agent

DSA [Yokoo & Hirayama, 1996] DBA [Fitzpatrick & Meeterns., 2005]
Max-Sum [Farinelli et al., 2009]

But not guarantee ....

Suboptimality: low-cost at not guarantees

Although suboptimal coordination returns good solutions on average ....

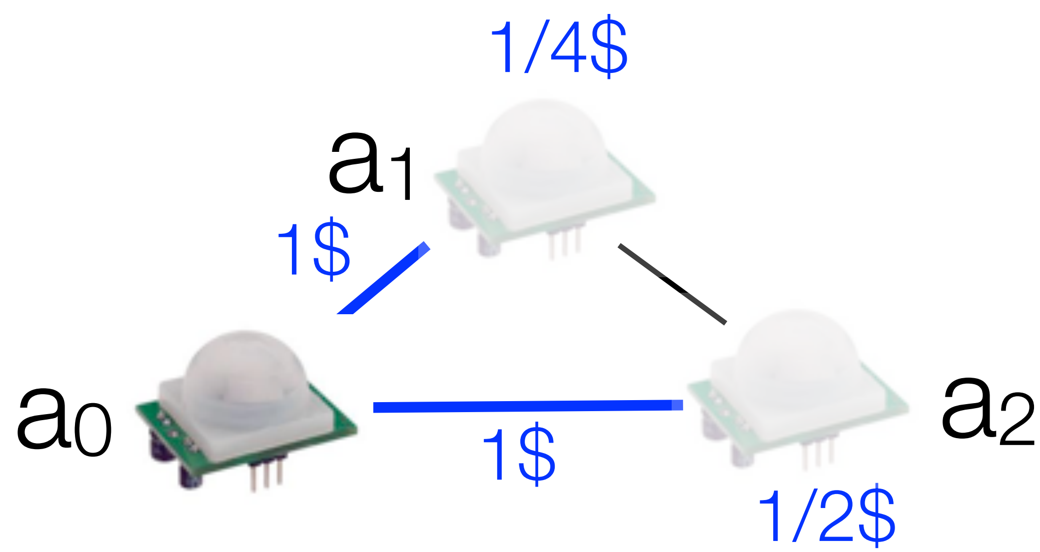

2.75$

(Optimal configuration 3$)

Suboptimality: low-cost at not guarantees

... it can also converge to very poor solutions

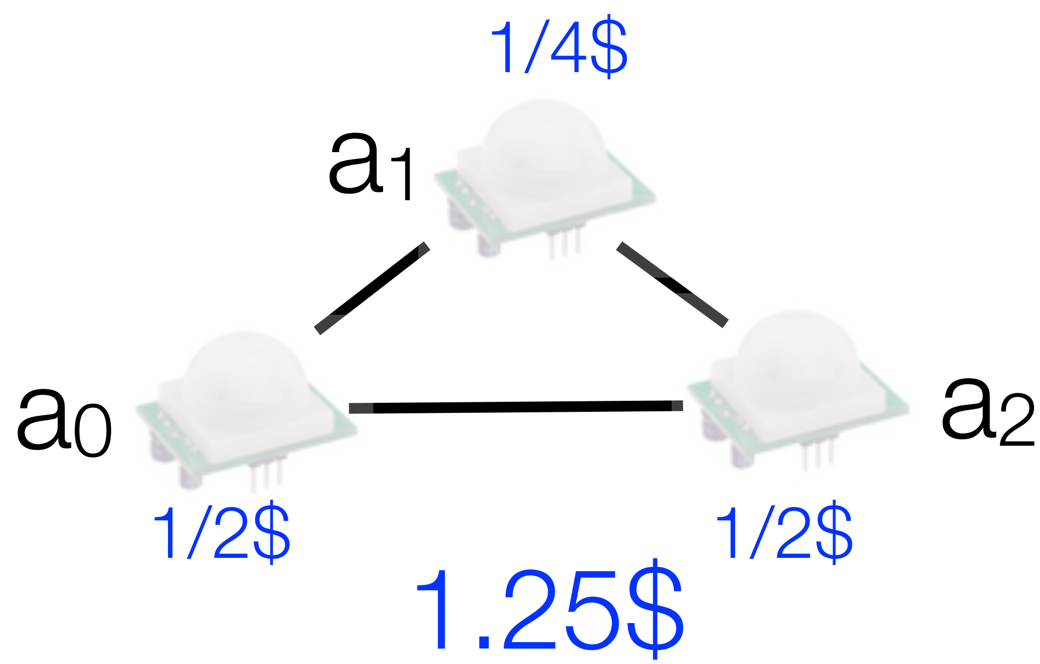

(Optimal configuration 3$)

Suboptimal coordination with quality guarantees

A quality guarantee ensures that the value of a solution is within a given distance δ from the optimal one

$$\delta \leq \frac{\text{solution value}}{\text{optimal value}}$$

Suboptimal coordination with quality guarantees

A quality guarantee ensures that the value of a solution is within a given distance δ from the optimal one

$\delta = 50\%$

Suboptimal coordination with **quality guarantees**

A quality guarantee ensures that the value of a solution is within a given distance δ from the optimal one

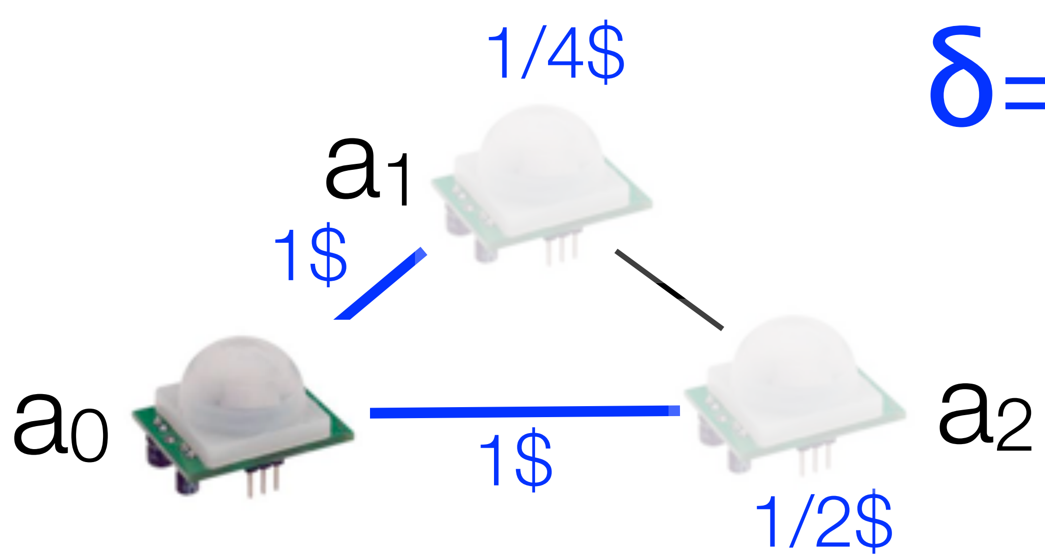

δ= 50%

2.75$

91%

Suboptimal coordination with **quality guarantees**

A quality guarantee ensures that the value of a solution is within a given distance δ from the optimal one

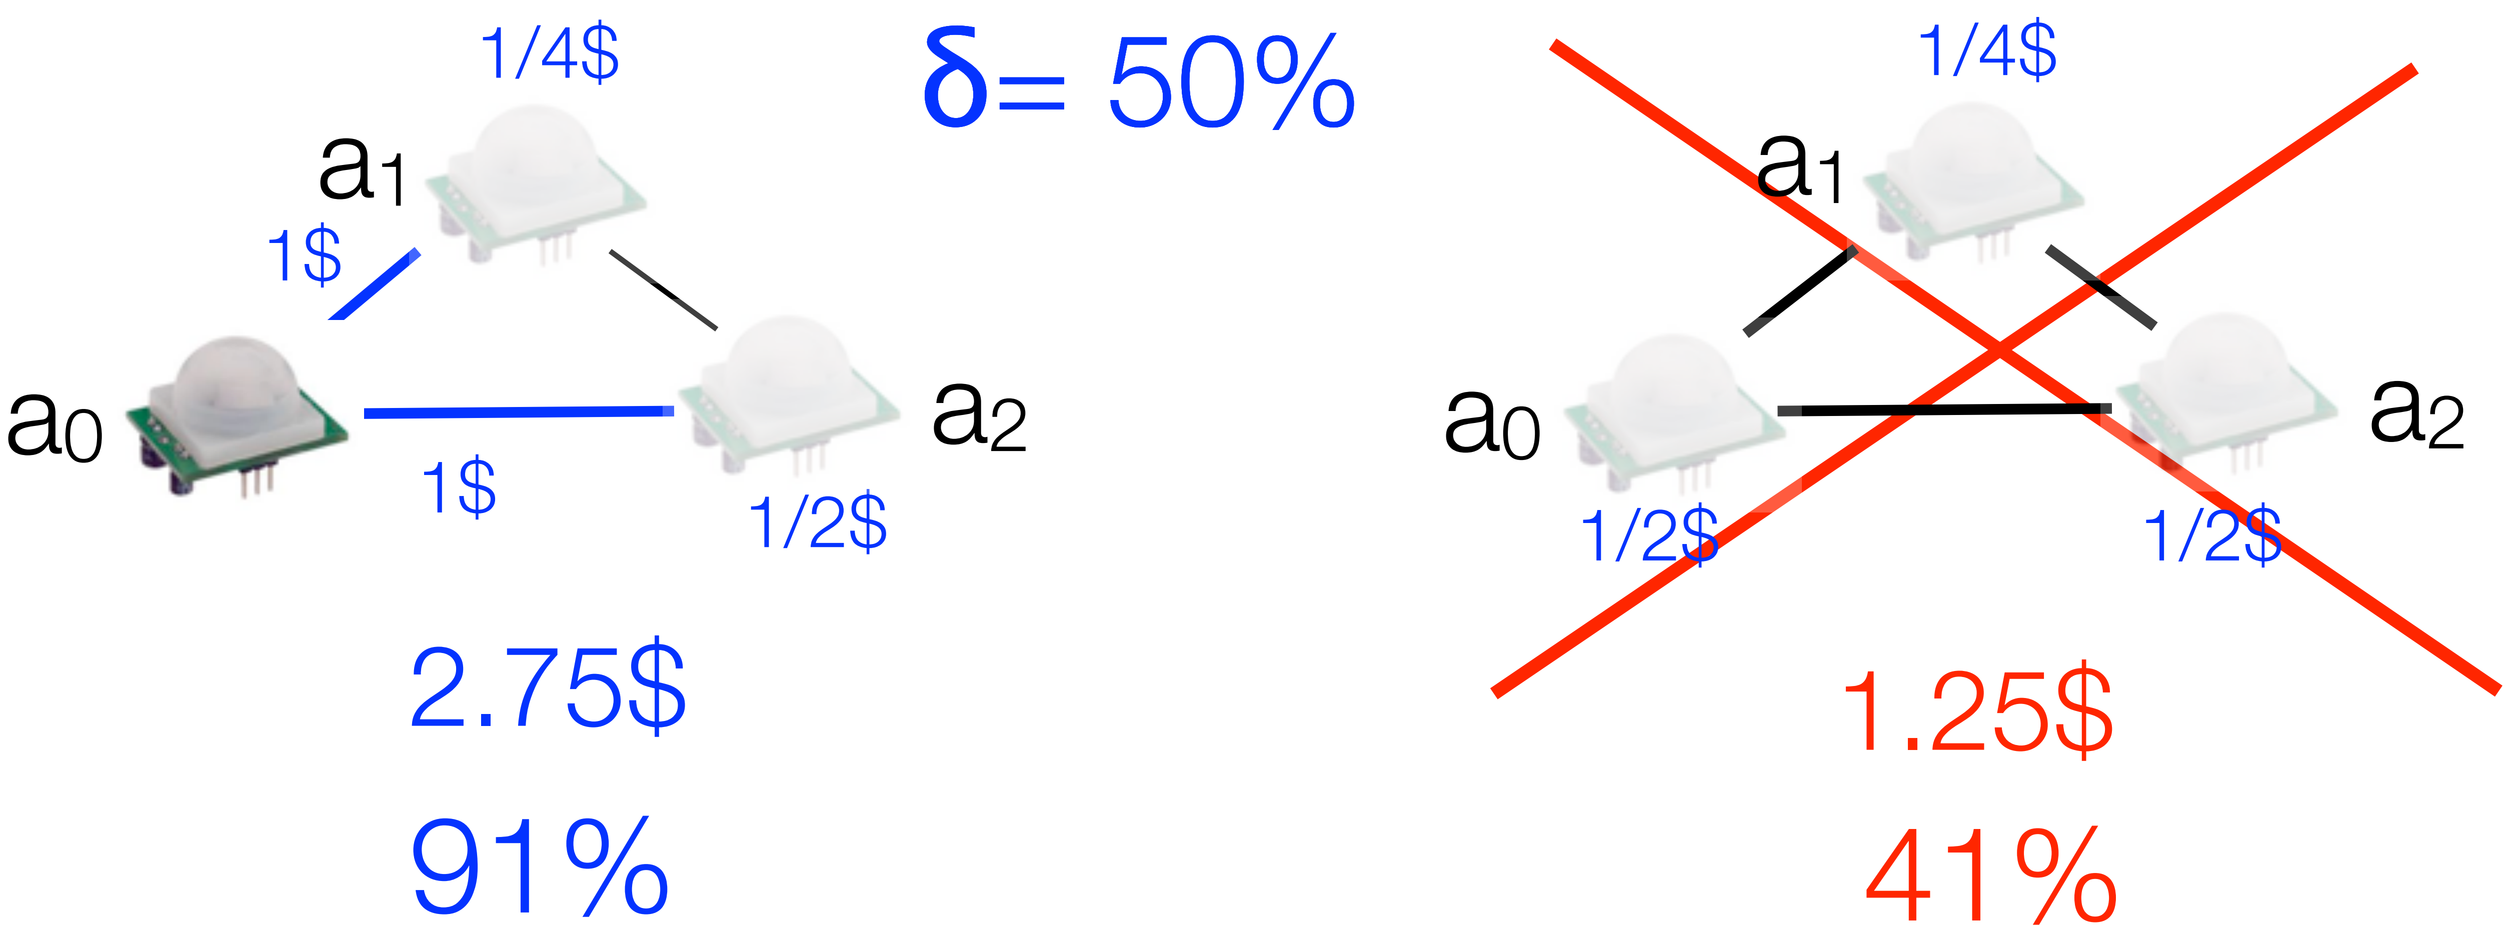

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: algorithm selection

e.g. in traffic control

Algorithm A   Algorithm B

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: algorithm selection

e.g. in traffic control

The best solution varies with traffic conditions ....

Algorithm A

Algorithm B

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: algorithm selection

e.g. in traffic control

Algorithm A   Algorithm B

... but the structure of dependencies is fixed and determined by the particular urban grid

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: algorithm selection

e.g. in traffic control

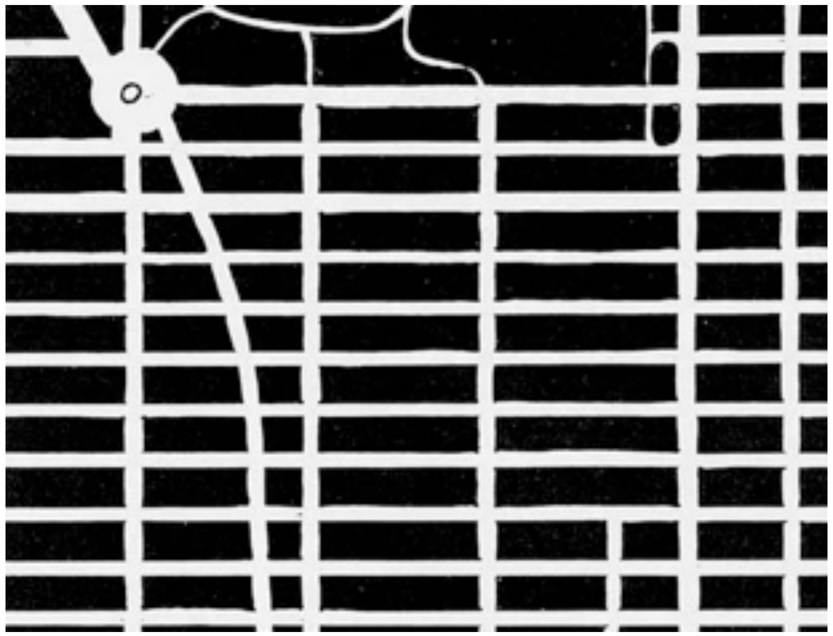

Algorithm A	Algorithm B
50%	25%

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: algorithm selection

e.g. in traffic control

Algorithm A
20%

Algorithm B
40%

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: configuration selection

e.g. in sensor networks

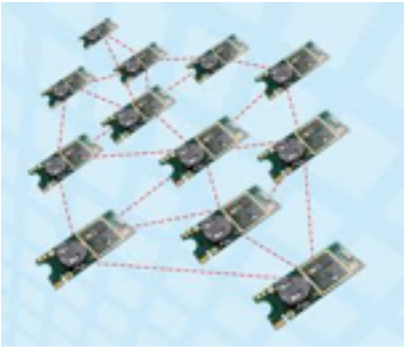

We can select a placement for sensors

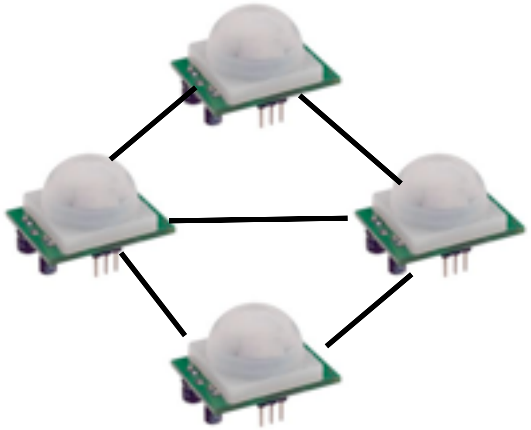

30%

Suboptimal coordination with quality guarantees

Quality guarantees allow system designer to evaluate different design alternatives: configuration selection

e.g. in sensor networks

We can select a placement for sensors

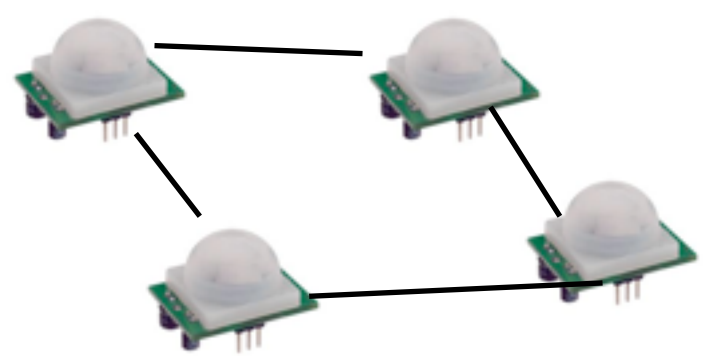

50%

Suboptimal coordination with quality guarantees: approaches

- Region optimal algorithms [AAMAS, 2011]

Suboptimal coordination with quality guarantees: approaches

A solution is region optimal when its value cannot be improved by changing the decision of any group of agents in the region

Suboptimal coordination with quality guarantees: approaches

A solution is region optimal when its value cannot be improved by changing the decision of any group of agents in the region

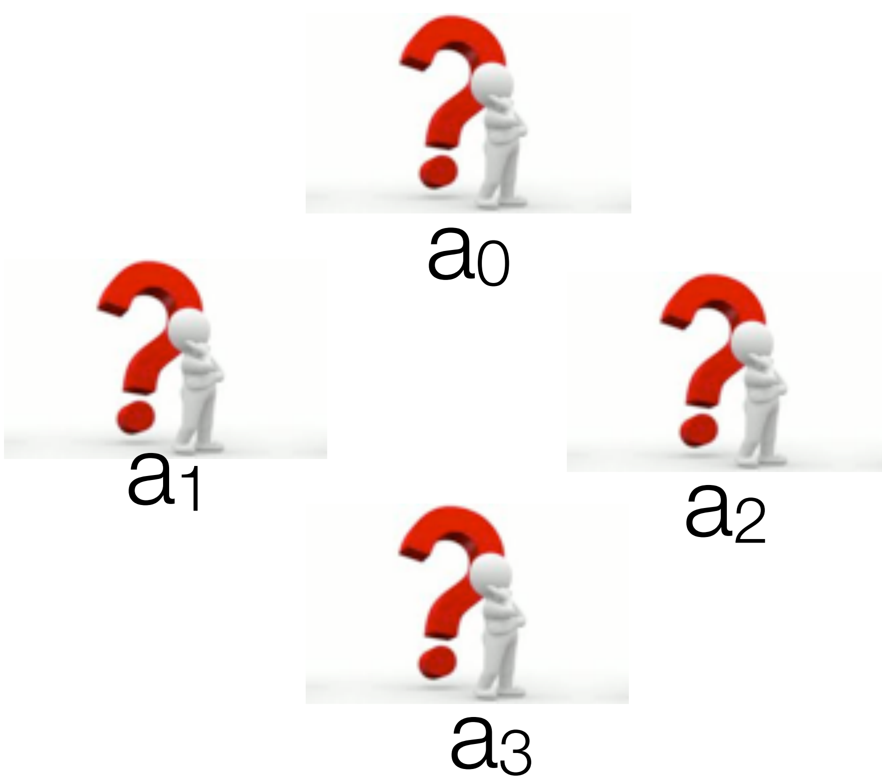

Suboptimal coordination with quality guarantees: approaches

A solution is region optimal when its value cannot be improved by changing the decision of any group of agents in the region

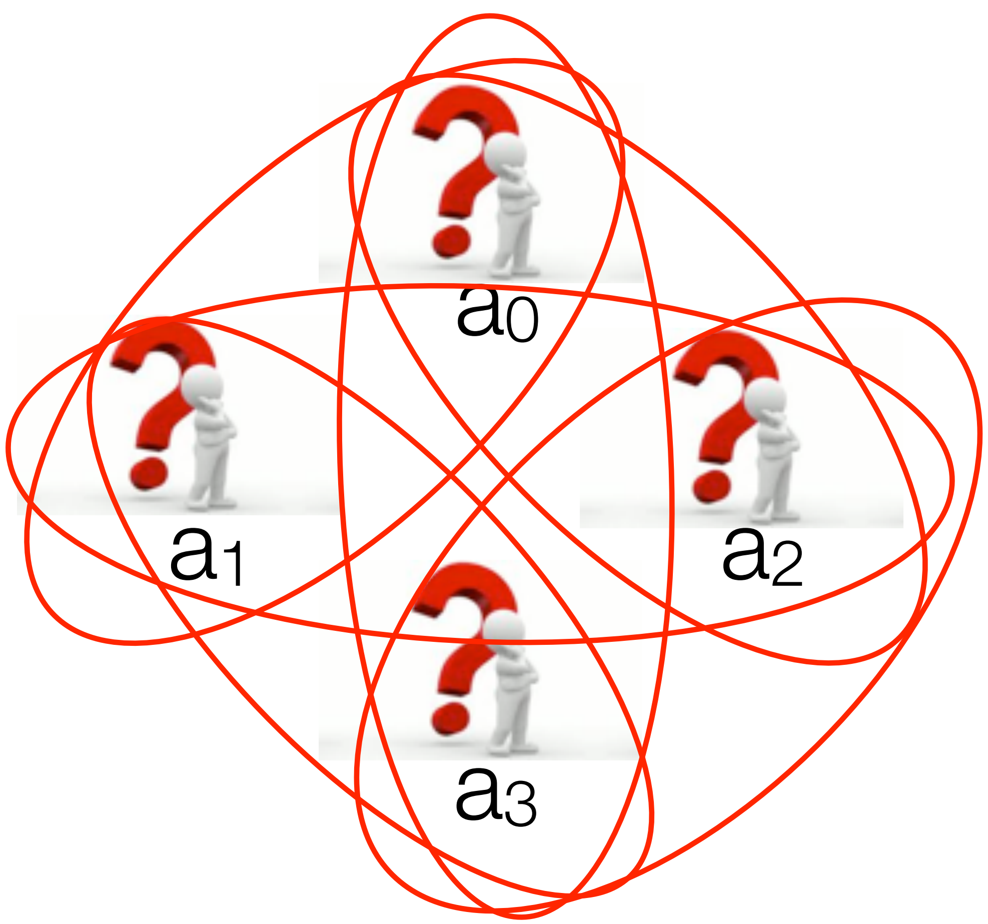

Groups in the region

$\{a_0, a_1\}$ $\{a_1, a_2\}$ $\{a_2, a_3\}$

$\{a_0, a_2\}$ $\{a_1, a_3\}$

$\{a_0, a_3\}$

Suboptimal coordination with quality guarantees: approaches

A solution is region optimal when its value cannot be improved by changing the decision of any group of agents in the region

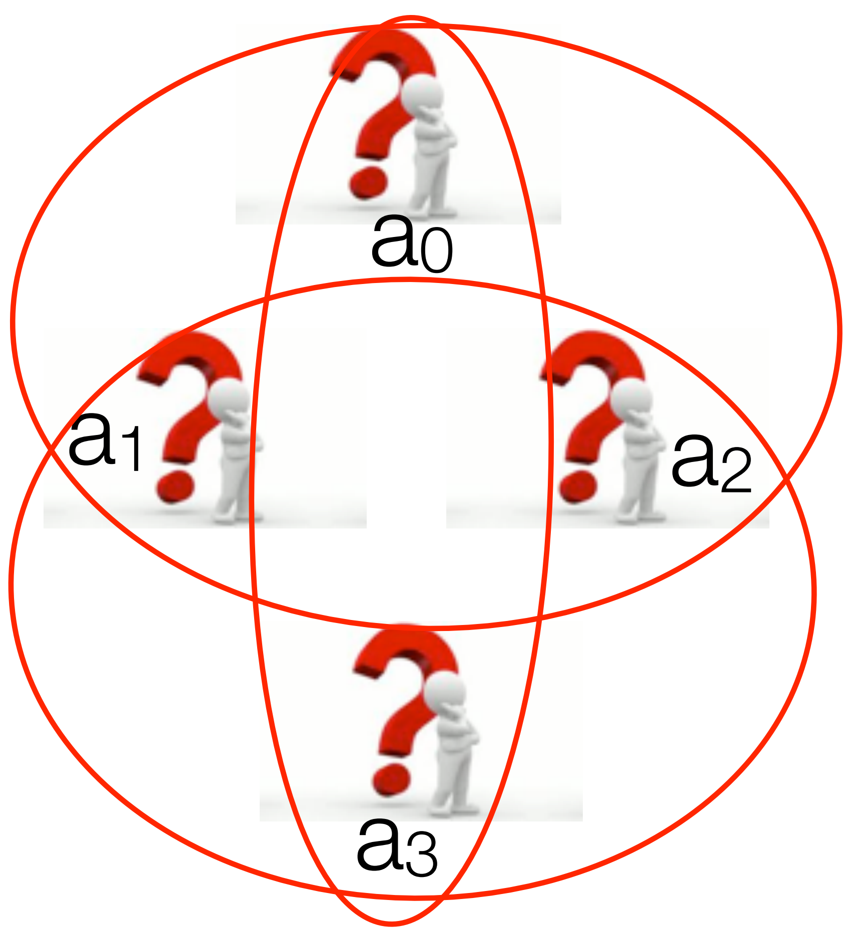

Groups in the region

$\{a_0, a_1, a_2\}$

$\{a_1, a_2, a_3\}$

$\{a_0, a_3\}$

Suboptimal coordination with quality guarantees: approaches

Region optimality [AAMAS, 2011] allows to assess quality guarantees for any region optimal

Suboptimal coordination with quality guarantees: approaches

Region optimality [AAMAS, 2011] allows to assess quality guarantees for any region optimal

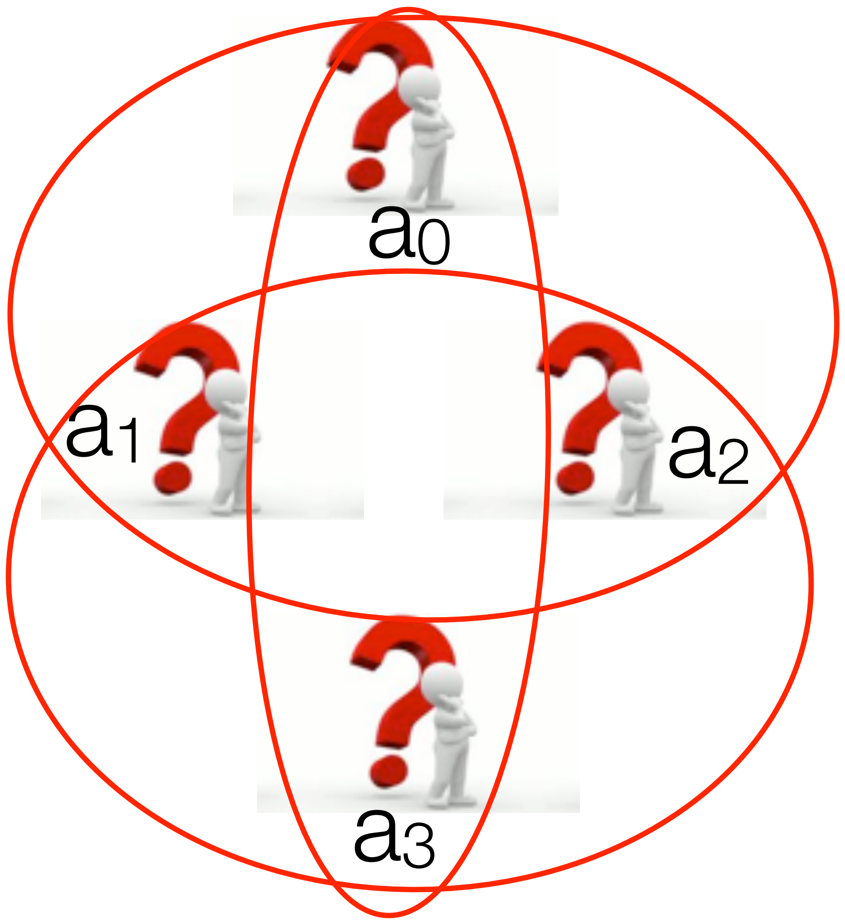

The quality of any region optimal in this region in any problem is guaranteed to be at least 33% the value of the optimal solution

Suboptimal coordination with quality guarantees: approaches

Region optimality [AAMAS, 2011] allows to assess quality guarantees for any region optimal

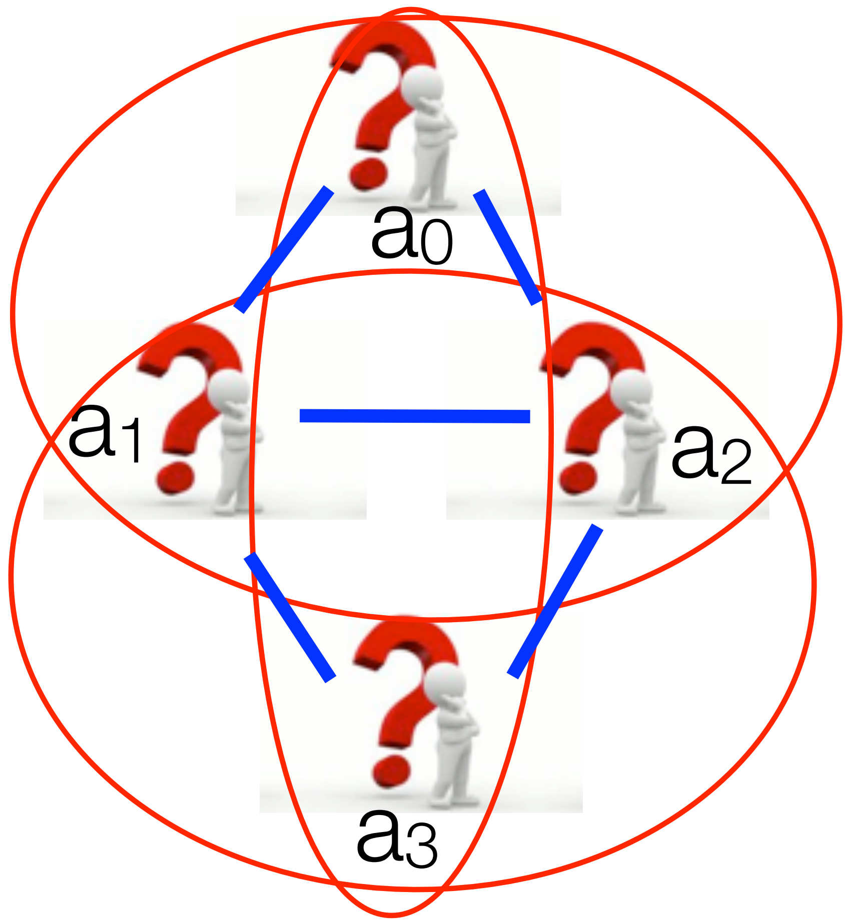

The quality of any region optimal in this region in any problem with this dependency graph is guaranteed to be at least 50% the value of the optimal solution

Suboptimal coordination with quality guarantees: approaches

Region optimality [AAMAS, 2011] allows to assess quality guarantees for any region optimal

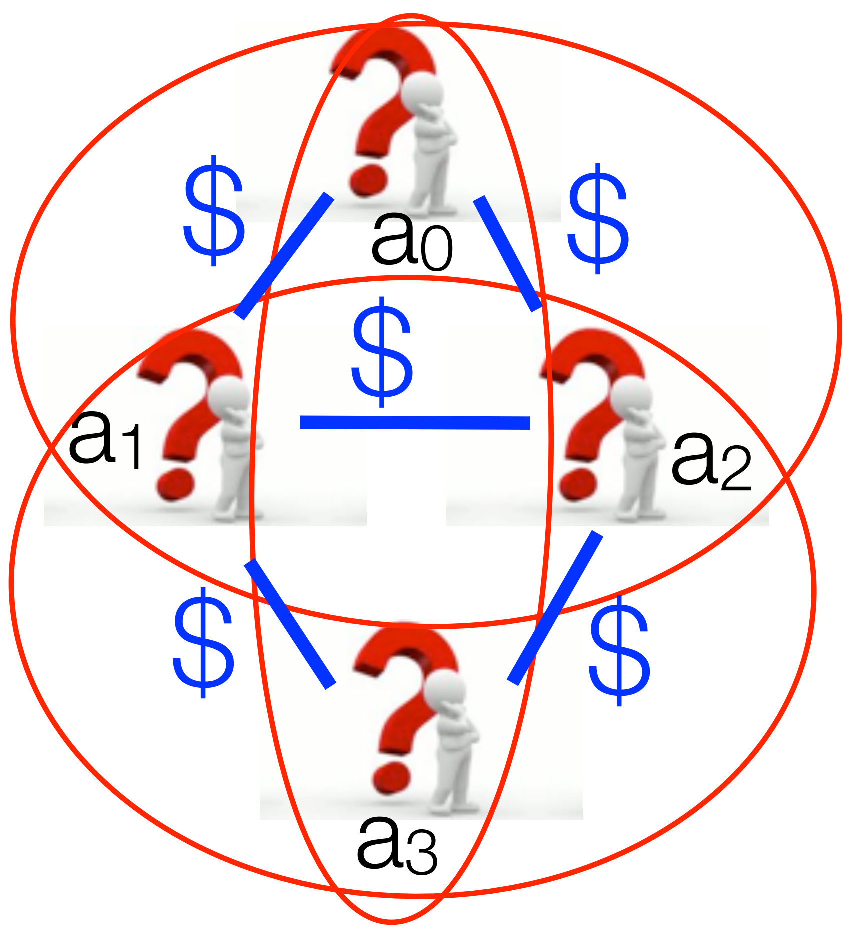

The quality of any region optimal in this region in any problem with this dependency graph and reward structure is guaranteed to be at least 75% the value of the optimal solution

Suboptimal coordination with quality guarantees: approaches

So far we have characterized quality guarantees
for region optimal solutions but ...
how agents find such region optimal solutions?

Suboptimal coordination with quality guarantees: approaches

So far we have characterized quality guarantees for region optimal solutions but ...

how agents find such region optimal solutions?

Suboptimal coordination with quality guarantees: approaches

A generic region optimal algorithm

Suboptimal coordination with quality guarantees: approaches

A generic region optimal algorithm

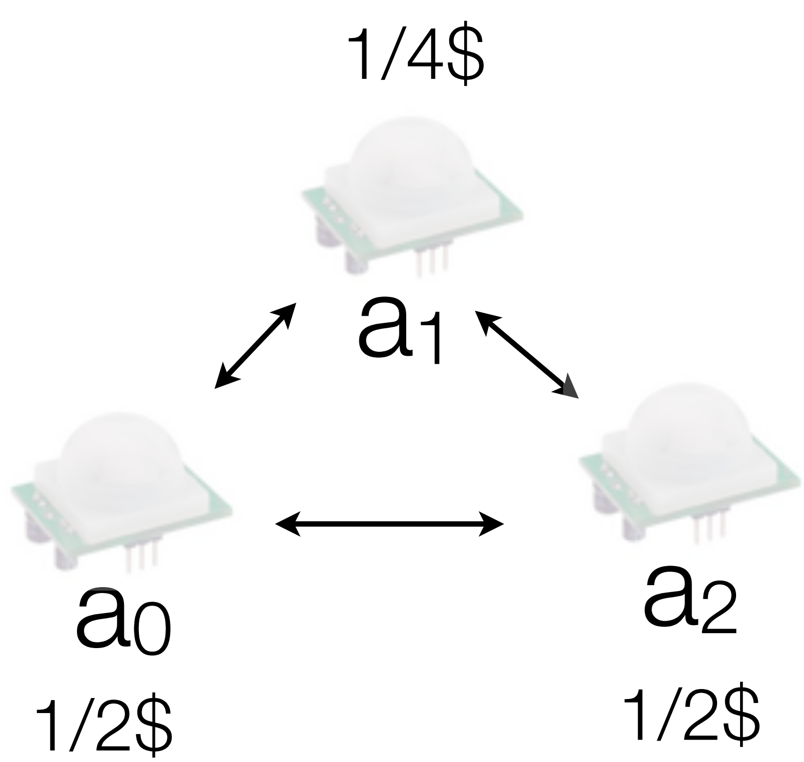

Agents select an initial action

Suboptimal coordination with quality guarantees: approaches

A generic region optimal algorithm

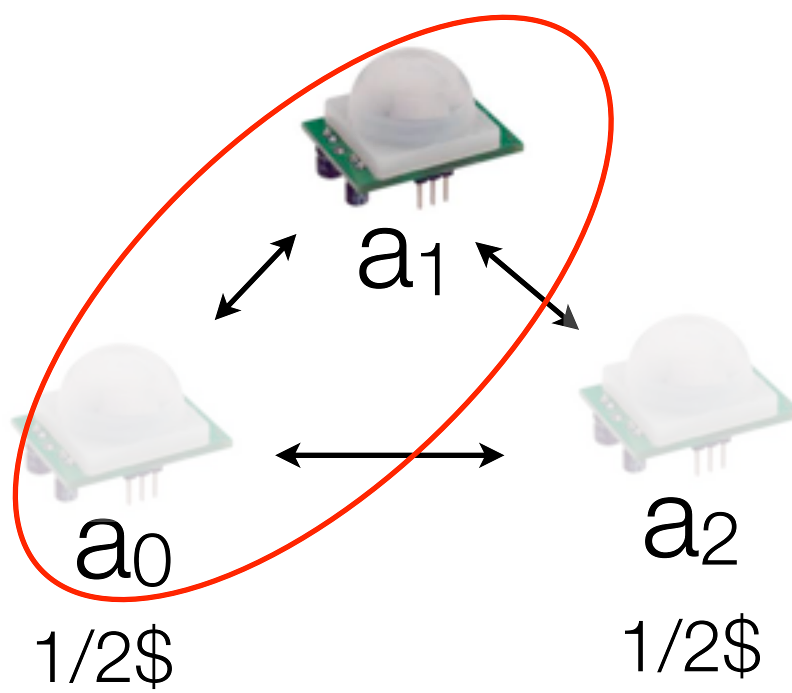

Each group of agents in the region optimizes its decision given other agents decisions.

Suboptimal coordination with quality guarantees: approaches

A generic region optimal algorithm

Initialization → Optimization ⇄ Implementation

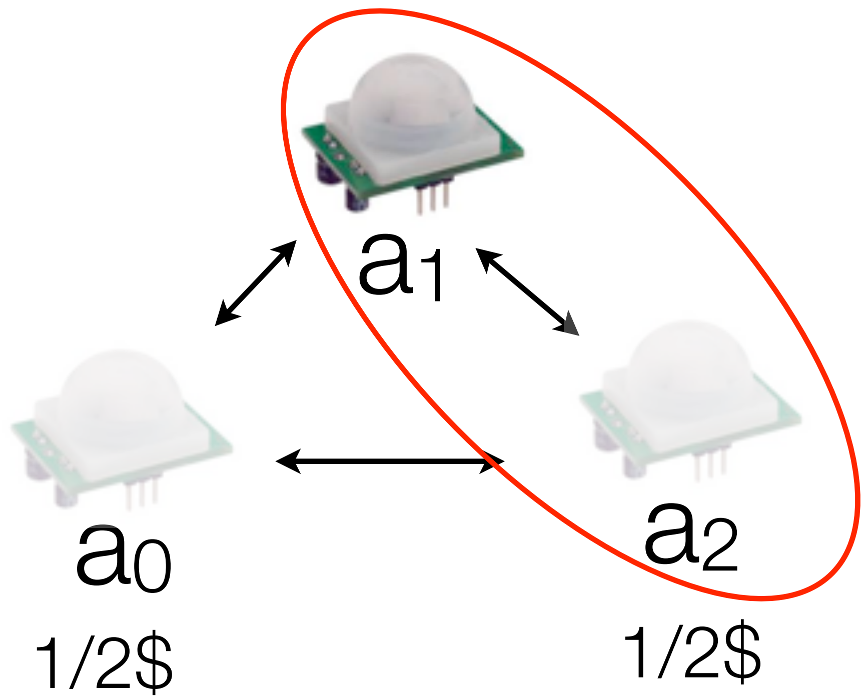

Suboptimal coordination with quality guarantees: approaches

A generic region optimal algorithm

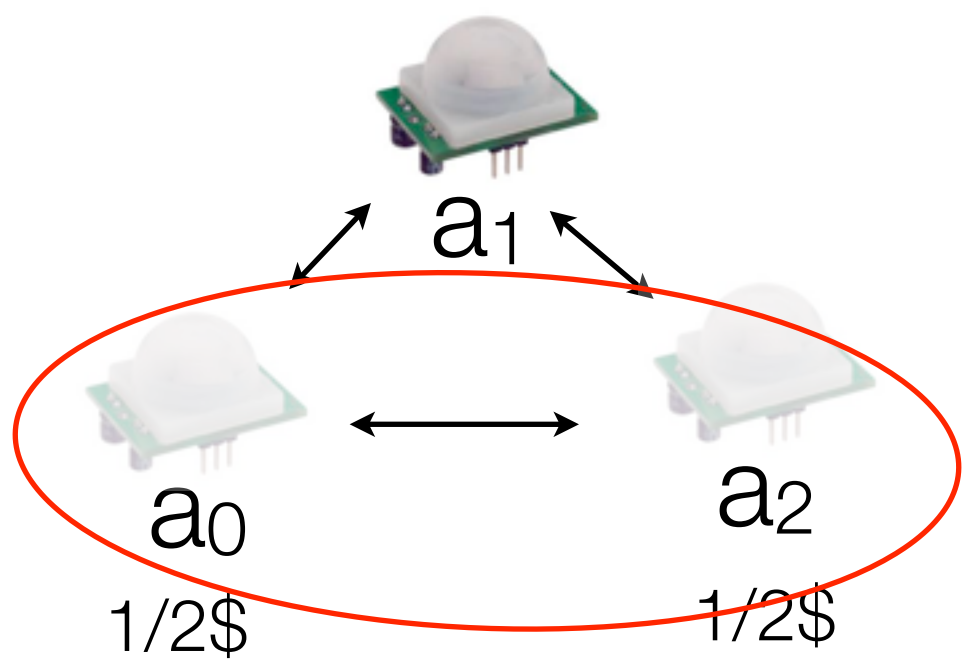

Suboptimal coordination with quality guarantees: approaches

A generic region optimal algorithm

Initialization → Optimization ⇄ Implementation

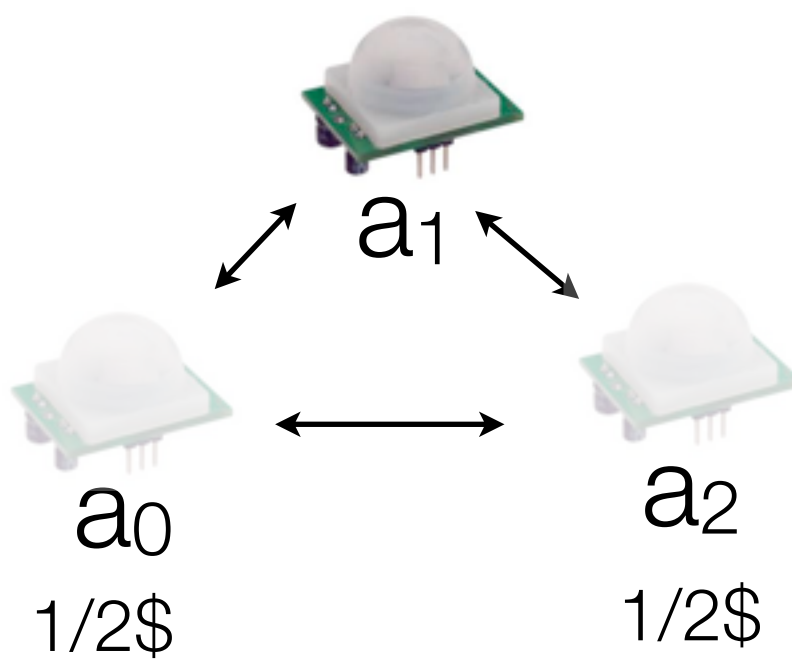

The quality of any region optimal in this region in any problem is guaranteed to be at least 50% the value of the optimal solution

... until stabilitzation

Message to take away

- Many real-world problems can be modeled as a network of agents that need to coordinate their actions to optimize system performance

- Optimality is not affordable in many of these emerging large-scale domains

- An open line of research is how to design suboptimal algorithms that provide quality guarantees over the agent's actions

Gracias por vuestra atención!!!