CP 2024

Ex-Ante Constraint Elicitation in Incomplete DCOP

Roie Zivan, Shiraz Regev and William Yeoh

Ben-Gurion University of the Negev

Washington University in St. Louis

Overview

Distributed Multi-Agent Systems:

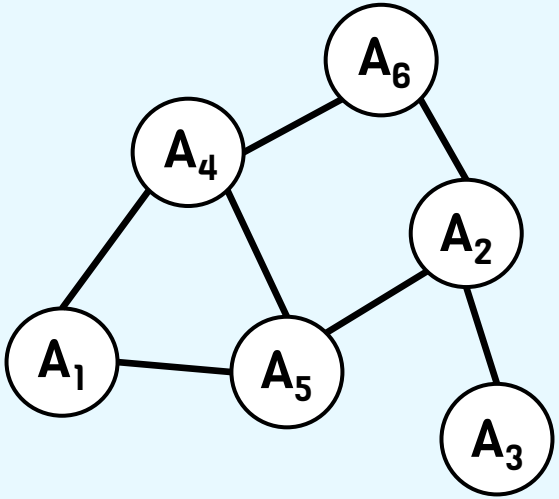

Agents cooperate and communicate

Mutual goal

Overview

Distributed Multi-Agent Systems:

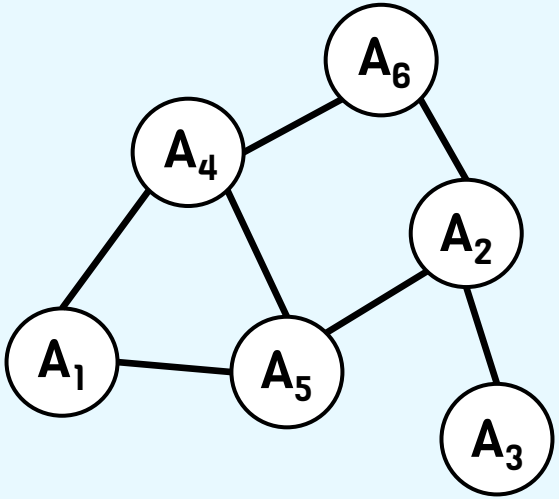

Meeting Scheduling Problem

Agents cooperate and communicate

Mutual goal

Overview

Distributed Multi-Agent Systems:

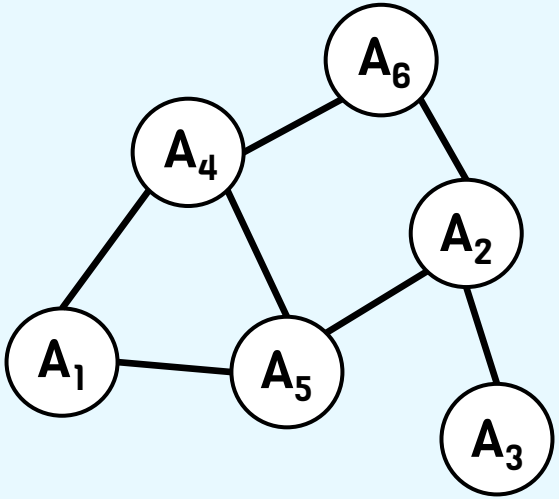

Application	Model
Users	Agents
Meetings	Variables
Possible time slots	Domain
Coordination of time slots that align with others	Constraints and costs

Agents cooperate and communicate

Mutual goal

DCOP

Overview

Not all preferences are known

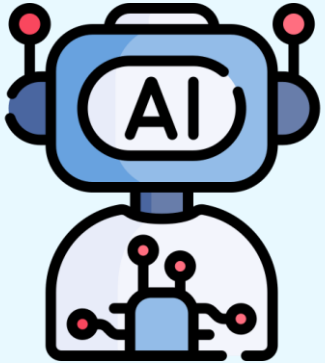

Agents cooperate and communicate

Mutual goal

DCOP

I-DCOP

Overview

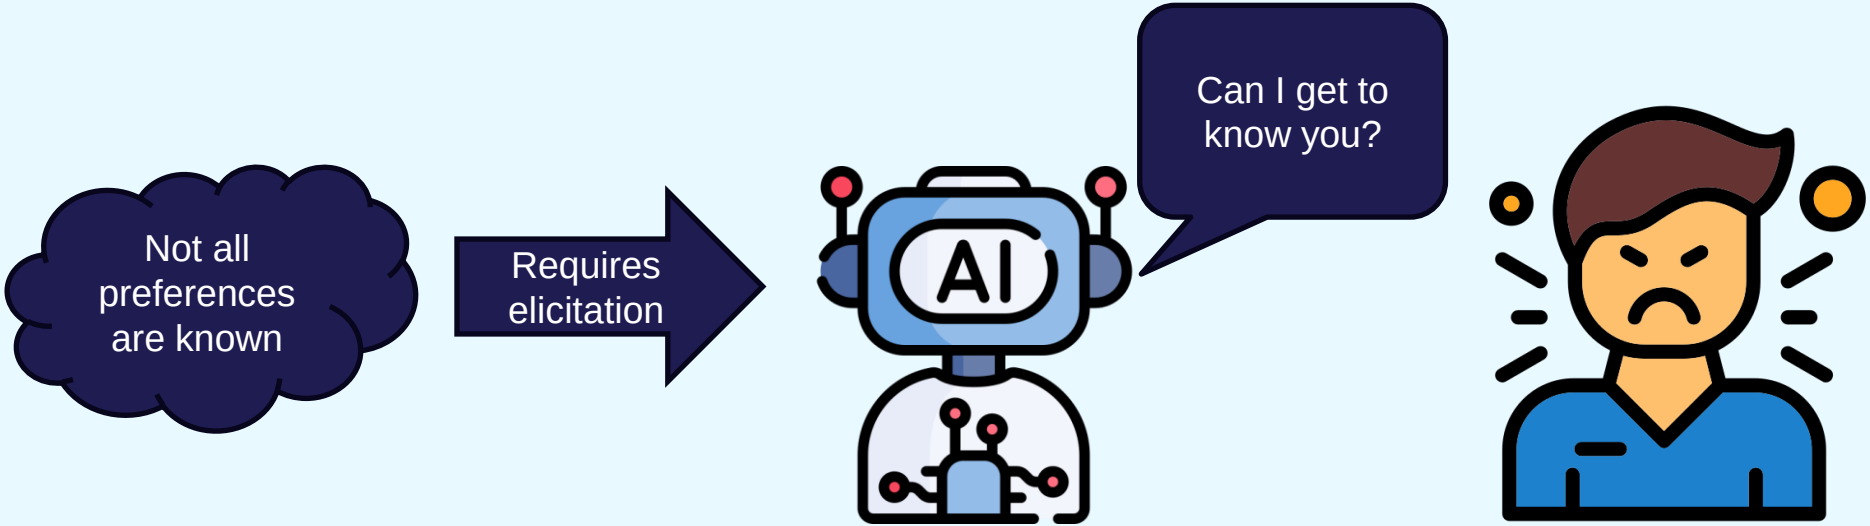

Agents cooperate and communicate

Mutual goal

DCOP

I-DCOP

Overview

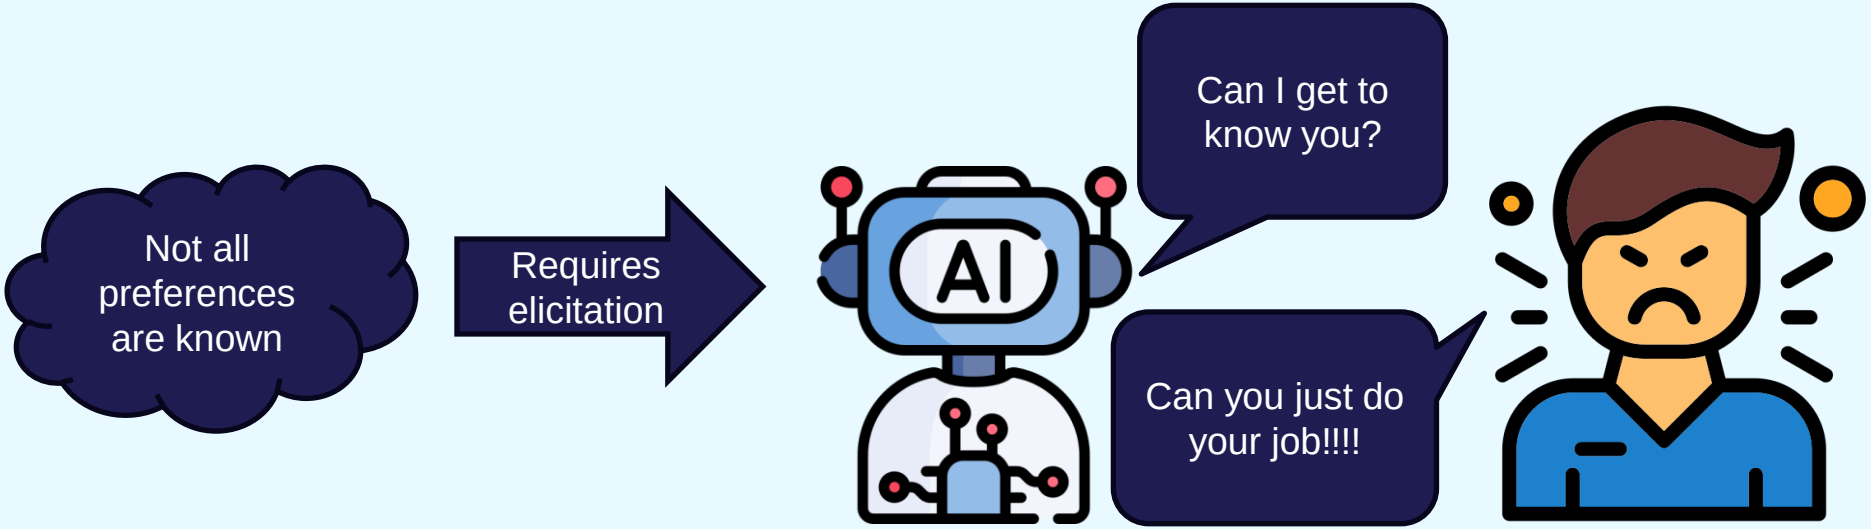

Agents cooperate and communicate

Mutual goal

DCOP

I-DCOP

Overview

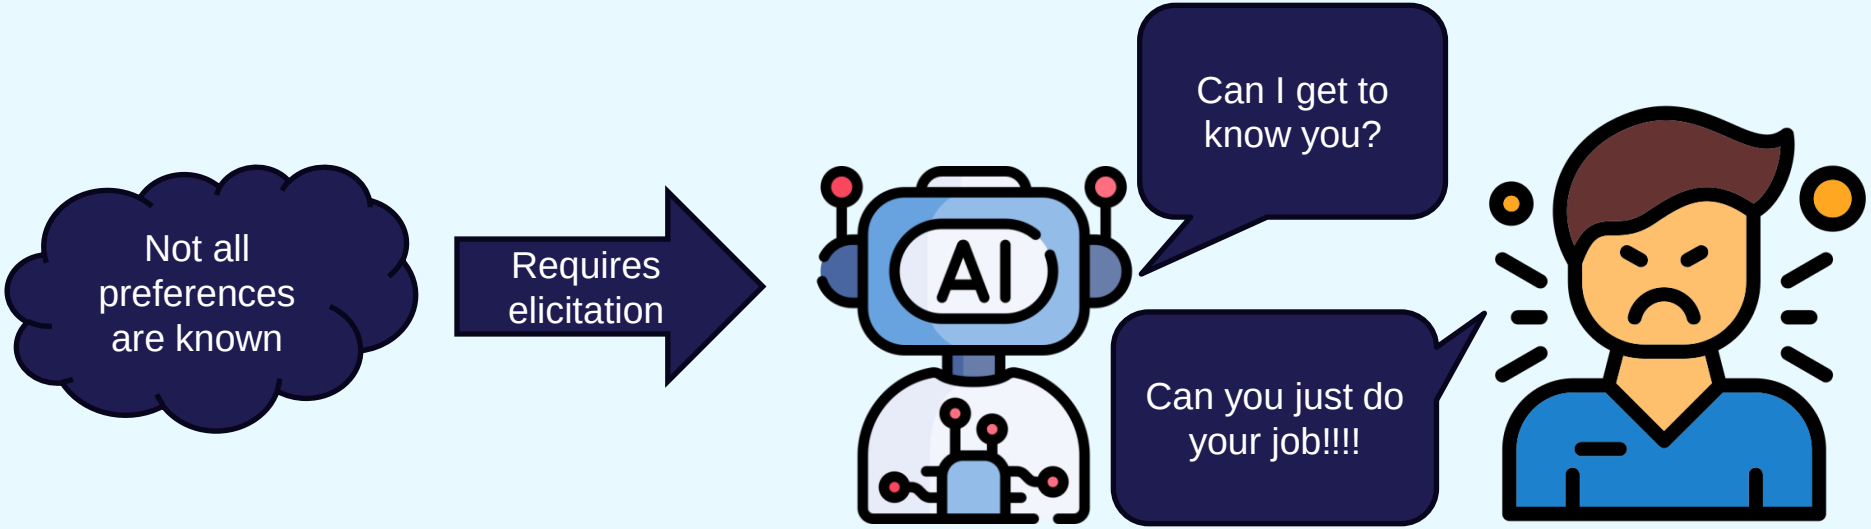

Agents cooperate and communicate

Mutual goal

DCOP

I-DCOP

Related Work

Tabakhi et al. (DAI 2021)

Ex-post

Past Approach

Include the cost of revealing information when deciding on an assignment.

Our Approach

Ask whether revealing information is worth it throughout the execution of the algorithm.

Related Work

Tabakhi et al. (DAI 2021)

Ex-post

Past Approach

Include the cost of revealing information when deciding on an assignment.

Ex-ante

Our Approach

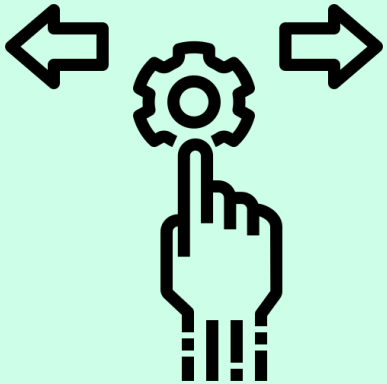

Ask whether revealing information is worth it throughout the execution of the algorithm.

Related Work

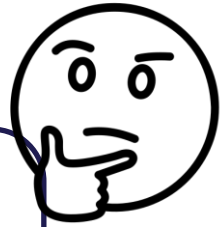

After found the hotel: “I spent a lot of time.” I won’t take it

	Happiness (Utility)	Time spent (Elicitation)
Option 1		
Option 2		

“Is the result worth the time?”

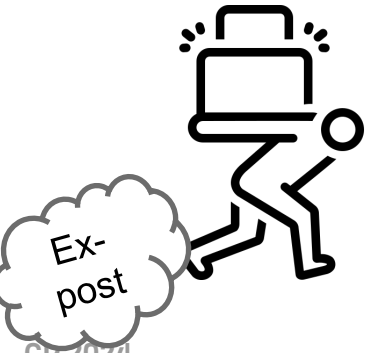

Related Work

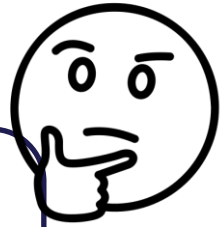

After found the hotel: “I spent a lot of time.” I won’t take it

	Happiness (Utility)	Time spent (Elicitation)
Option 1		
Option 2		

“Is the result worth the time?”

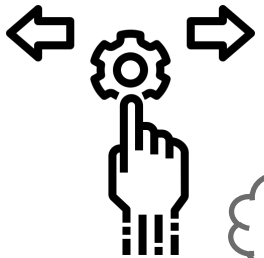

Background

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1, \ldots, A_n\}$

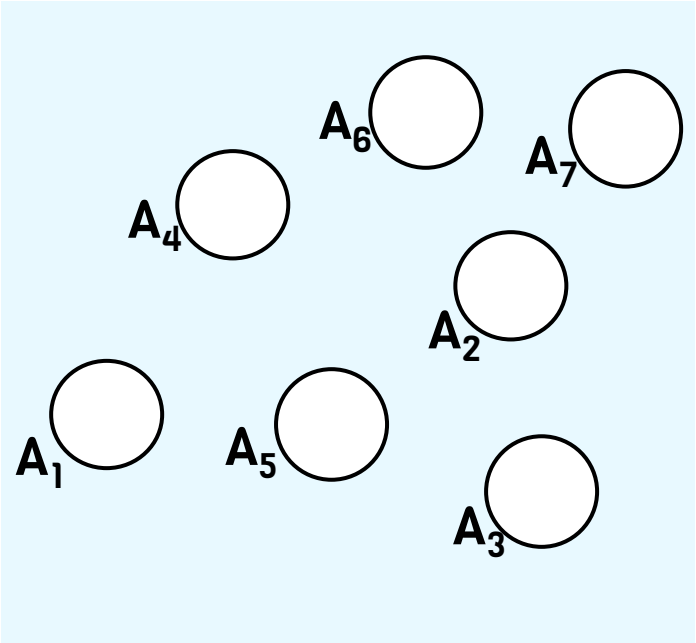

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1, \ldots, A_n\}$

X - Variables $\{X_1, \ldots, X_m\}$

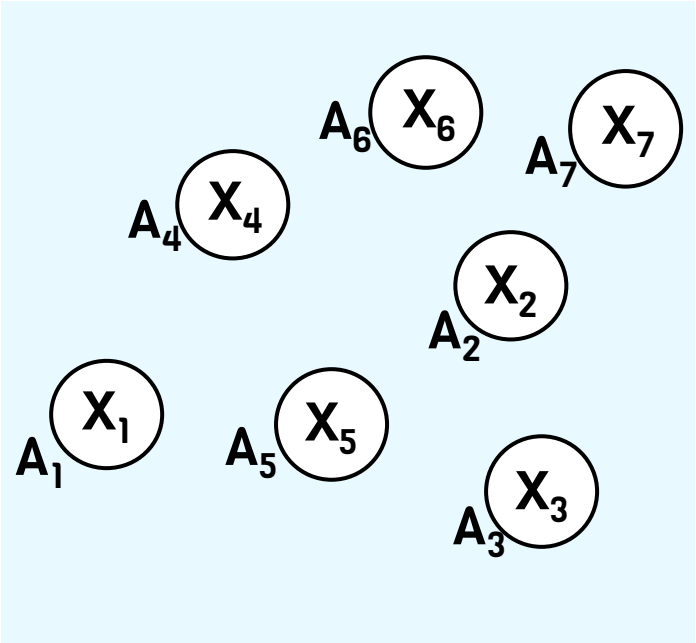

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1, \ldots, A_n\}$

X - Variables $\{X_1, \ldots, X_m\}$

D - Domain $\{D_1, \ldots, D_m\}$

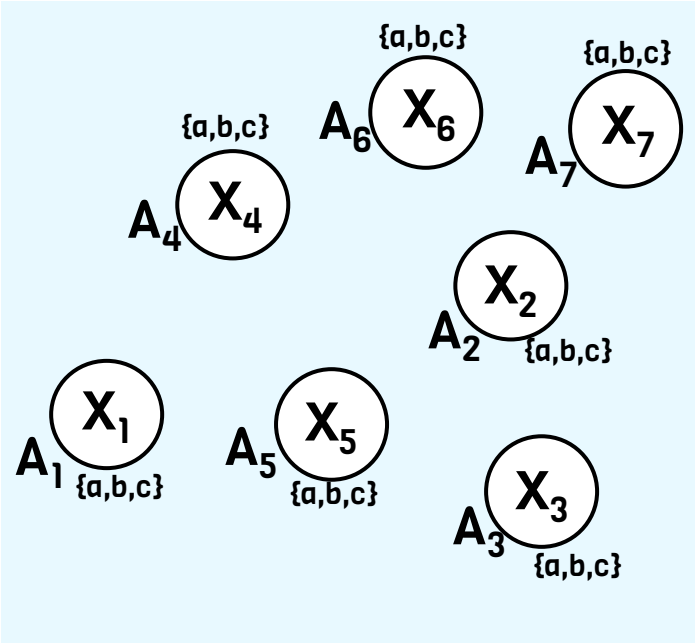

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1,\dots, A_n\}$

X - Variables $\{X_1,\dots, X_m\}$

D - Domain $\{D_1,\dots, D_m\}$

C- Constraints

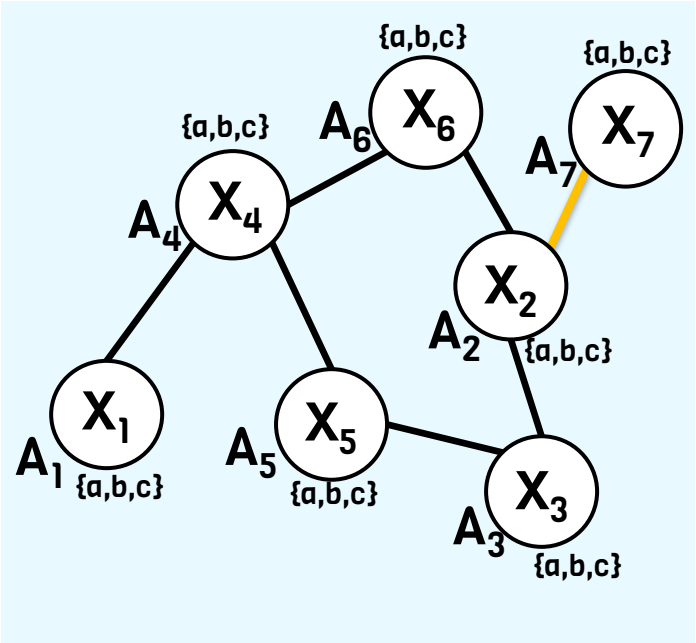

$C_{2,7}$ - Cost Table

X_2/ X_7	a	b	c
a	8	3	2
b	1	5	4
c	6	2	4

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1, \ldots, A_n\}$

X - Variables $\{X_1, \ldots, X_m\}$

D - Domain $\{D_1, \ldots, D_m\}$

C- Constraints

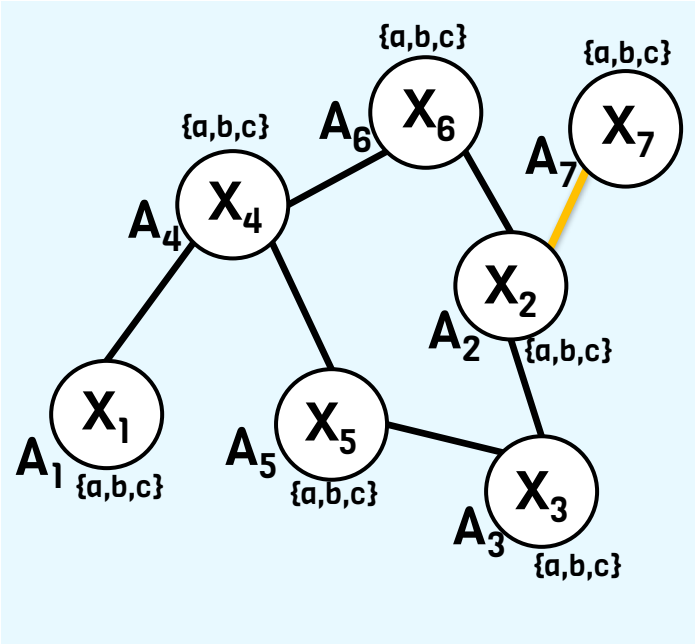

$C_{2,7}$ - Cost Table

X_2/ X_7	a	b	c
a	8	3	2
b	1	5	4
c	6	2	4

Goal:
finding a complete assignment with minimal global cost

Incomplete

Distributed Constraint Optimization Problem (DCOP)

A DCOP is a tuple:

A - Agents $\{A_1, \ldots, A_n\}$

X - Variables $\{X_1, \ldots, X_m\}$

D - Domain $\{D_1, \ldots, D_m\}$

~~C- Constraints~~

C- Partially Specified Constraint

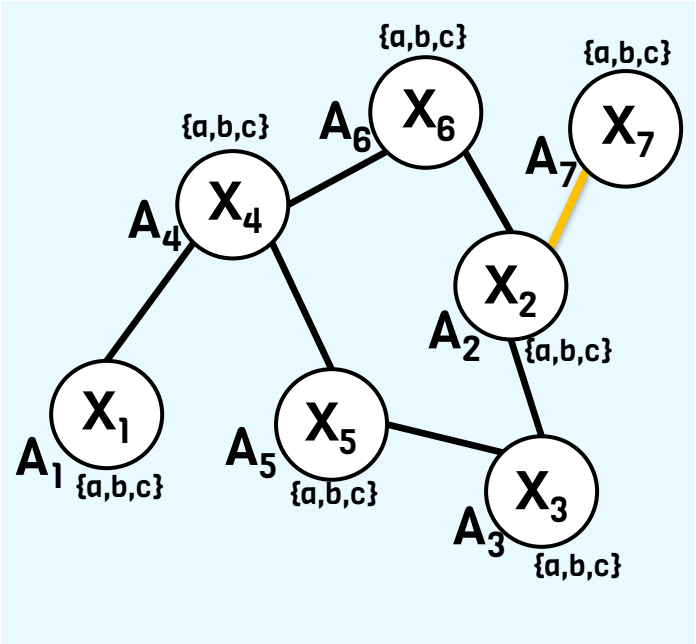

Some of the costs are unknown

$C_{2,7}$ - Cost Table

X_2/ X_7	a	b	c
a	8	3	2
b	?	5	4
c	6	2	?

Goal:
finding a complete assignment with minimal global cost

Incomplete

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1, \ldots, A_n\}$

X - Variables $\{X_1, \ldots, X_m\}$

D - Domain $\{D_1, \ldots, D_m\}$

C- Partially Specified Constraint

E- Elicitation Cost

B – Budget (global?)

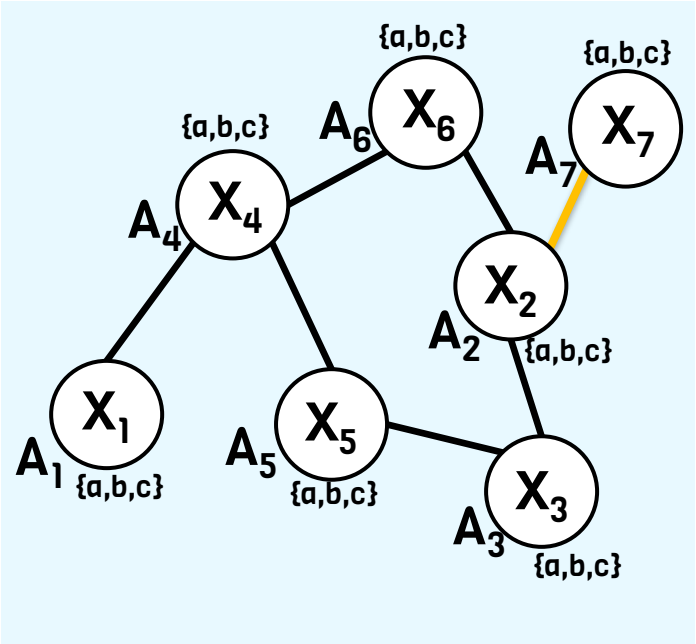

$C_{2,7}$ - Cost Table

X_2/ X_7	a	b	c
a	8	3	2
b	?	5	4
c	6	2	?

$E_{2,7}$ - Elicitation Table

X_2/ X_7	a	b	c
a	0	0	0
b	15	0	0
c	0	0	10

Incomplete

Distributed Constraint Optimization Problem (DCOP)

DCOP is a tuple:

A - Agents $\{A_1,\dots, A_n\}$

X - Variables $\{X_1,\dots, X_m\}$

D - Domain $\{D_1,\dots, D_m\}$

C- Partially Specified Constraint

E- Elicitation Cost

B – Budget (global?)

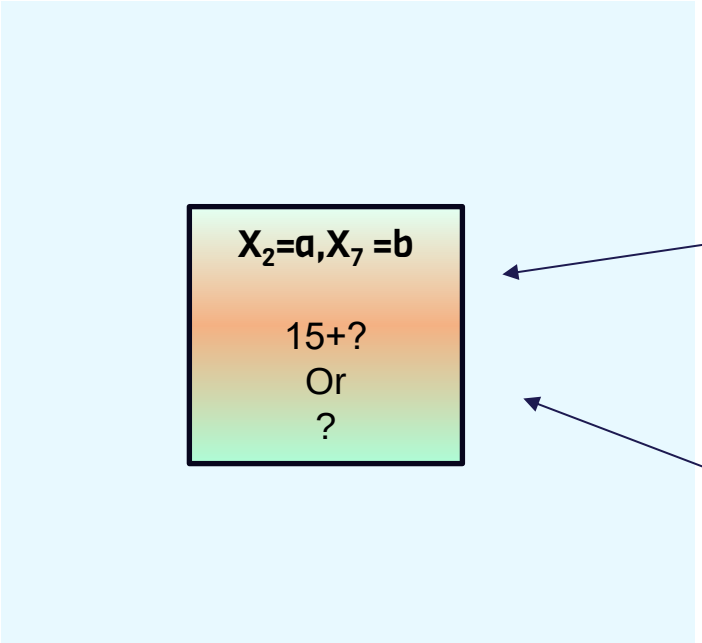

$C_{2,7}$ - Cost Table

X_2/ X_7	a	b	c
a	8	3	2
b	?	5	4
c	6	2	?

$E_{2,7}$ - Elicitation Table

X_2/ X_7	a	b	c
a	0	0	0
b	15	0	0
c	0	0	10

Solving DCOP

Complete

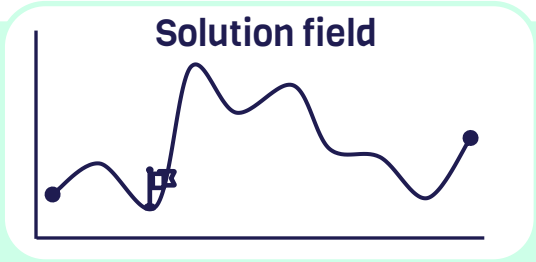

Solving DCOP

Complete

Synchronous Branch and Bound

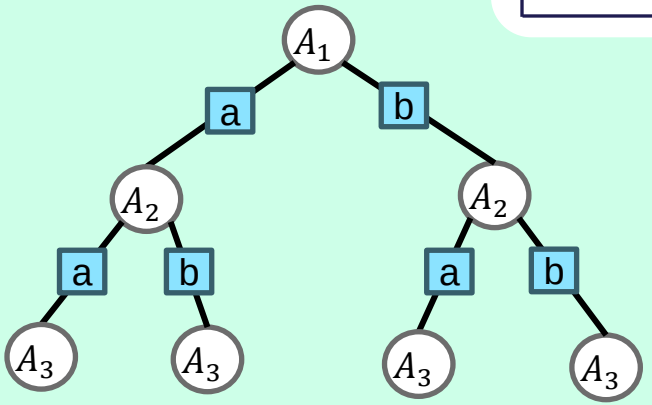

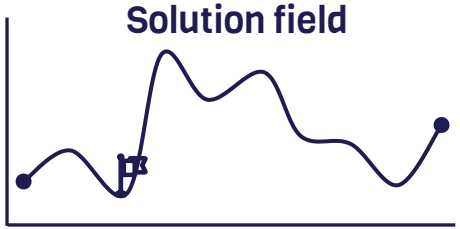

Solving DCOP

Complete

Synchronous Branch and Bound

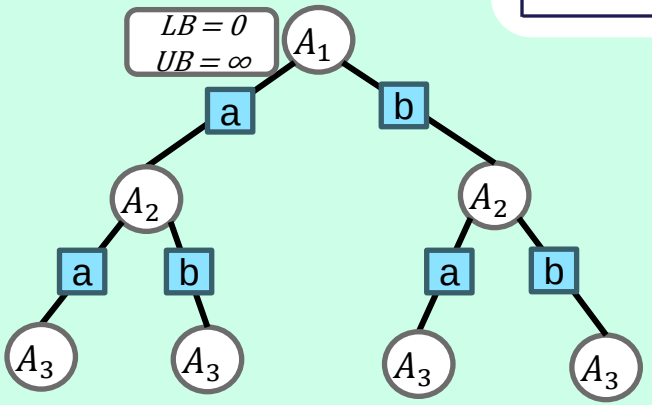

CPA is passed in a pre-set order with partial assignment and aggregated cost as long as LB<UB

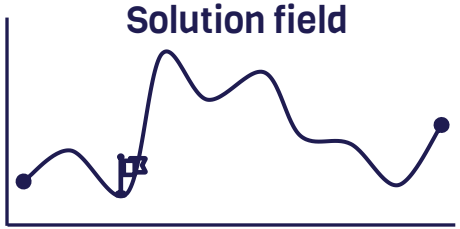

Solving DCOP

Complete

Synchronous Branch and Bound

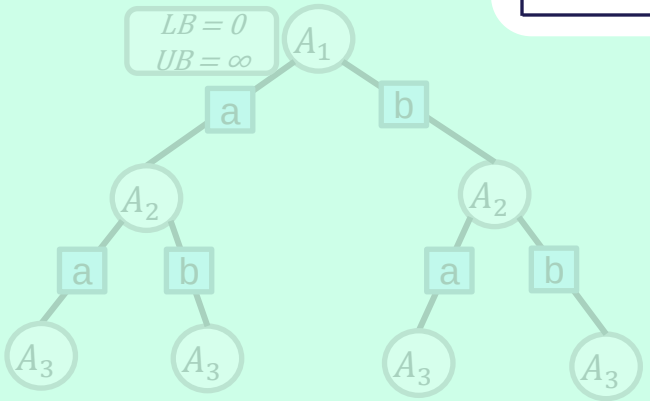

CPA is passed in a pre-set order with partial assignment and aggregated cost as long as LB<UB

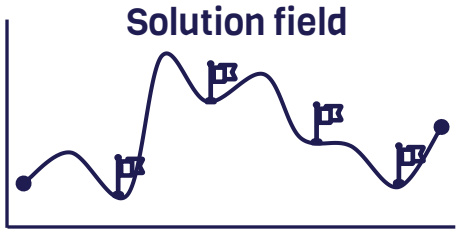

Incomplete

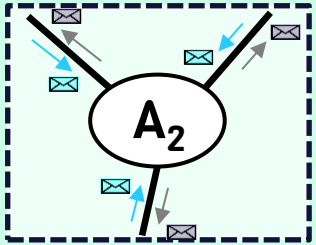

Agent sends and collects its neighbors' value assignment.

Solving DCOP

Complete

Synchronous Branch and Bound

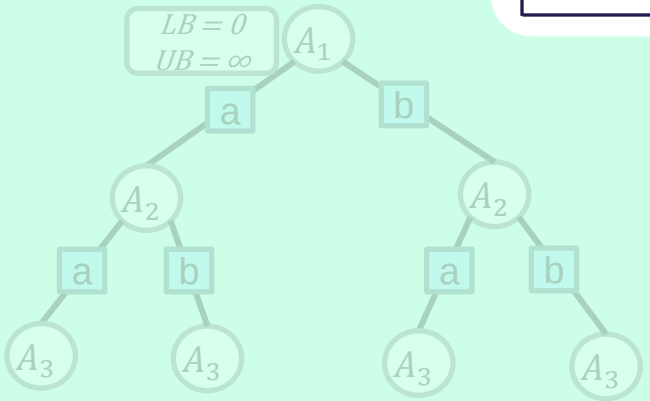

CPA is passed in a pre-set order with partial assignment and aggregated cost as long as LB<UB

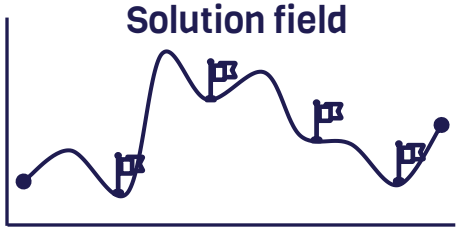

Incomplete

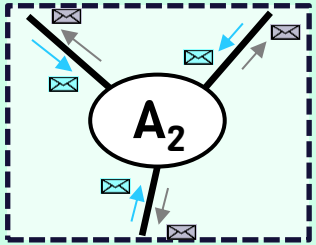

MGM

Agents exchanging assignments and gains; only the highest gain results in replacement.

Agent sends and collects its neighbors' value assignment.

Ex-Ante I DCOP

Main Idea

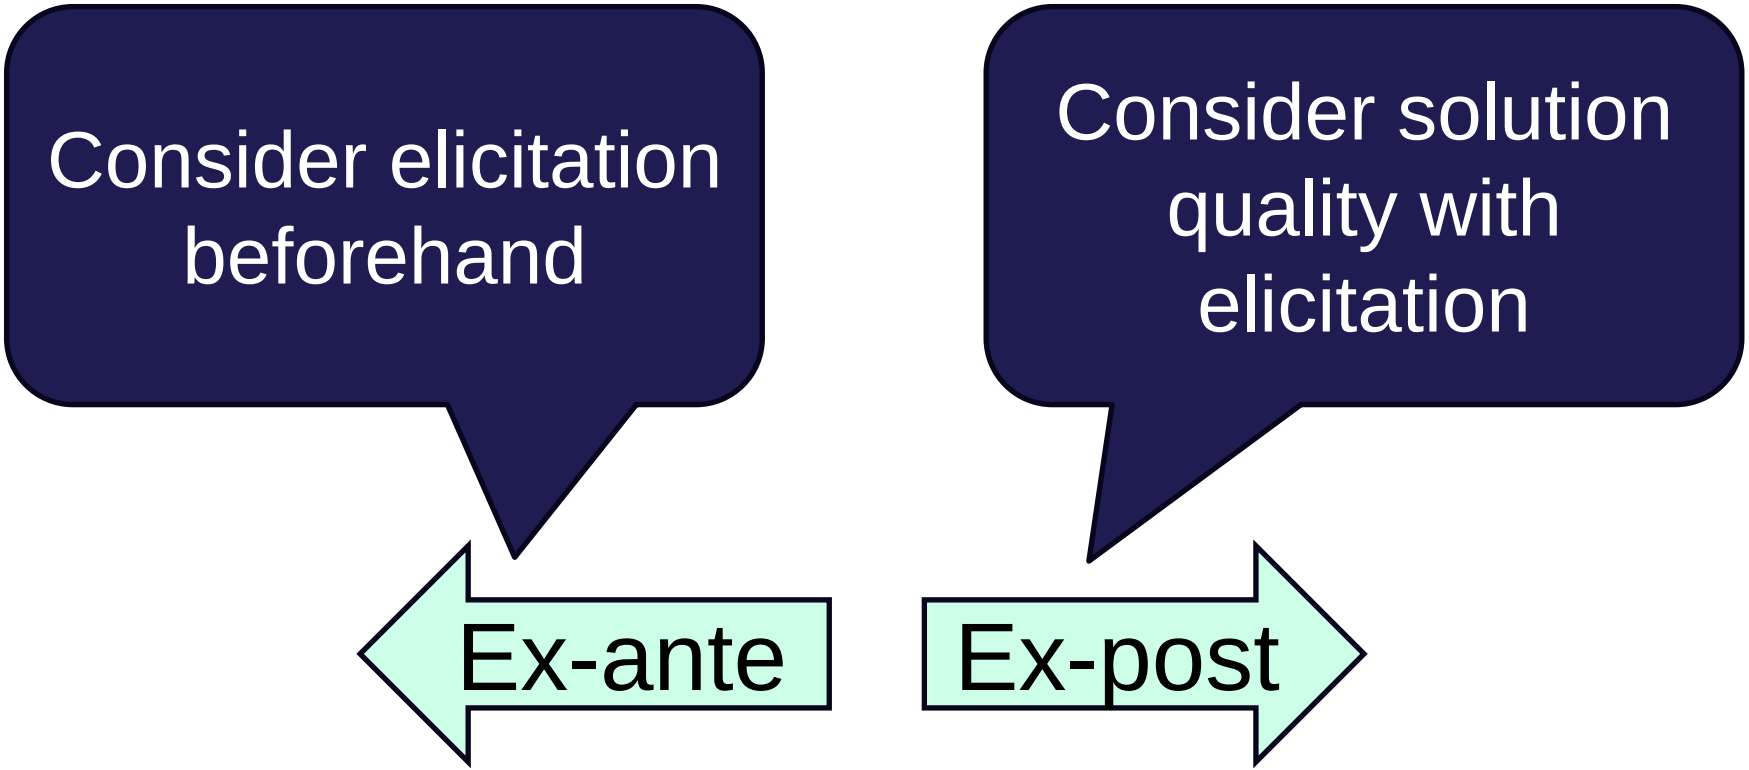
Consider elicitation beforehand
Consider solution quality with elicitation
Ex-ante
Ex-post

Ex-Ante I DCOP

Complete

Synchronous Branch and Bound

Incomplete

DSA and MGM

Ex-Ante I DCOP

A_2 Decides which assignment to select?

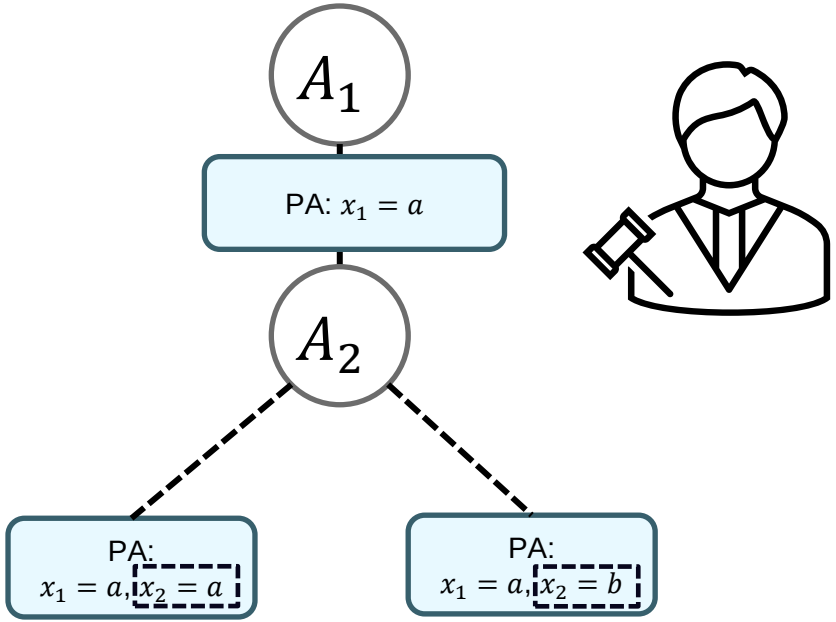

Ex-Ante I DCOP

$E(x_1 = a, x_2 = a) = 0$
$E(x_1 = a, x_2 = b) > 0$

Ex-Ante I DCOP

Select assignment without elicitation cost

$E(x_1 = a, x_2 = a) = 0$
$E(x_1 = a, x_2 = b) > 0$

Pay?
$E(x_1 = a, x_2 = b) > 0$

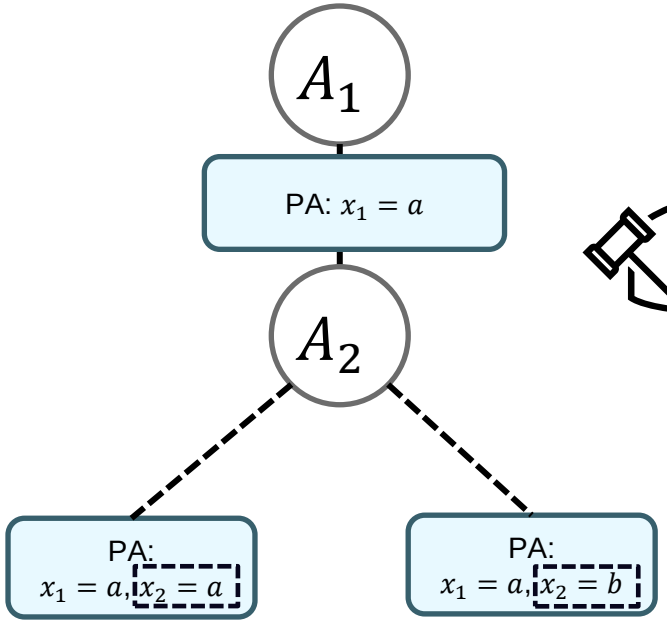

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

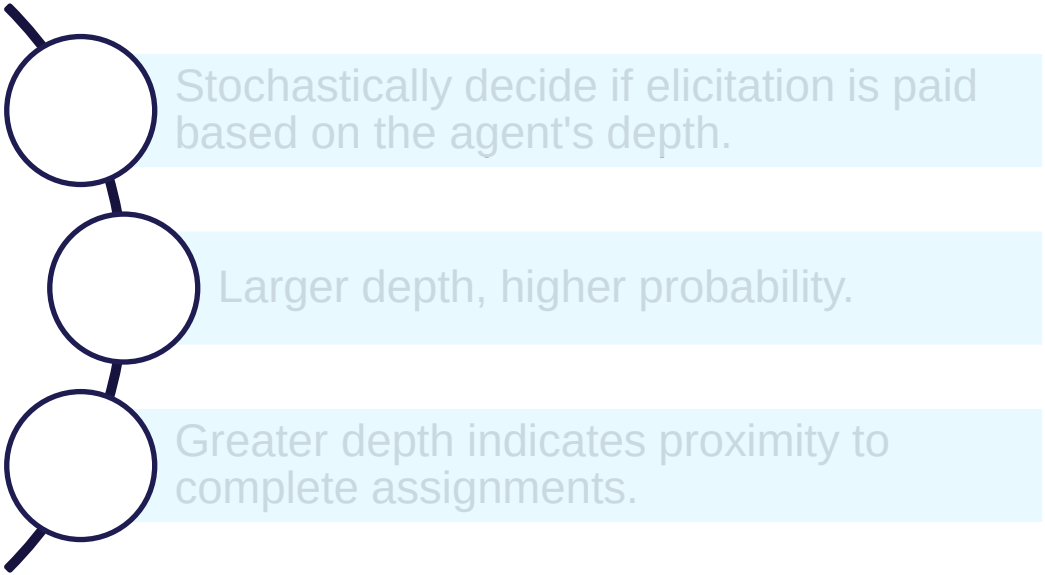

Stochastically decide if elicitation is paid based on the agent's depth.

Larger depth, higher probability.

Greater depth indicates proximity to complete assignments.

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Stochastically decide if elicitation is paid based on the agent's depth.

Larger depth, higher probability.

Greater depth indicates proximity to complete assignments.

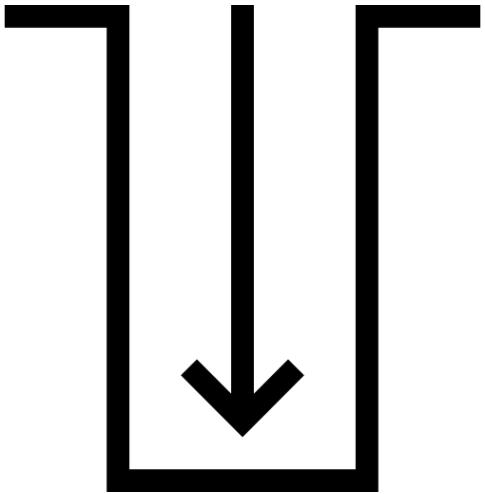

Heuristics

Ex-ante

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Stochastically decide if elicitation is paid based on the agent's depth.

Larger depth, higher probability.

Greater depth indicates proximity to complete assignments.

$$p(i) = \frac{e^{i}}{e^{n/2} + e^{i}}$$

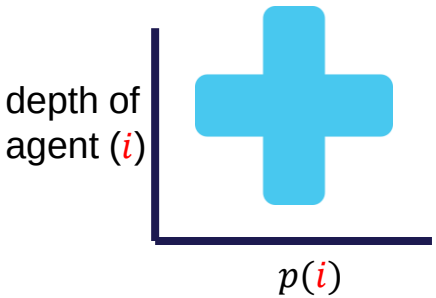

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Stochastically decide if elicitation is paid based on the agent's depth.

Larger depth, higher probability.

Greater depth indicates proximity to complete assignments.

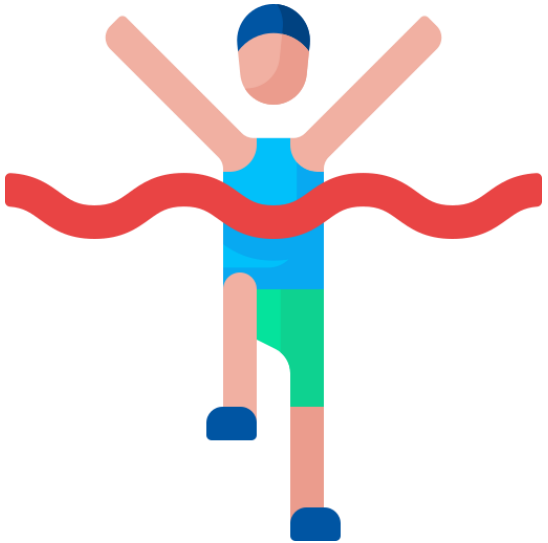

Heuristics

Ex-ante

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

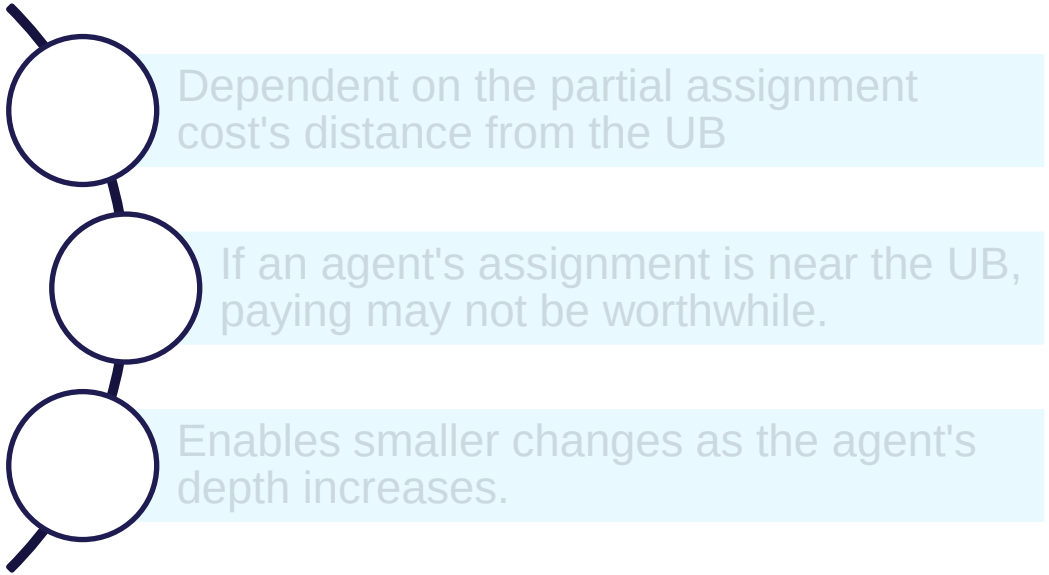

Dependent on the partial assignment cost's distance from the UB

If an agent's assignment is near the UB, paying may not be worthwhile.

Enables smaller changes as the agent's depth increases.

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

- Dependent on the partial assignment cost's distance from the UB
- If an agent's assignment is near the UB, paying may not be worthwhile.
- Enables smaller changes as the agent's depth increases.

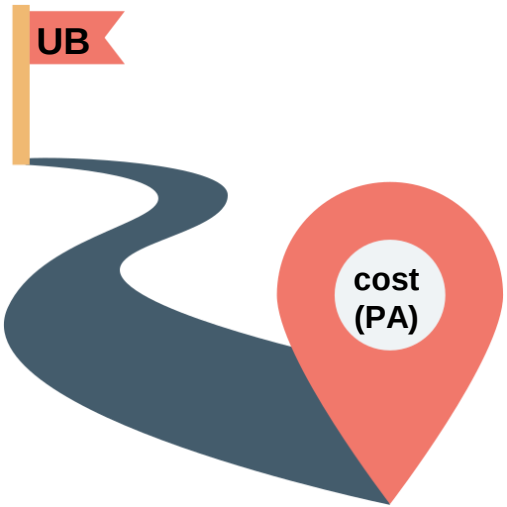

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Dependent on the partial assignment cost's distance from the UB

If an agent's assignment is near the UB, paying may not be worthwhile.

Enables smaller changes as the agent's depth increases.

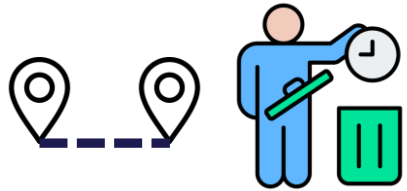

Heuristics

Ex-ante

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Dependent on the partial assignment cost's distance from the UB

If an agent's assignment is near the UB, paying may not be worthwhile.

Enables smaller changes as the agent's depth increases.

Agent calculates:

$$t(i) = g * (1 - \frac{e^i}{e^{\frac{n}{2}} + e^i})$$

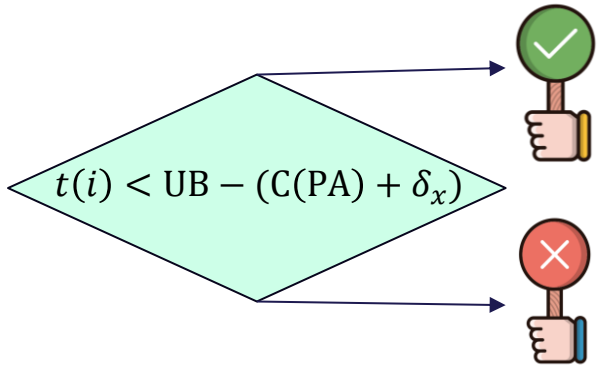

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

- Pre-process count amount of unknown costs throughout the solution space
- Prioritize subtrees where you have more knowledge.
- When an agent elicits an unknown constraint, it updates the count and notifies constrained agents.

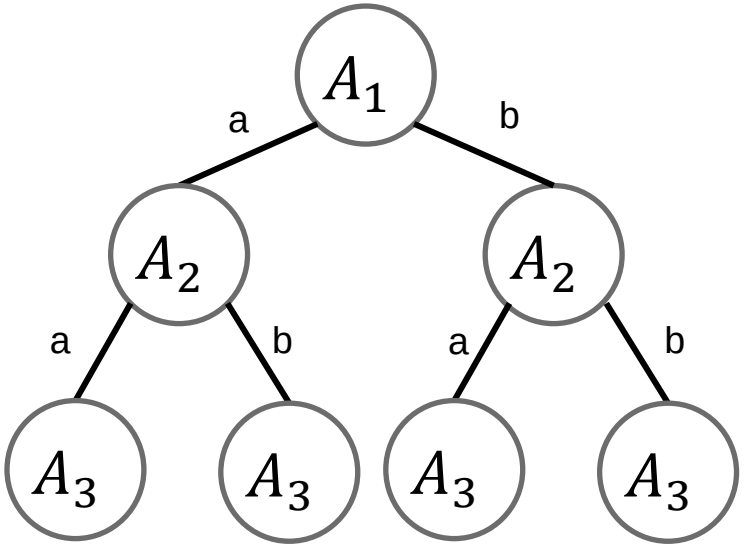

Heuristics

Ex-ante

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

- Pre-process count amount of unknown costs throughout the solution space
- Prioritize subtrees where you have more knowledge.
- When an agent elicits an unknown constraint, it updates the count and notifies constrained agents.

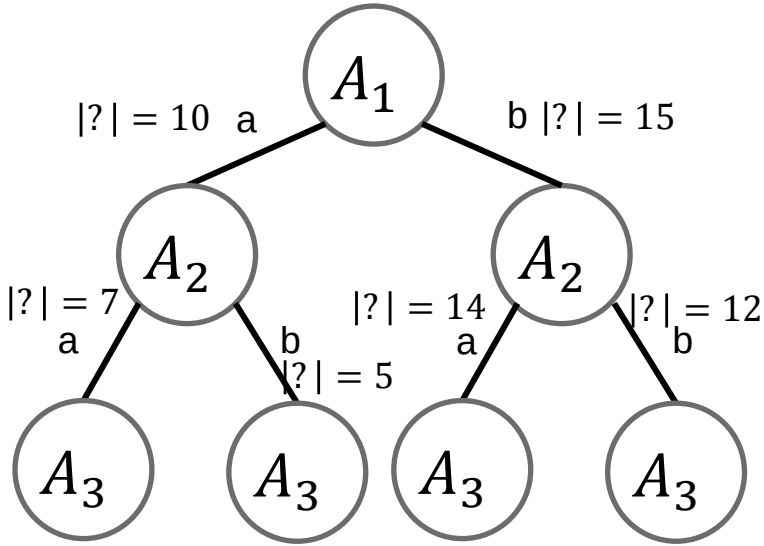

Heuristics

Ex-ante

Complete\Incomplete

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

- Pre-process count amount of unknown costs throughout the solution space
- Prioritize subtrees where you have more knowledge.
- When an agent elicits an unknown constraint, it updates the count and notifies constrained agents.

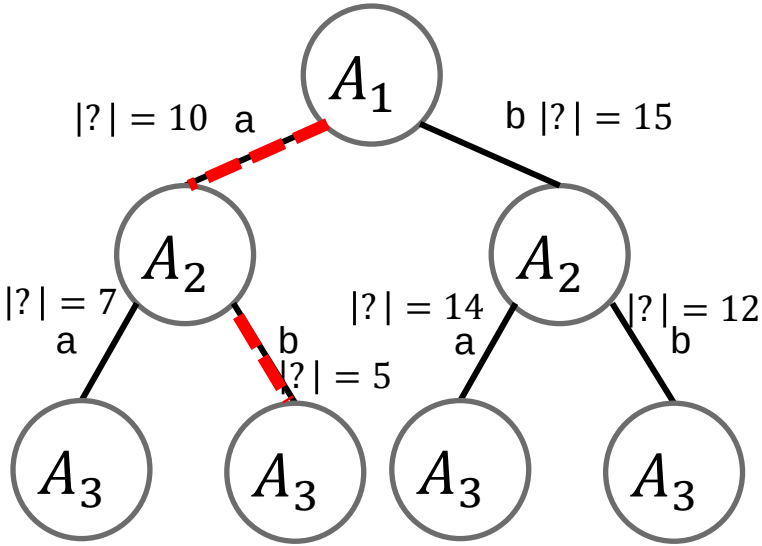

Heuristics

Ex-ante

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Pre-process count amount of unknown costs throughout the solution space

Prioritize subtrees where you have more knowledge.

When an agent elicits an unknown constraint, it updates the count and notifies constrained agents.

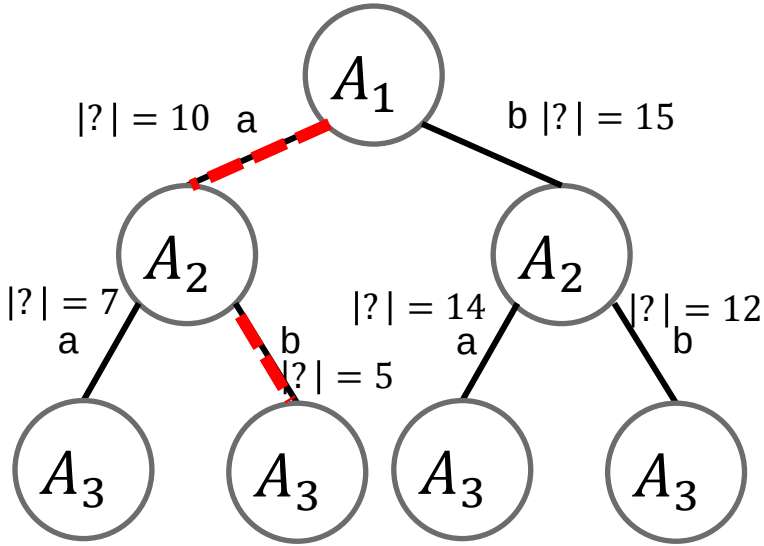

Ex-Ante I DCOP

Complete

Synchronous Branch and Bound

Incomplete

MGM

Heuristics

Ex-ante

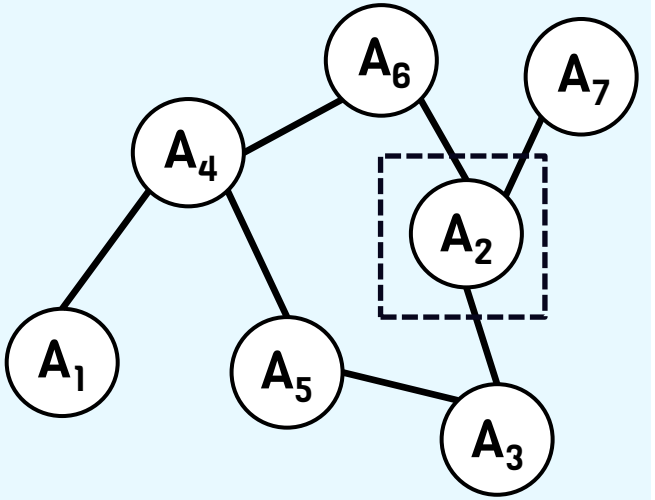

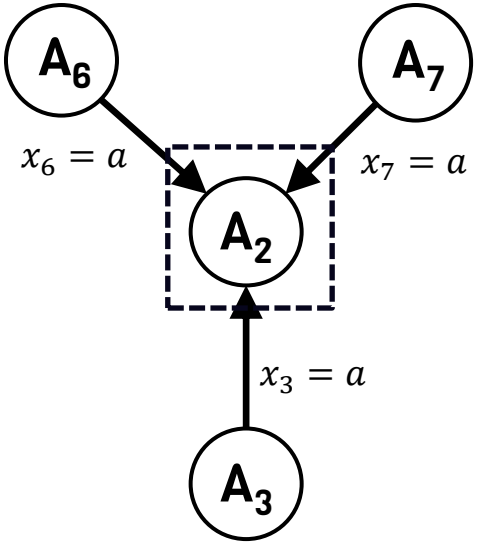

Heuristics

Ex-ante

A_2

$D_2 = \{a, b, c\}$

Potential value	Neighbors' values	cost
$x_2 = a$	$x_3 = a$	
	$x_6 = a$	
	$x_7 = a$	

Potential value	Neighbors' values	cost
$x_2 = b$	$x_3 = a$	
	$x_6 = a$	
	$x_7 = a$	

Potential value	Neighbors' values	cost
$x_2 = c$	$x_3 = a$	
	$x_6 = a$	
	$x_7 = a$	

Heuristics

Ex-ante

A_2

$D_2 = \{a, b, c\}$

Potential value	Neighbors' values	cost
$x_2 = a$	$x_3 = a$	Known or Unknown
	$x_6 = a$	
	$x_7 = a$	

Potential value	Neighbors' values	cost
$x_2 = b$	$x_3 = a$	Known or Unknown
	$x_6 = a$	
	$x_7 = a$	

Potential value	Neighbors' values	cost
$x_2 = c$	$x_3 = a$	Known or Unknown
	$x_6 = a$	
	$x_7 = a$	

Heuristics

Ex-ante

Examine the uncertainty for elicitation cost

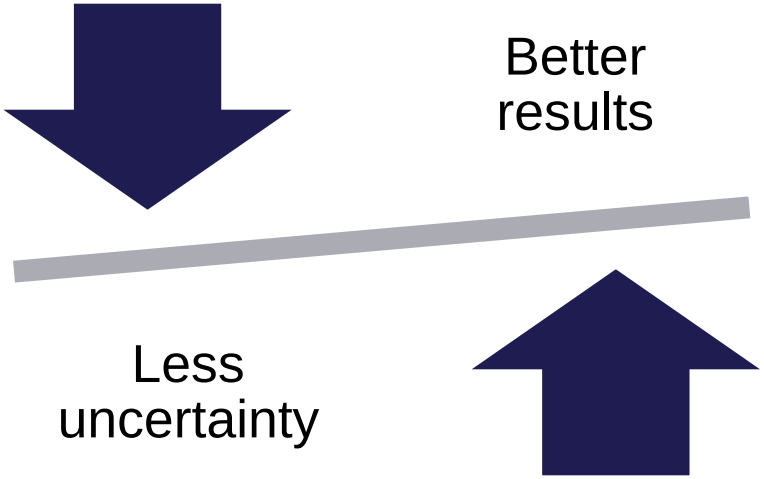

Heuristics

Ex-ante

Examine the uncertainty for elicitation cost

conditions

Is there less than 50% uncertainty in the examined value?

Compare the best-known cost with the cost of the examined value and see if the difference exceeds a defined threshold.

Heuristics

Ex-ante

Examine the uncertainty for elicitation cost

conditions

Is there less than 50% uncertainty in the examined value?

Compare the best-known cost with the cost of the examined value and see if the difference exceeds a defined threshold.

$$\#unknown\ constraints_{alternative} \leq \frac{\#neighbors}{2}$$

Heuristics

Ex-ante

Examine the uncertainty for elicitation cost

conditions

Is there less than 50% uncertainty in the examined value?

Compare the best-known cost with the cost of the examined value and see if the difference exceeds a defined threshold.

$$f(x) = \left(1 - \frac{e^i}{e^{\frac{n}{2}} + e^i}\right) * g(constraint\ costs)$$

Experimental Evaluation

Experimental Design

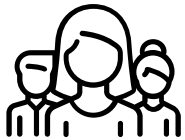

7 Agents
(1 variable)

4 Values in the domain

Complete Algorithm

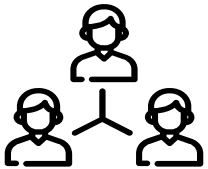

3 neighbors on average

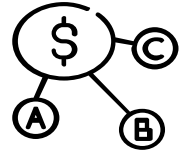

Constraint costs [2,5]

Elicitation costs [0,20]

Experimental Evaluation

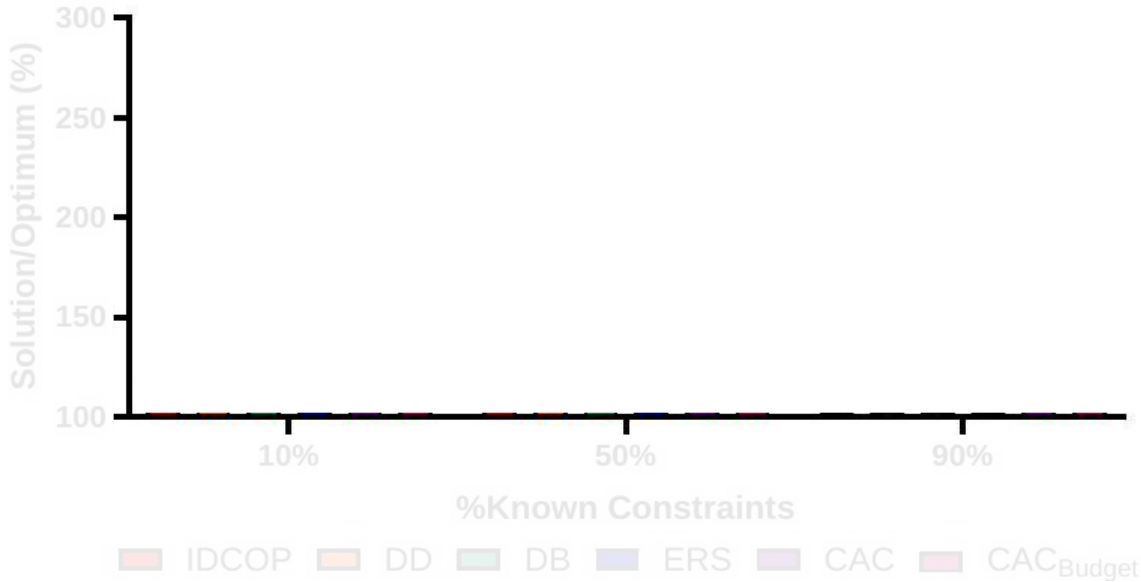

Experimental Evaluation

$$\frac{Solution}{Optimum} \geq 1$$

Small value

Large Value

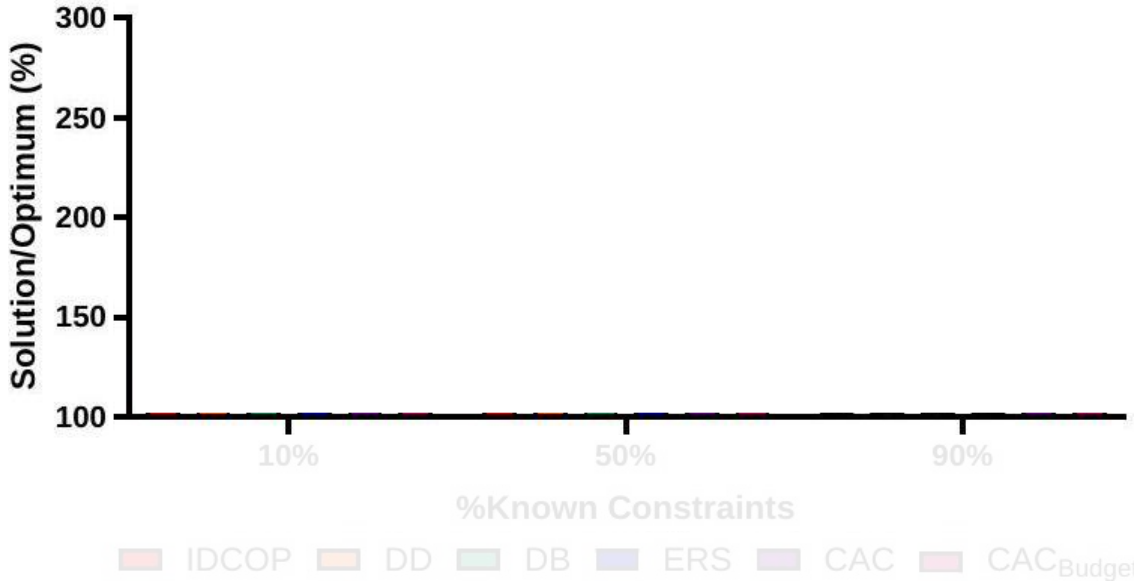

Experimental Evaluation

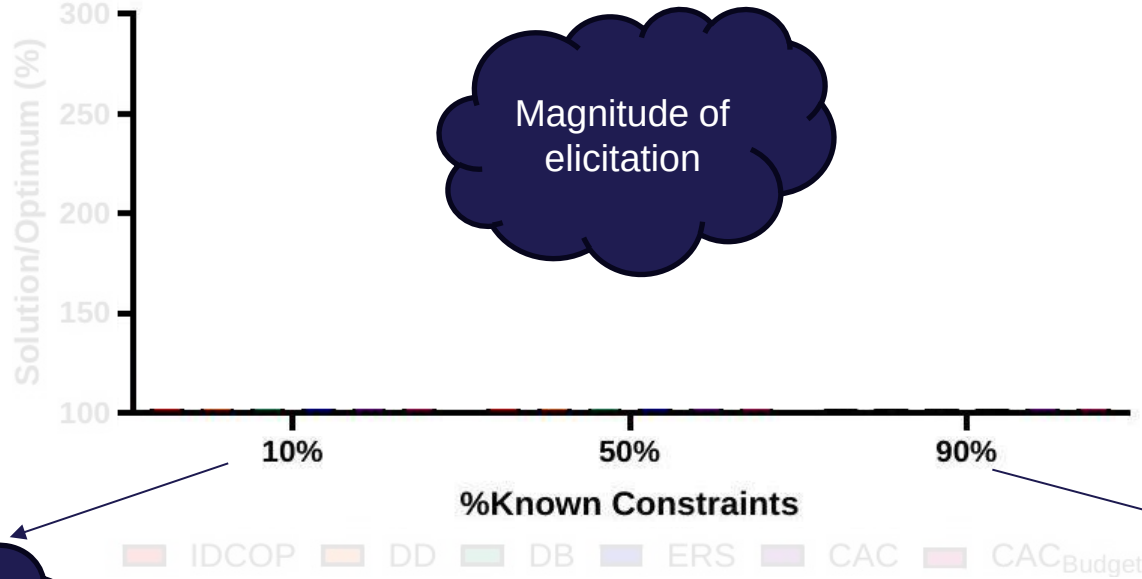

Magnitude of elicitation

Most cost tables are empty

Very close to the original problem

Experimental Evaluation

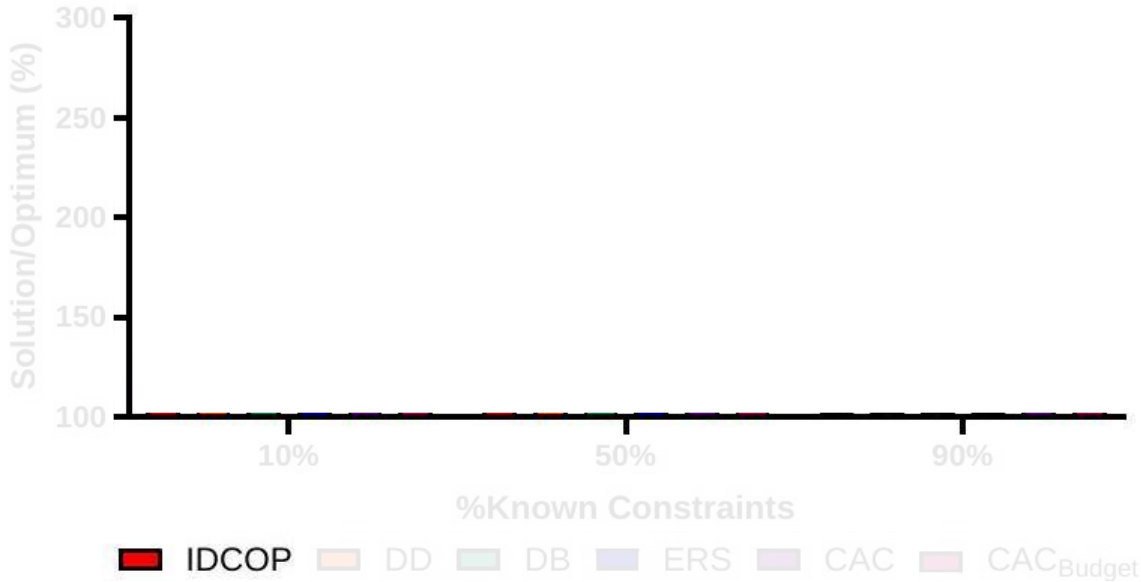

Go through branches with no elicitation cost

Experimental Evaluation

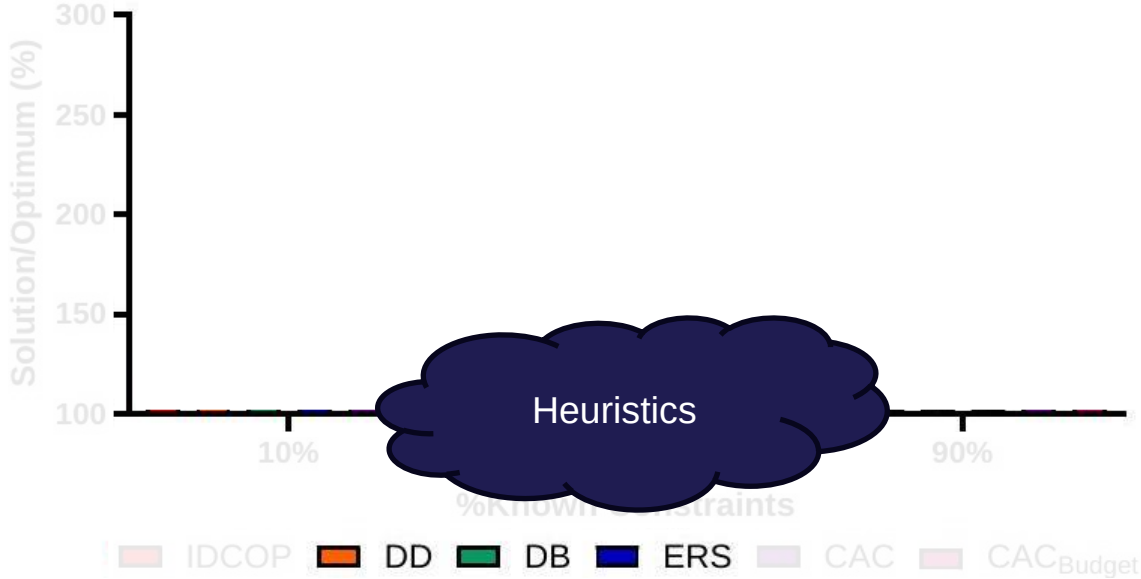

1) Depth Dependent (DD)

2) Distance From Bound (DB)

3) Elicitation Required in Sub-Search-Tree (ERS)

Experimental Evaluation

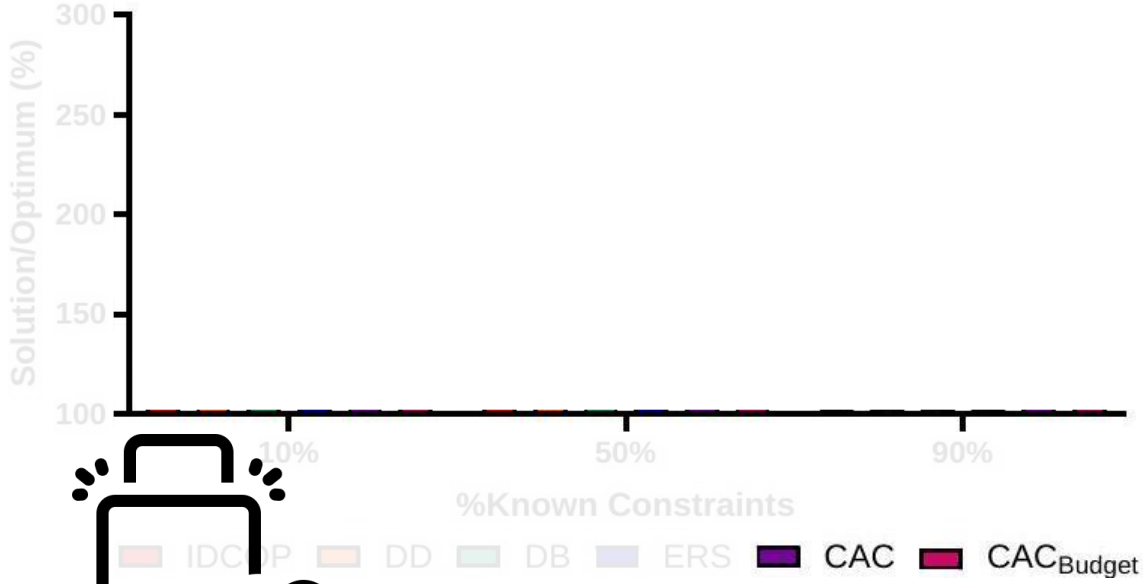

Carry elicitation costs with and without budget limitations

Experimental Evaluation

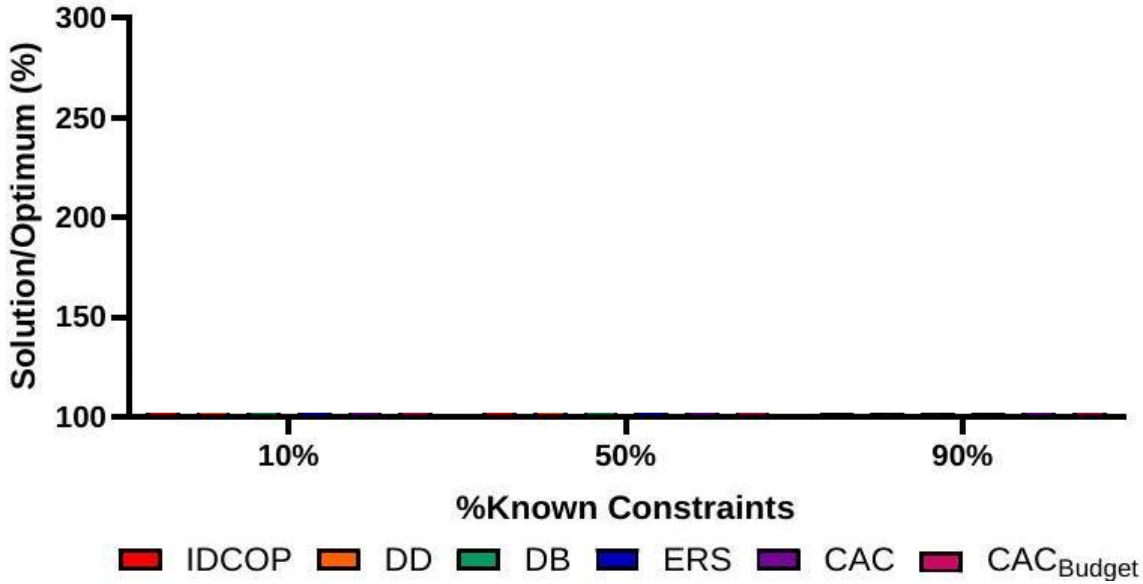

Solution cost for problems in which agents had a budget of 105

Experimental Evaluation

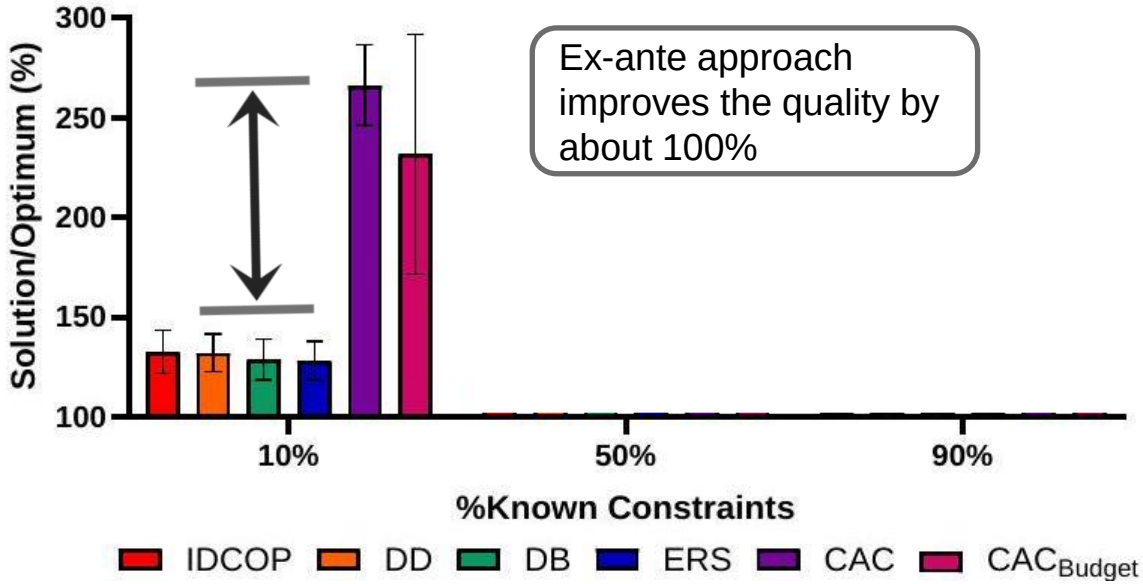

Solution cost for problems in which agents had a budget of 105

Experimental Evaluation

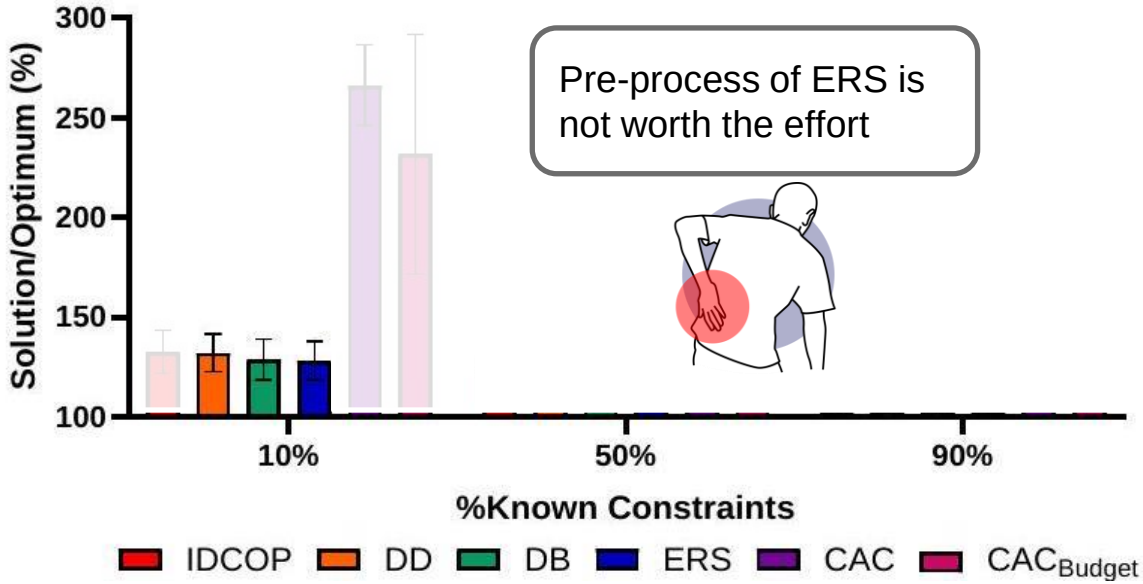

Solution cost for problems in which agents had a budget of 105

Experimental Evaluation

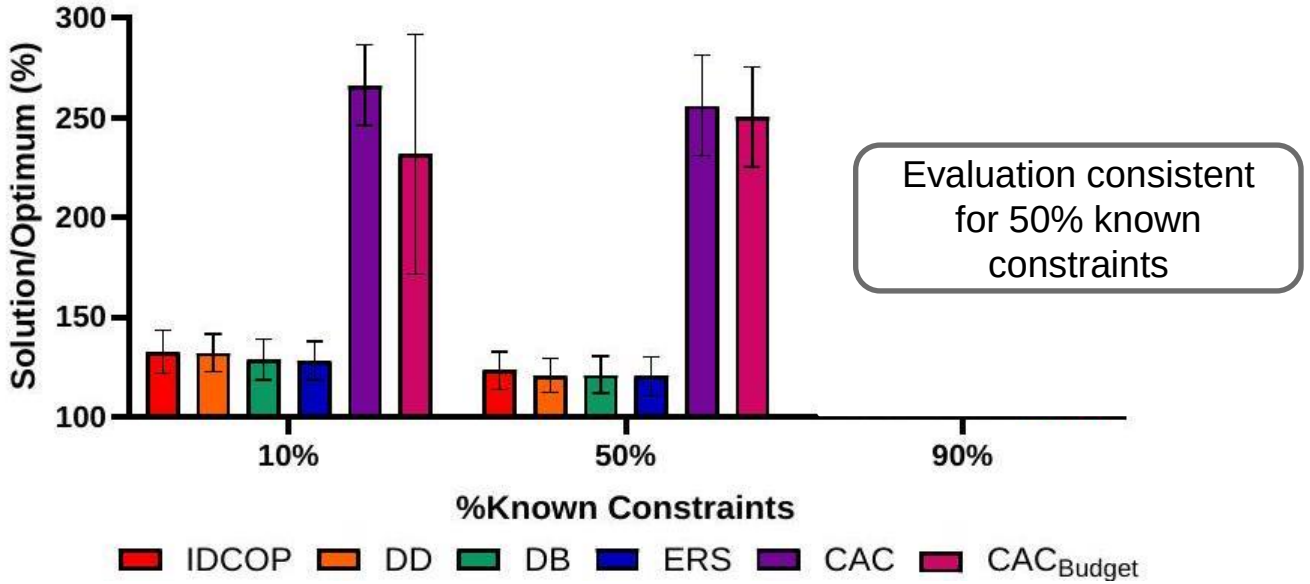

Solution cost for problems in which agents had a budget of 105

Experimental Evaluation

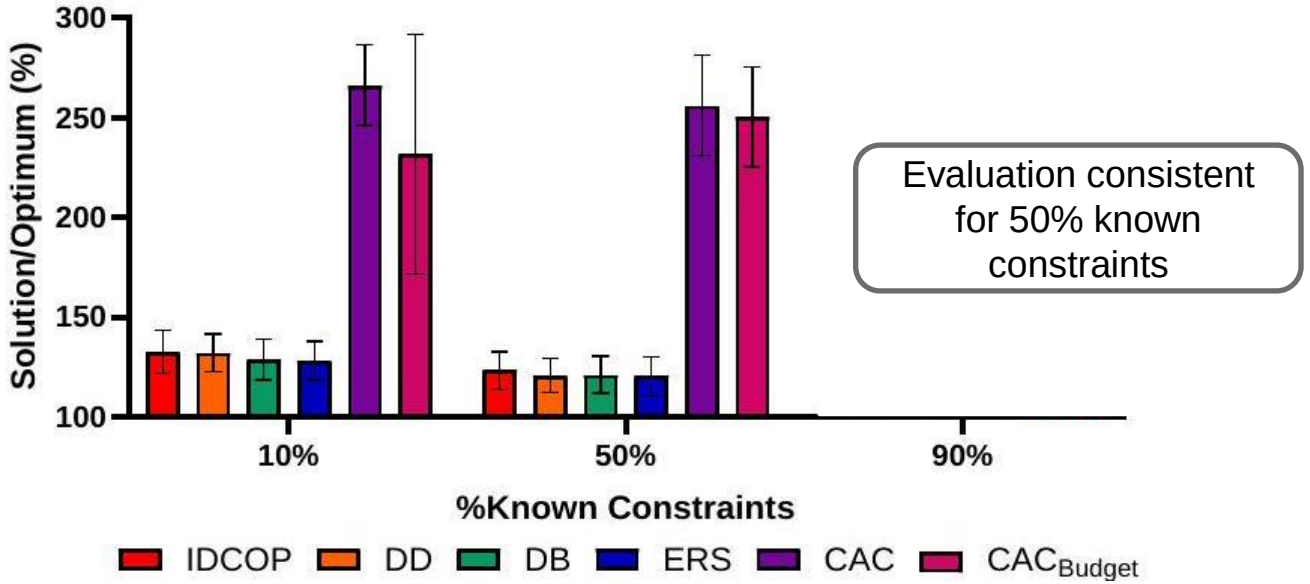

Solution cost for problems in which agents had a budget of 105

Experimental Evaluation

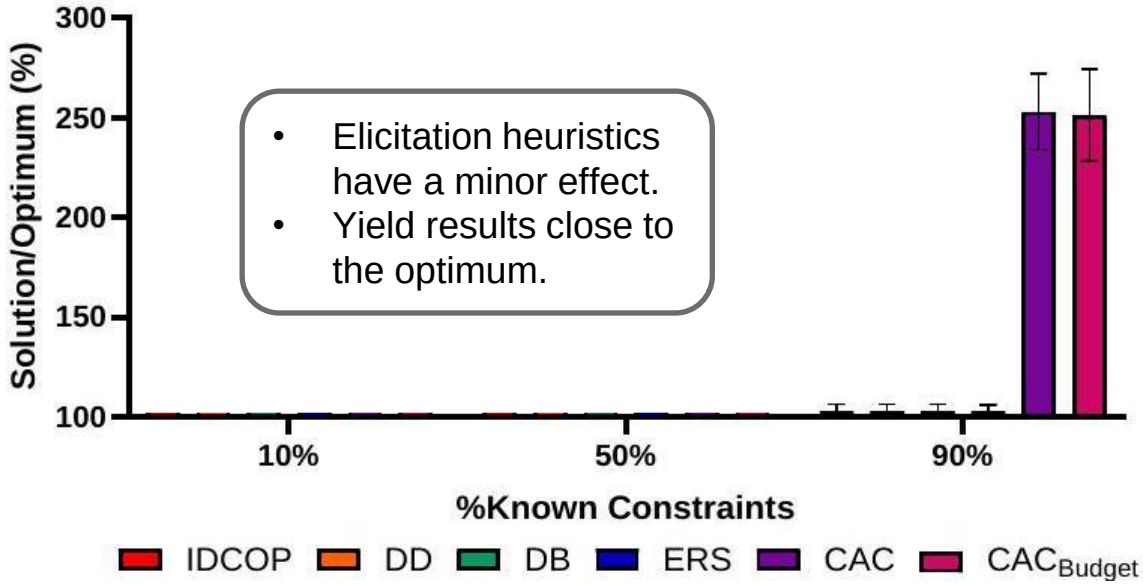

- Elicitation heuristics have a minor effect.
- Yield results close to the optimum.

Solution cost for problems in which agents had a budget of 105

Experimental Evaluation

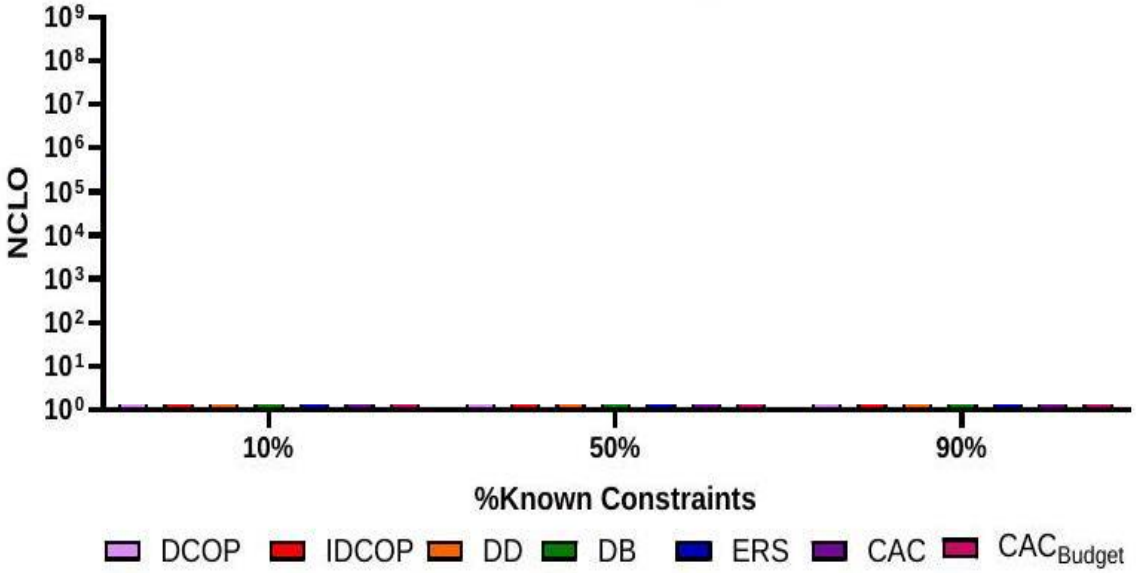

Experimental Evaluation

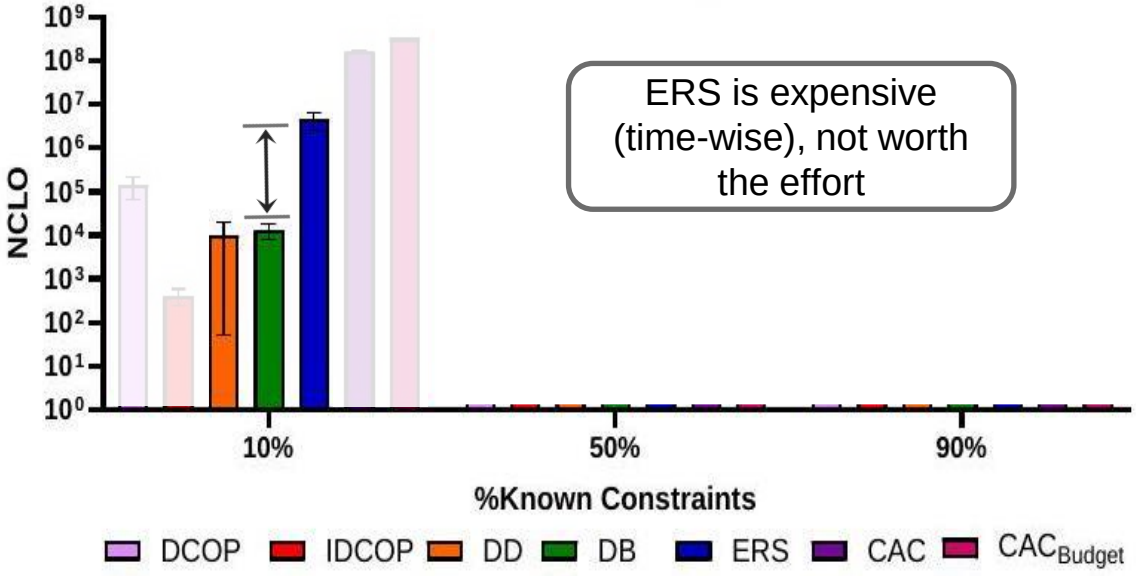

ERS is expensive (time-wise), not worth the effort

Experimental Evaluation

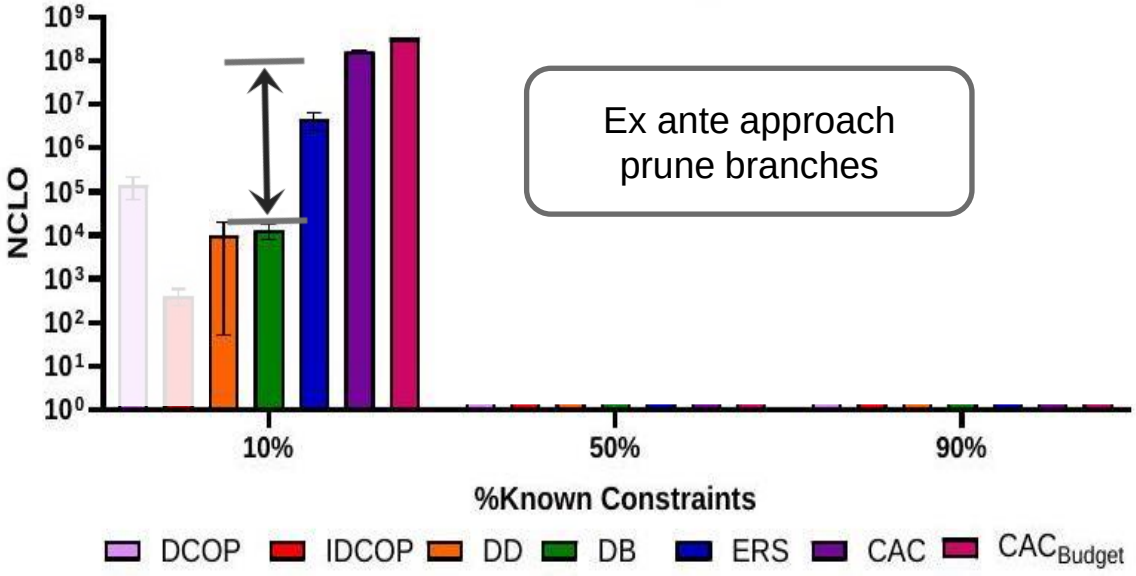

Ex ante approach
prune branches

Experimental Evaluation

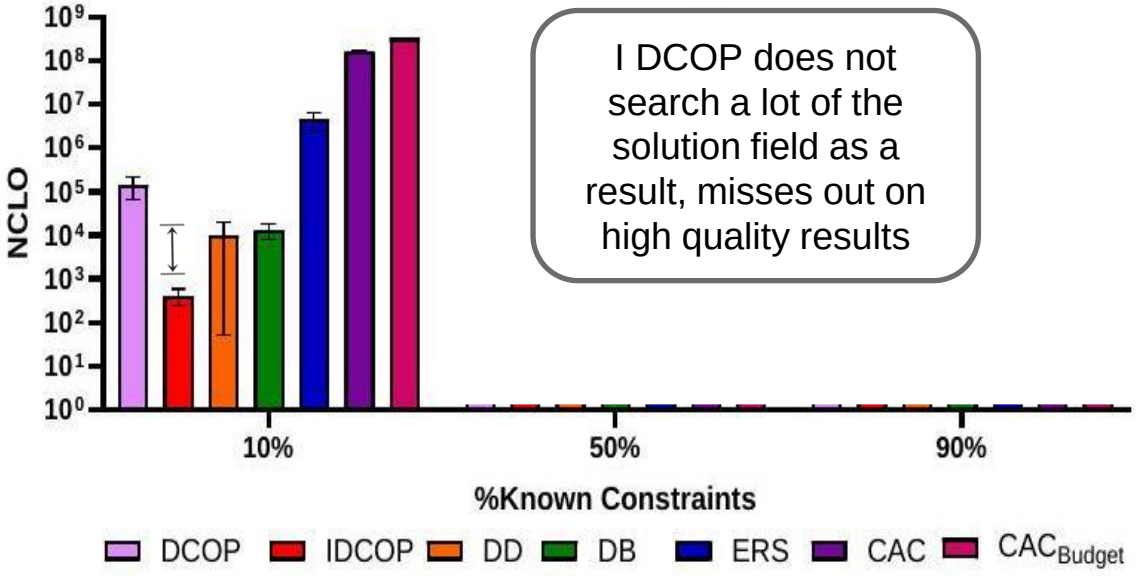

I DCOP does not search a lot of the solution field as a result, misses out on high quality results

Experimental Evaluation

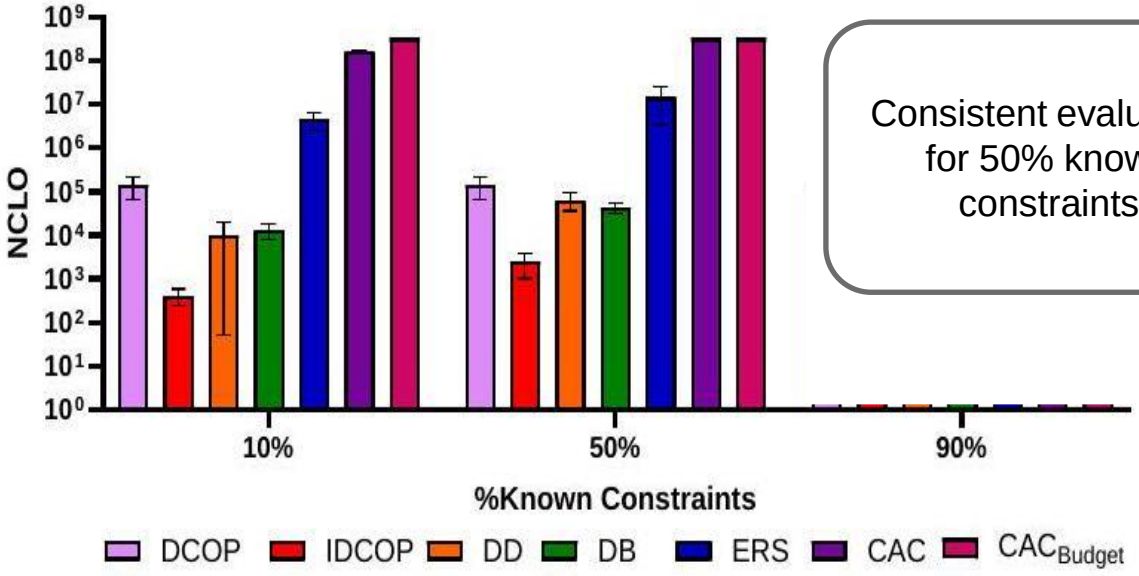

Consistent evaluation for 50% known constraints

Experimental Evaluation

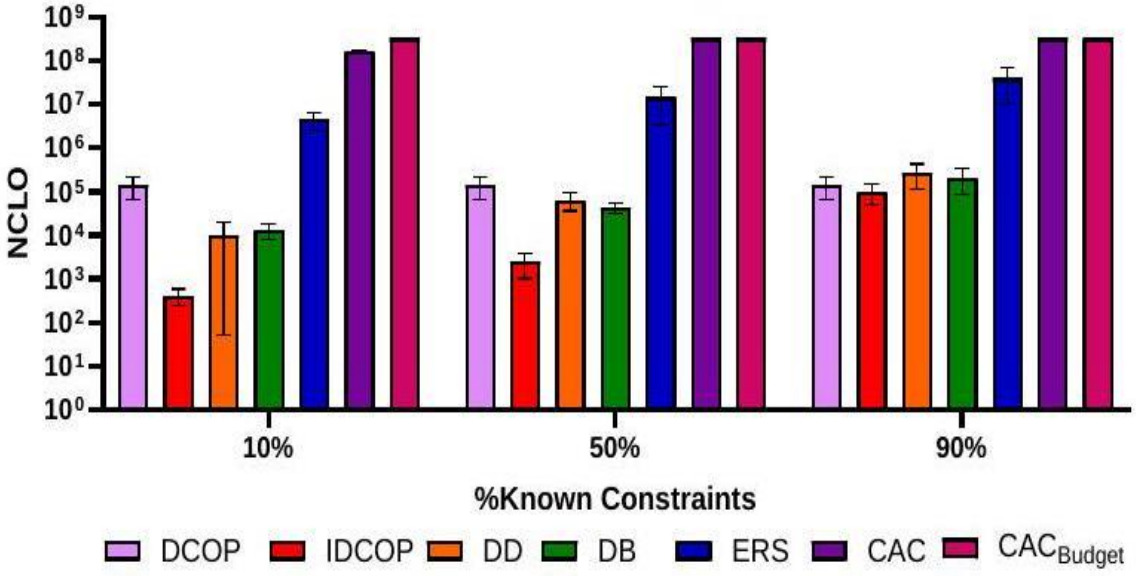

For 90% known constraints, heuristics take longer as they can afford more exploration.

Experimental Design

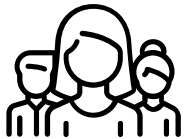

50 Agents
(1 variable)

10 Values in
the domain

Incomplete Algorithms

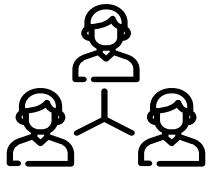

20 neighbors
on average

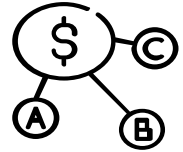

Constraint
costs [2,5]

Elicitation
costs [0,20]

Experimental Evaluation

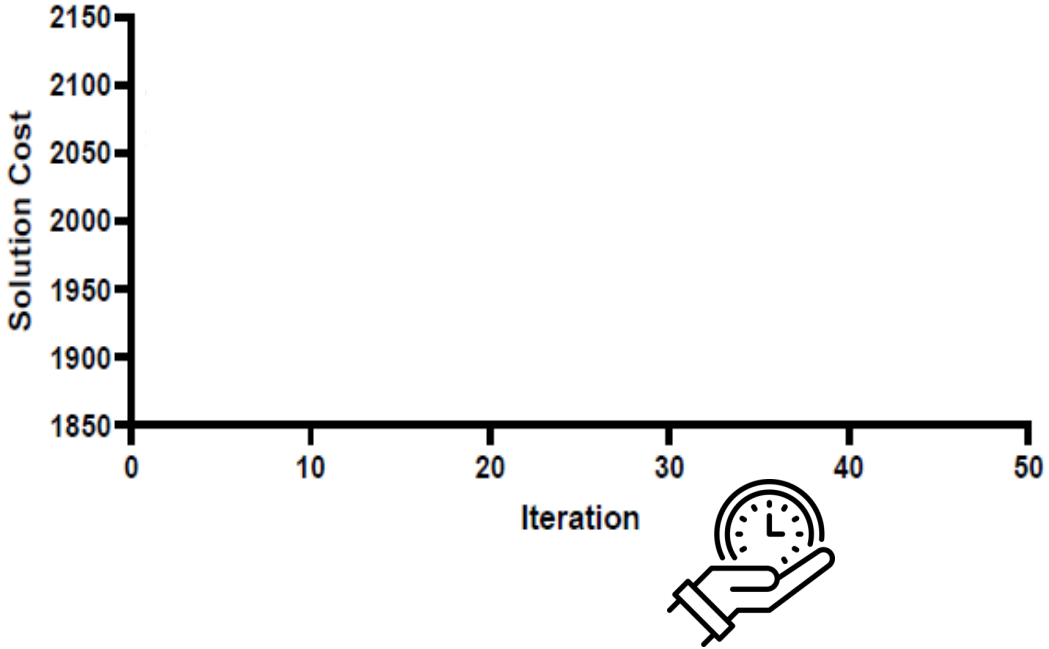

Experimental Evaluation

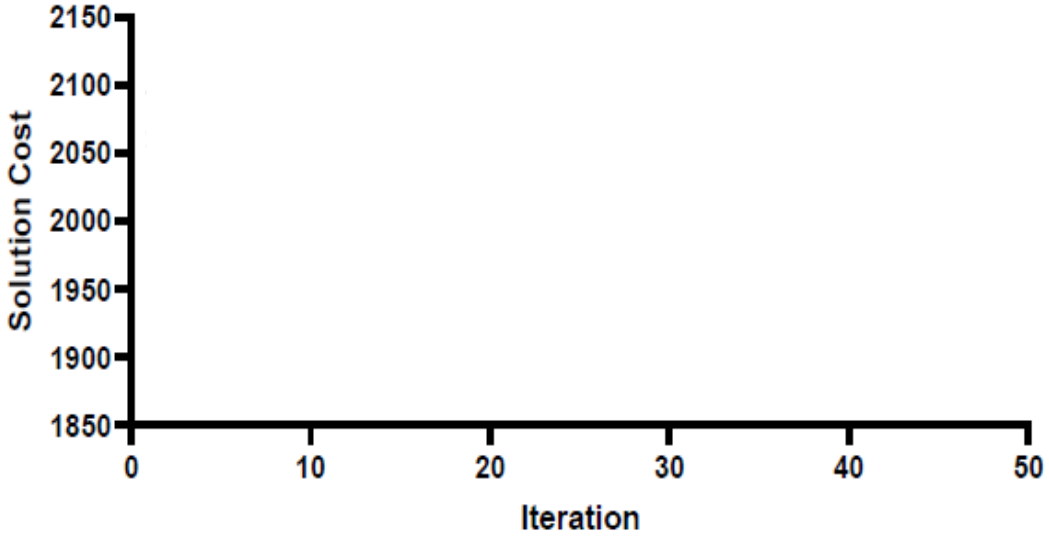

IDCOP $NHC_{10\%}$
IDCOP $NHC_{50\%}$
IDCOP $NHC_{90\%}$

Ex-post

Experimental Evaluation

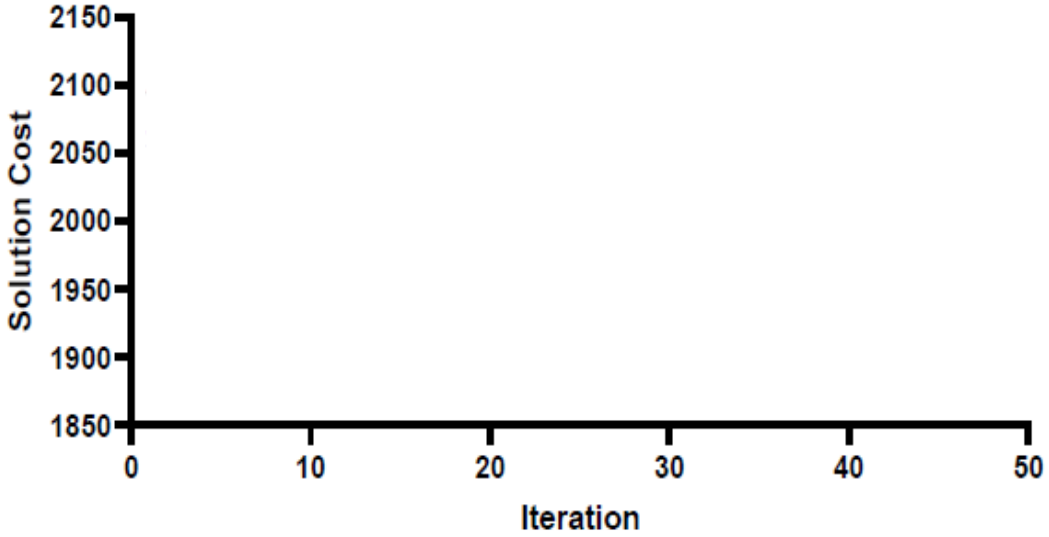

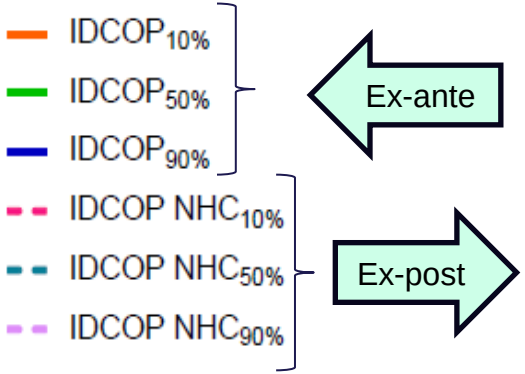

Experimental Evaluation

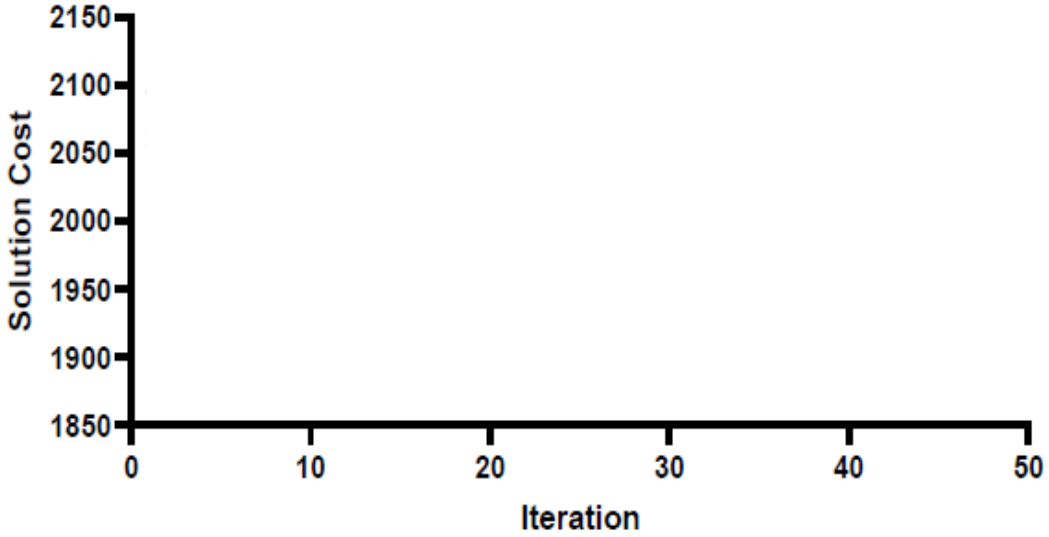

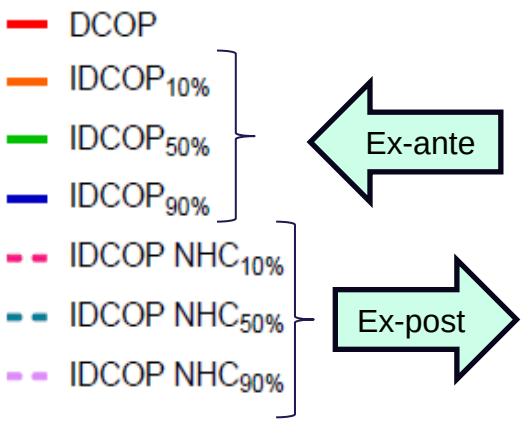

Experimental Evaluation

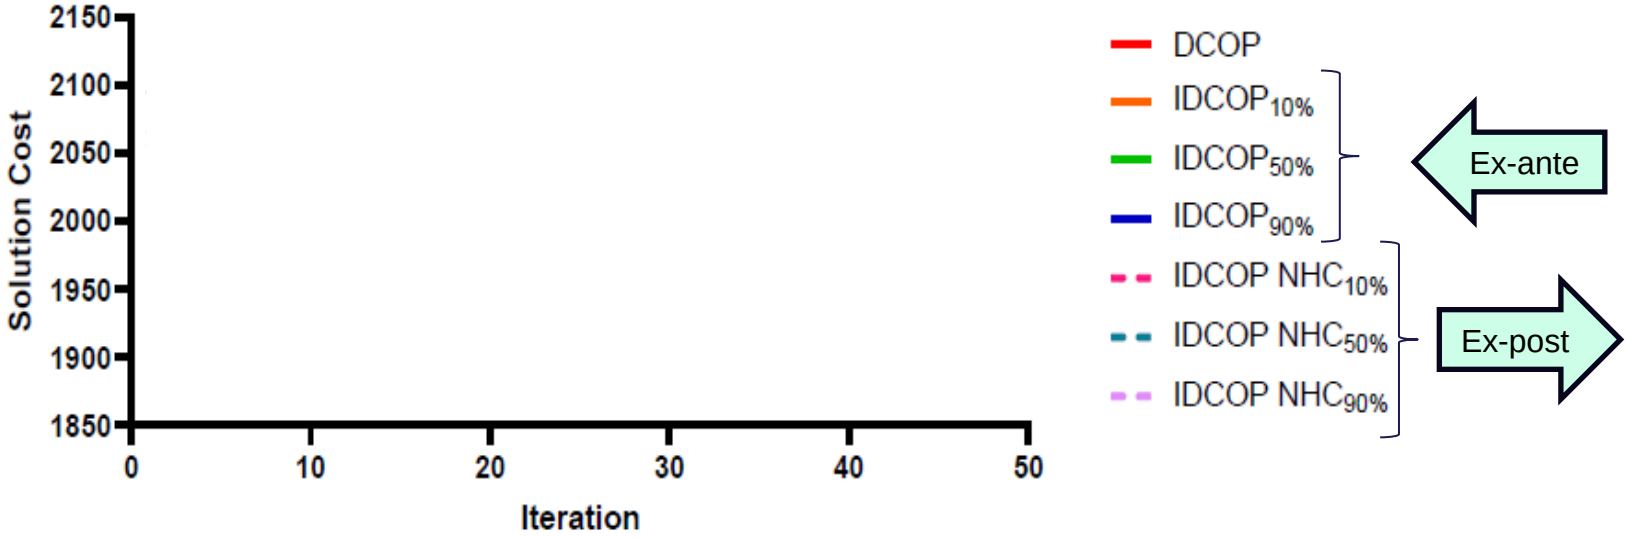

Solution cost as a function of the number of iterations with k=180

Experimental Evaluation

MGM

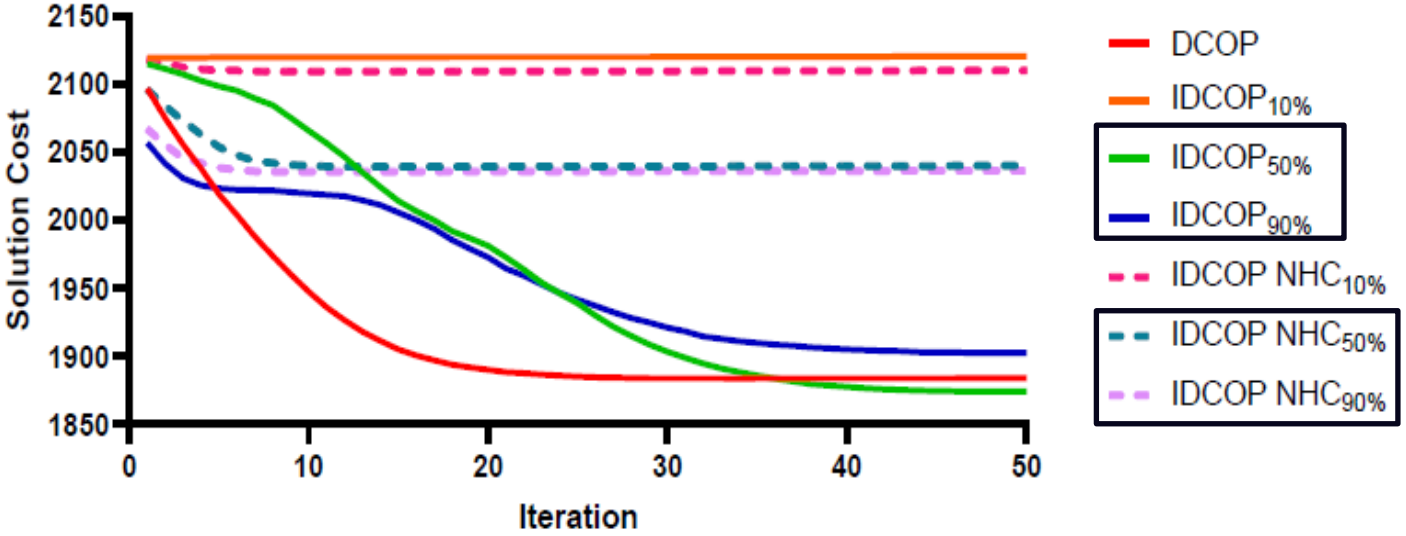

IDCOP50%
IDCOP90%

Ex-ante Approach outperforms Ex-post

IDCOP NHC50%
IDCOP NHC90%

Experimental Evaluation

MGM

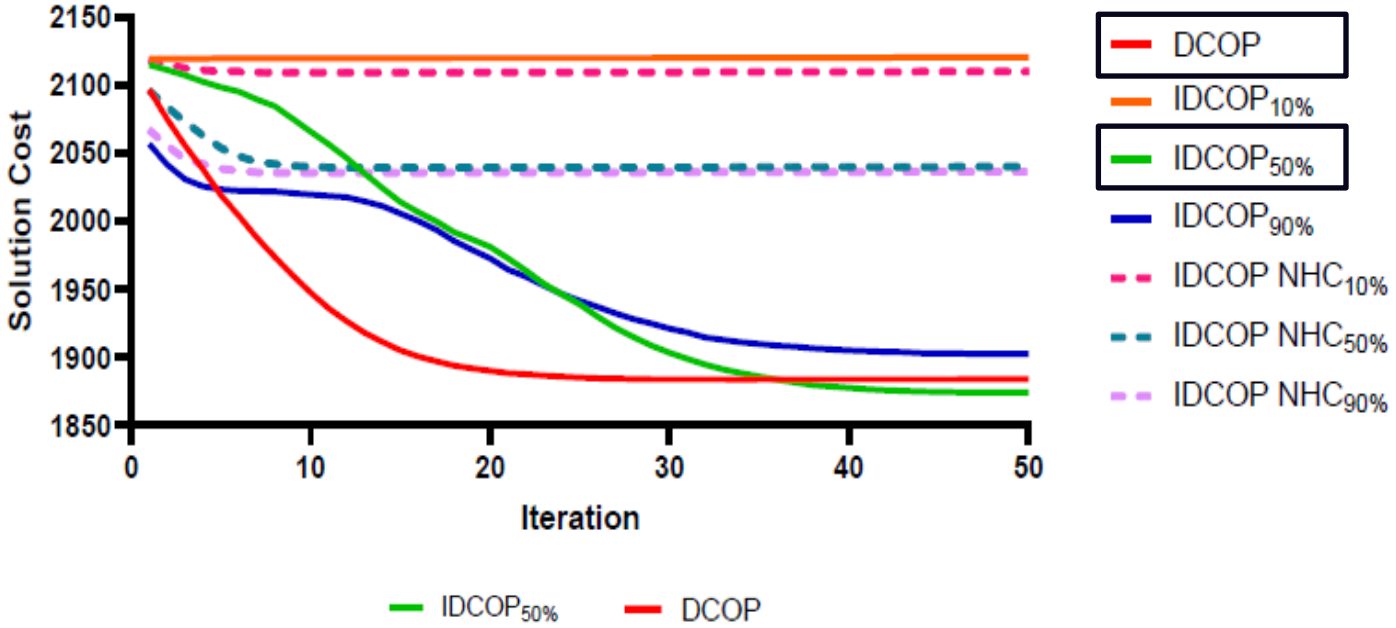

Unknown costs, can lead to positive explorative effect

Conclusions

Conclusions

Novel Approach

Introduced a new method for solving I-DCOPs, considering elicitation costs before hand.

Conclusions

Novel Approach

Introduced a new method for solving I-DCOPs, considering elicitation costs before hand.

Realistic and Efficient

The approach is more realistic and finds higher quality solutions with reduced runtime.

Conclusions

Novel Approach

Introduced a new method for solving I-DCOPs, considering elicitation costs before hand.

Realistic and Efficient

The approach is more realistic and finds higher quality solutions with reduced runtime.

Comprehensive Improvement

Outperforms previous methods across practicality, solution quality, and runtime.

Paper's QR code

Roie Zivan

Ben-Gurion University of the Negev

zivanr@bgu.ac.il

William Yeoh

Washington University in St. Louis

wyeoh@wustl.edu

Any Questions

THANK YOU!

Ben Rachmut

rachmut@post.bgu.ac.il