Tightly and Loosely Coupled Decision Paradigms in Multiagent Expedition

Yang Xiang & Frank Hanshar
University of Guelph
Ontario, Canada

PGM 2008
September 18, 2008

Outline

- Introduction
- What is Multiagent Expedition?
- Collaborative Design Network
- Graphical Model for Multiagent Expedition
- Recursive Model for Multiagent Expedition
- Experimental Results & Discussion
- Conclusion

Introduction

- We consider frameworks for online decision making:
    - Loosely-coupled frameworks (LCF): do not communicate, rely on observing other agents actions to discern state and coordinate with each other
    - Tightly-coupled frameworks (TCF): agents communicate through messages over interfaces that are rigourously defined

Introduction cont.

- Relevant computational advantages of each paradigm are poorly understood.
- We wish to understand the tradeoffs in LCFs and TCFs for multiagent planning.
- In this work we select one example framework from LCFs (RMM) and one from TCFs (CDN).
- We resolve technical issues encountered, and compare them experimentally on a test problem called multiagent expedition.

What is Multiagent Expedition (MAE)?

- Agents have no prior knowledge of how rewards are distributed in the environment.
- Multiple alternative goals with varying rewards are present.
- Coordination problem - objective is for agents to cooperate to maximize team reward.
- Possible applications: multi-robot exploration of Mars, sea-floor exploration, disaster rescue, ...

Instance of MAE

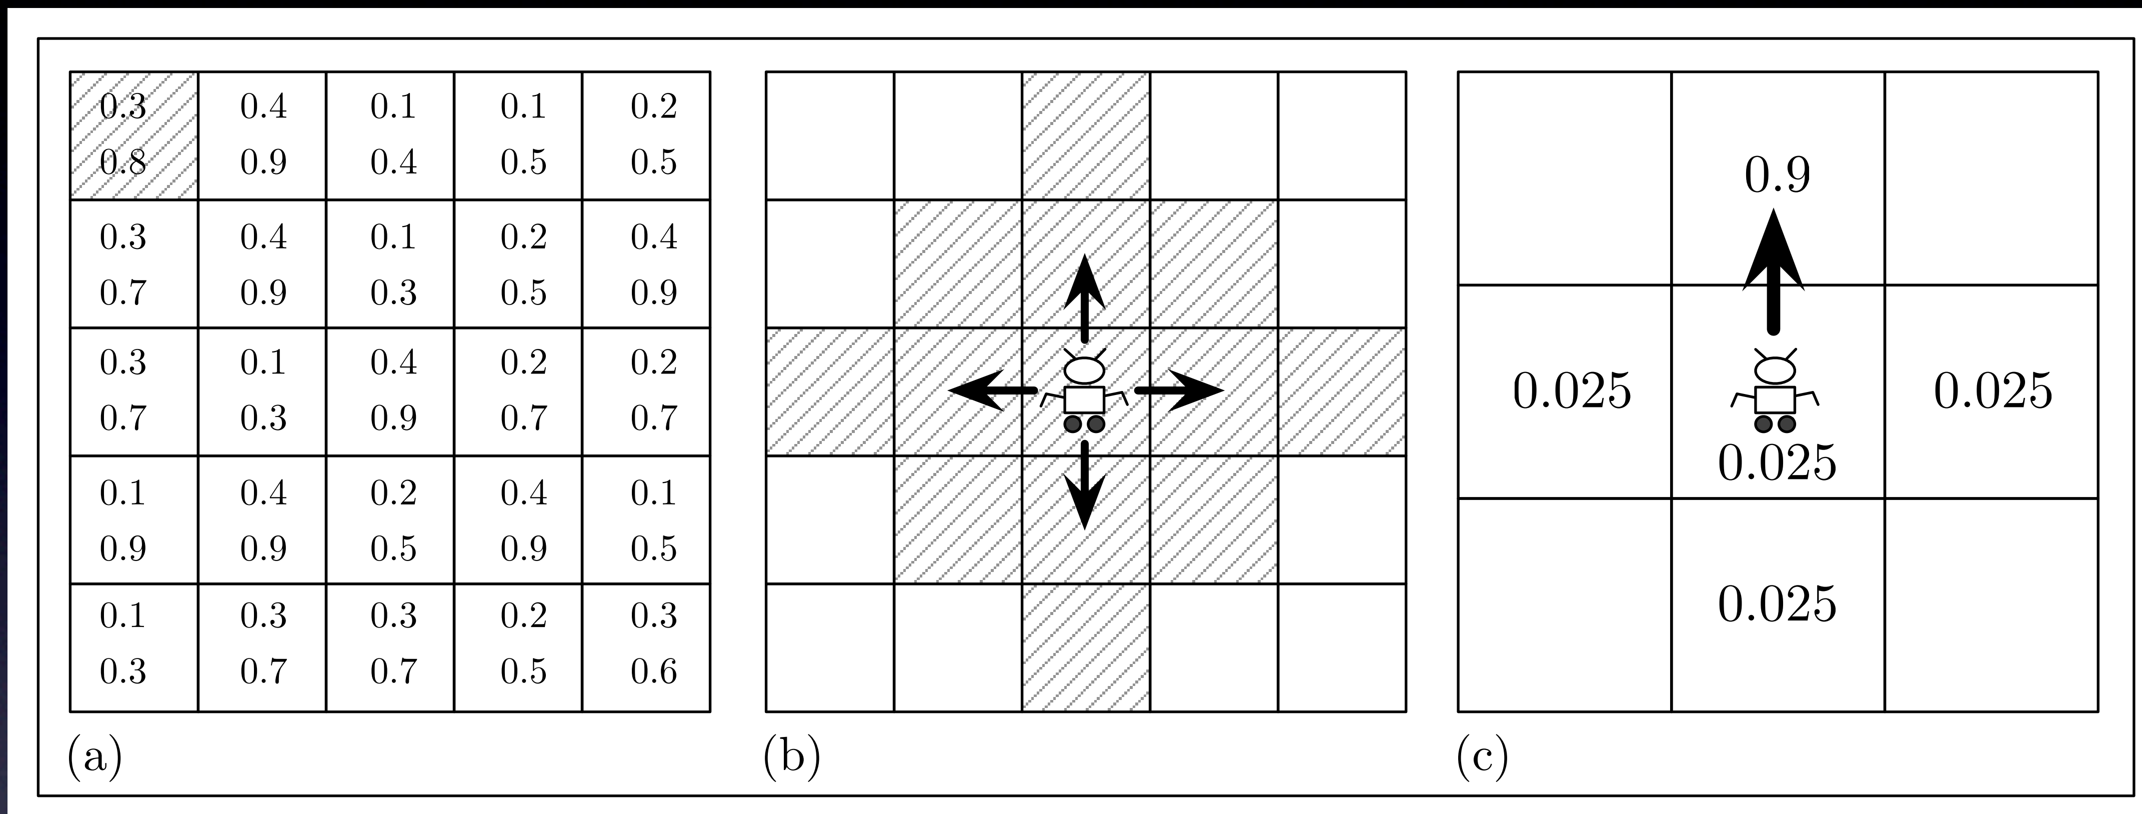

- Each cell has a reward pair (a).
- Observations are local (b).
  - Agent can observe the 13 cells around it.
- Effect of an action uncertain (c).

Instance of MAE

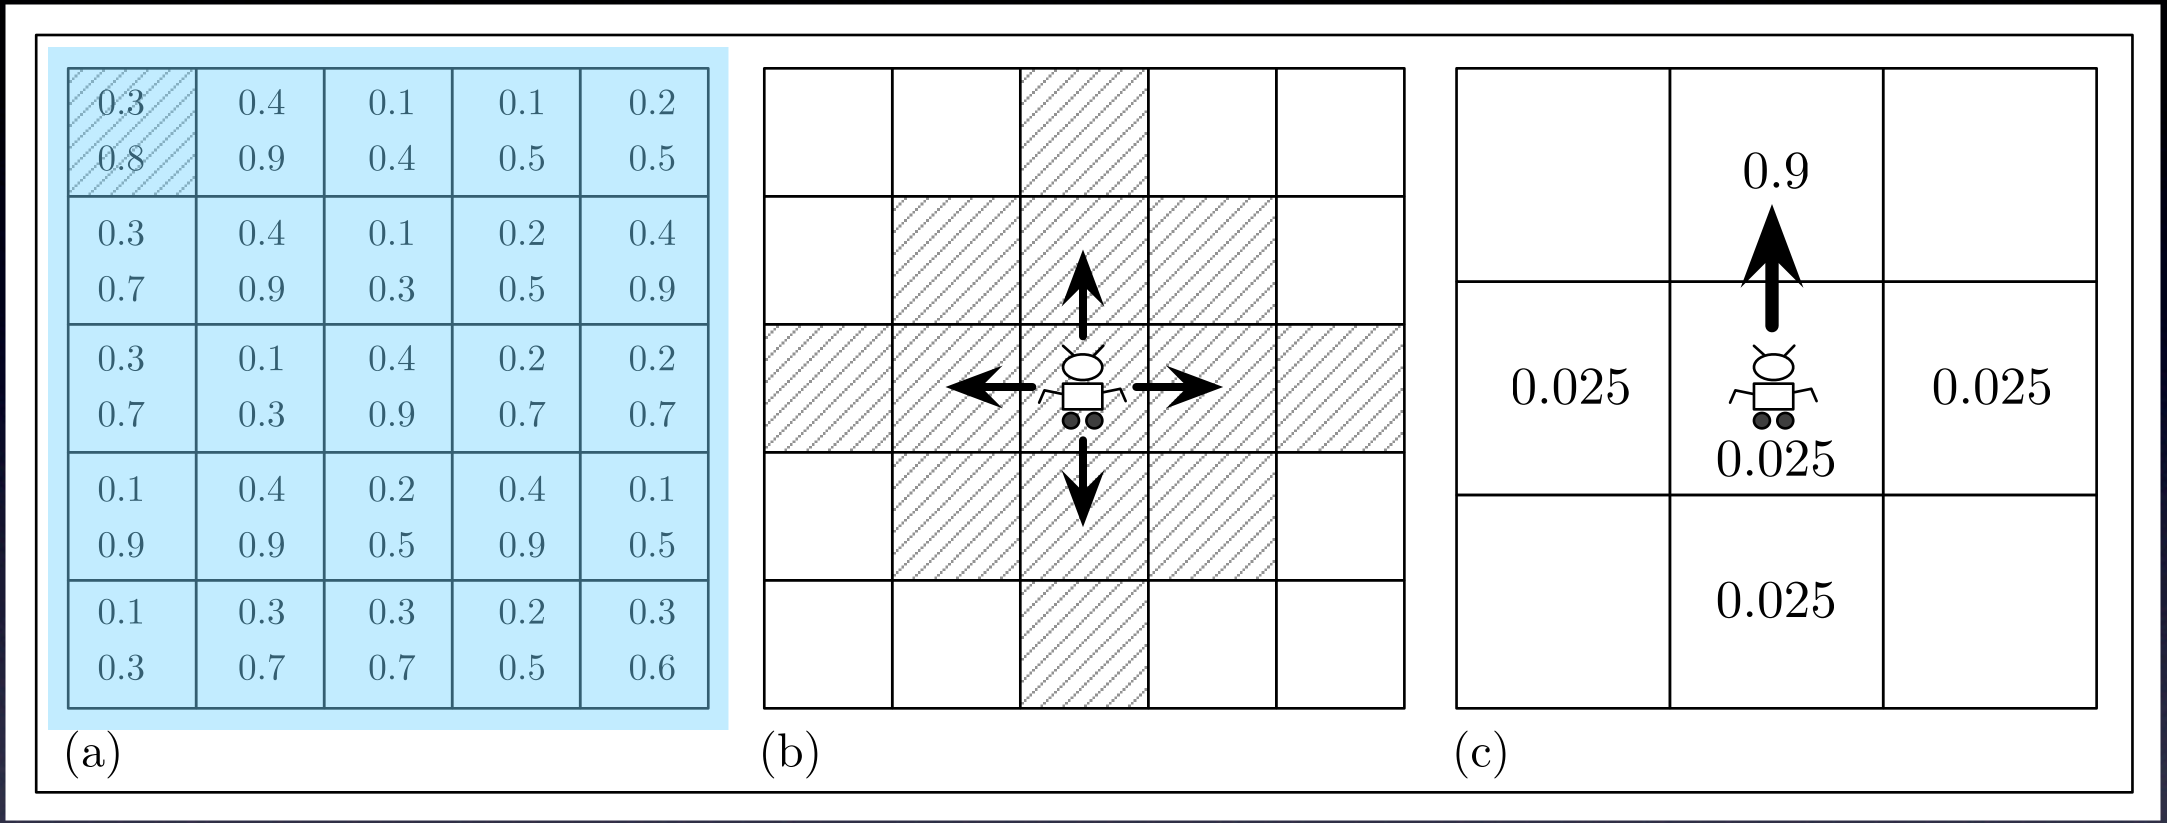

- Each cell has a reward pair (a).
- Observations are local (b).
  - Agent can observe the 13 cells around it.
- Effect of an action uncertain (c).

Instance of MAE

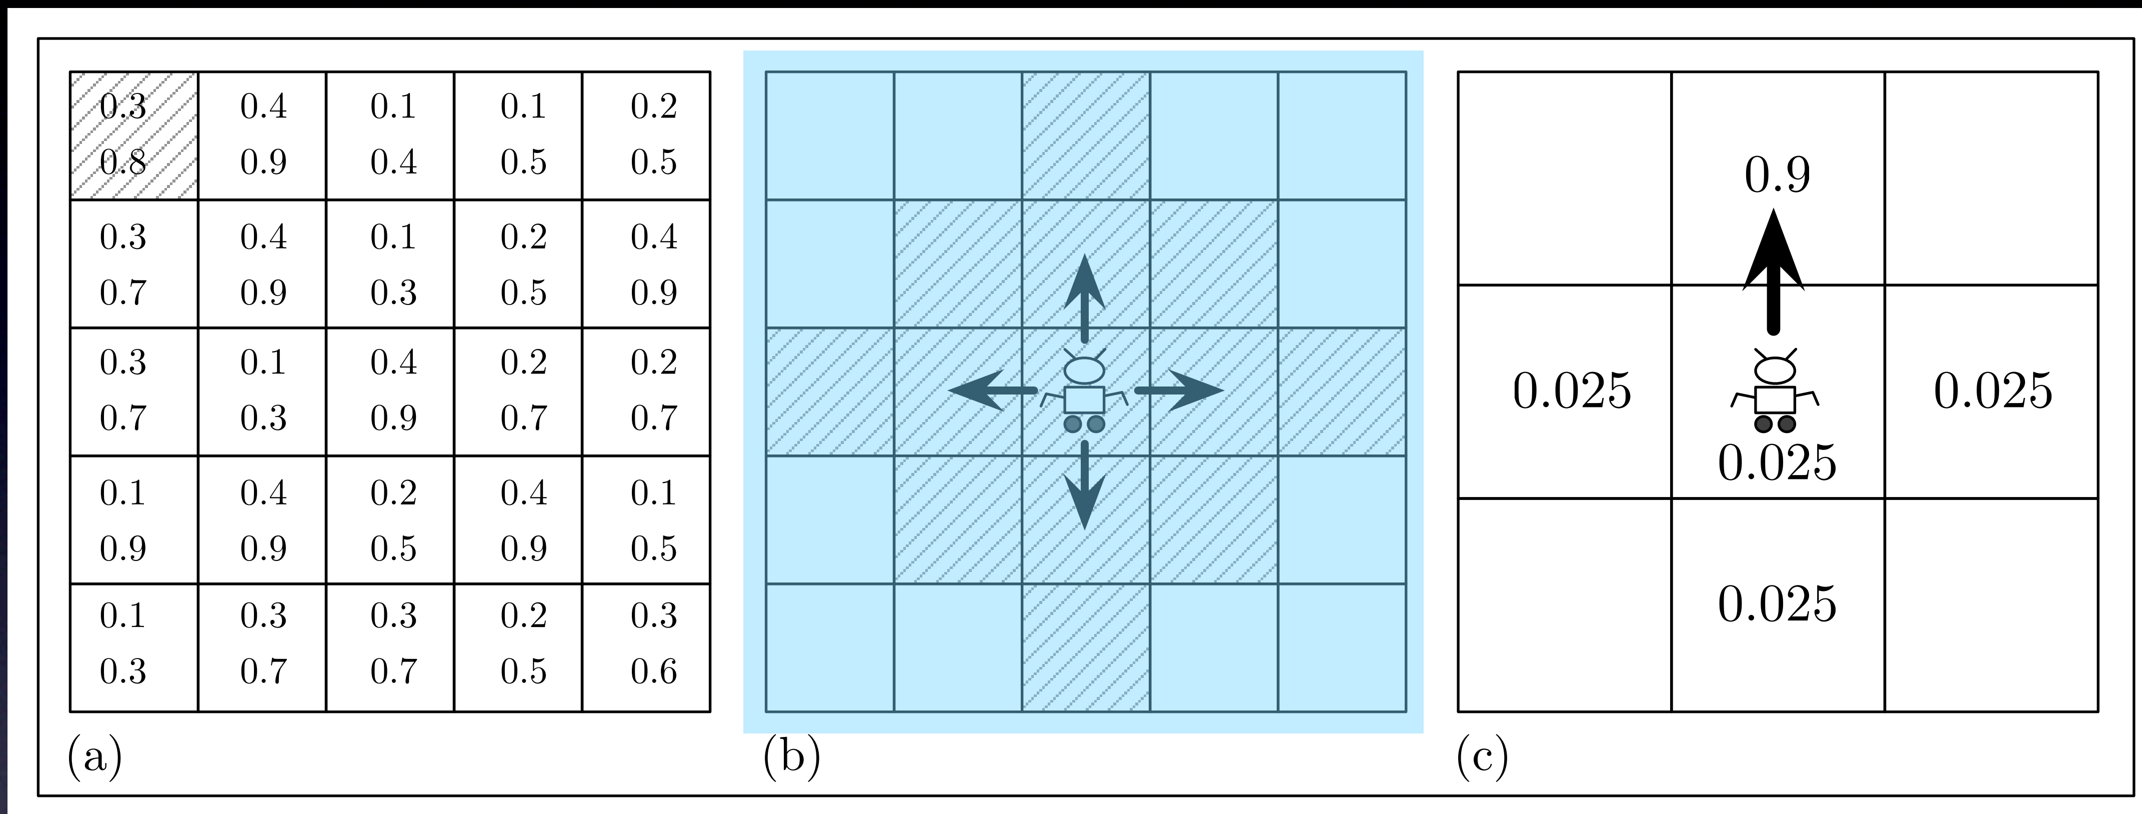

- Each cell has a reward pair (a).
- Observations are local (b).
  - Agent can observe the 13 cells around it.
- Effect of an action uncertain (c).

Instance of MAE

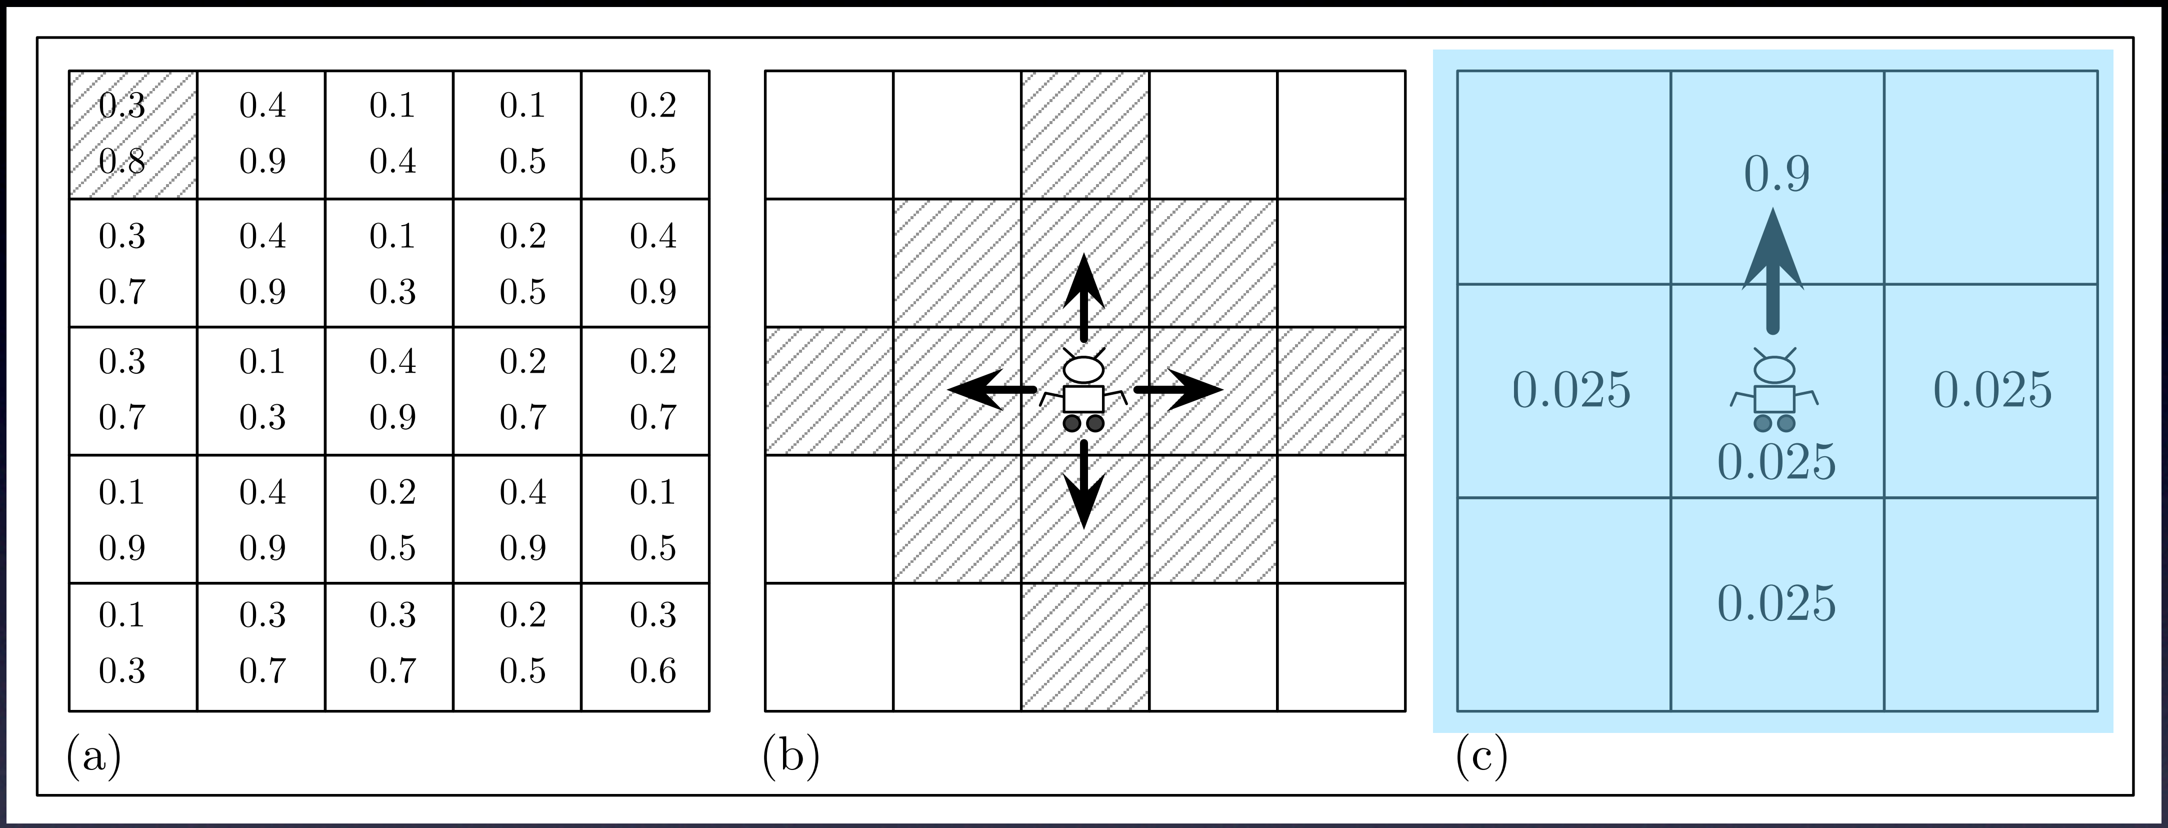

- Each cell has a reward pair (a).
- Observations are local (b).
  - Agent can observe the 13 cells around it.
- Effect of an action uncertain (c).

MAE Rewards

- Physical interaction between agents has some optimal level.
  - Above or below this level will reduce the reward.
  - We set this level at 2 agents, but other levels could be used.
- Thus each cell has a reward pair $(r_1, r_2), r_1, r_2 \in [0, 1]$
  - r_1 denotes unilateral reward, r_2 bilateral reward.

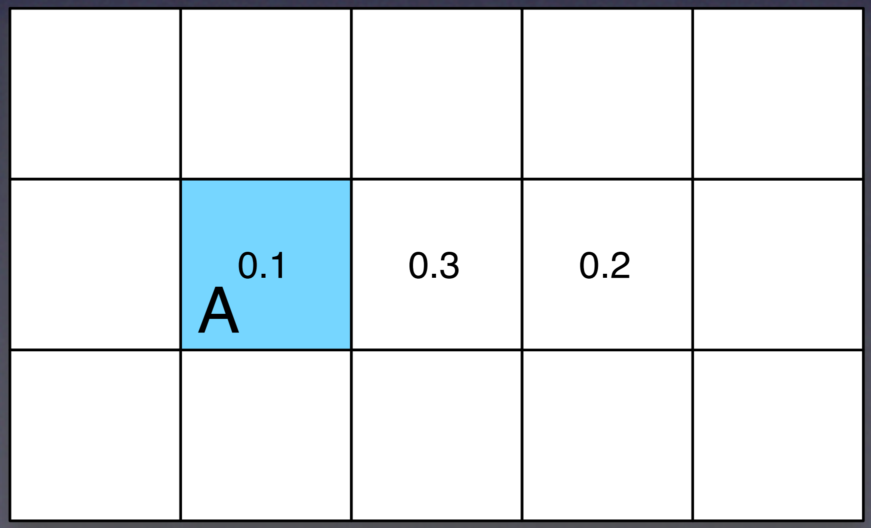

- Agents move and collect utility.
- Cells revert to **default** rewards d after they are visited.

$$d = (r_1, r_2) = (0.1, 0.2)$$

MAE Rewards

- Physical interaction between agents has some optimal level.
  - Above or below this level will reduce the reward.
  - We set this level at 2 agents, but other levels could be used.
- Thus each cell has a reward pair $(r_1, r_2), r_1, r_2 \in [0, 1]$
  - r_1 denotes unilateral reward, r_2 bilateral reward.

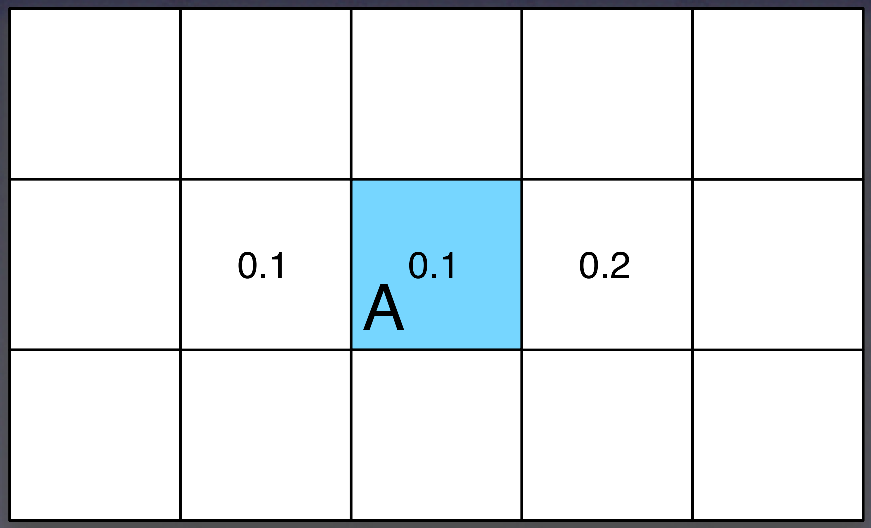

- Agents move and collect utility.
- Cells revert to **default** rewards d after they are visited.

$$d = (r_1, r_2) = (0.1, 0.2)$$

MAE Rewards

- Physical interaction between agents has some optimal level.
  - Above or below this level will reduce the reward.
  - We set this level at 2 agents, but other levels could be used.
- Thus each cell has a reward pair $(r_1, r_2), r_1, r_2 \in [0, 1]$
  - r_1 denotes unilateral reward, r_2 bilateral reward.

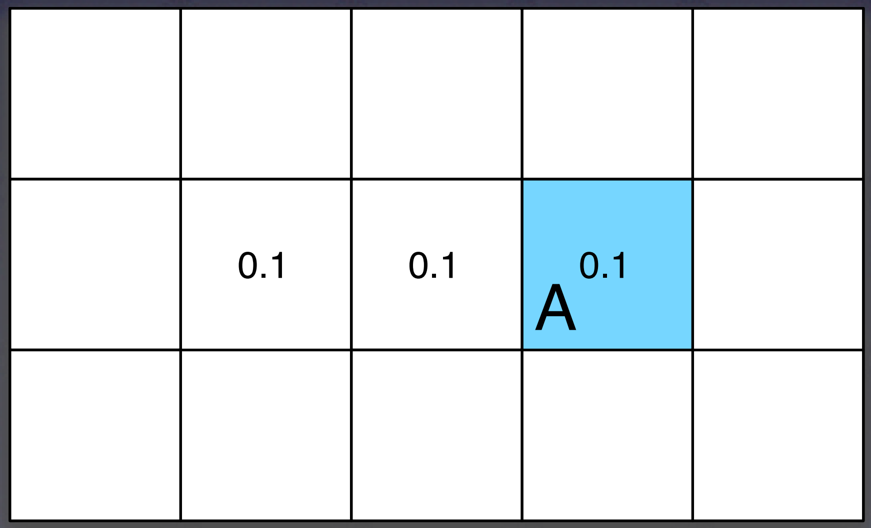

- Agents move and collect utility.
- Cells revert to **default** rewards d after they are visited.

$$d = (r_1, r_2) = (0.1, 0.2)$$

Unilateral / Bilateral Rewards

Initial

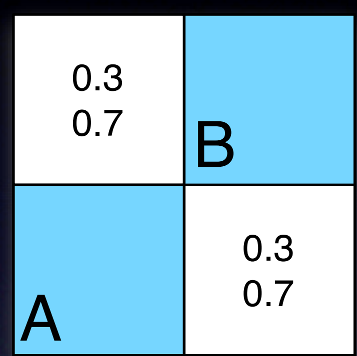

Unilateral

A → North
B → South
A Reward = 0.3
B Reward = 0.3
Total = 0.6

Bilateral

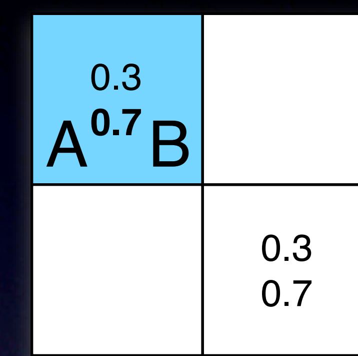

A → North
B → West
A Reward = 0.7/2=0.35
B Reward = 0.7/2=0.35
Total = 0.7

Unilateral / Bilateral Rewards

Initial

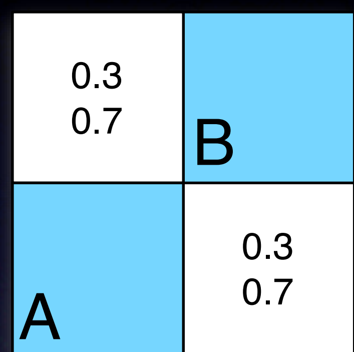

Unilateral

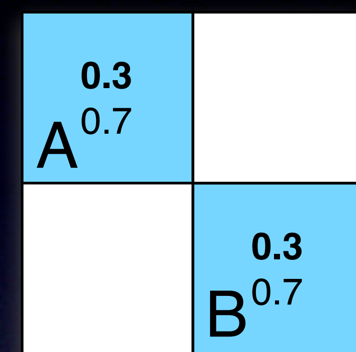

A → North
B → South
A Reward = 0.3
B Reward = 0.3
Total = 0.6

Bilateral

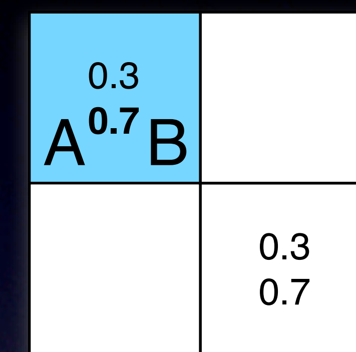

A → North
B → West
A Reward = 0.7/2=0.35
B Reward = 0.7/2=0.35
Total = 0.7

- If 3 agents cooperate
  - Two receive bilateral reward
  - One receives default unilateral reward

MAE (DEC-W-POMDP)

- Instance of DEC-W-POMDP (NEXP-complete)
  - **stochastic** since effects uncertain.
  - **Markovian** since new state is conditionally independent of the history given the current state and joint action of agents.
  - **partially observable** agents cannot perceive other agents neighbourhoods
  - **w-weakly** agents can perceive absolute location and their own local neighbourhood.
- For 6 agents, and horizon 2, each agent needs to evaluate $5^{24} = 6 \times 10^{16}$ possible effects.

Outline

- Introduction
- What is Multiagent Expedition?
- Collaborative Design Network
- **Graphical Model for Multiagent Expedition**
- Recursive Model for Multiagent Expedition
- Experimental Results & Discussion
- Conclusion & Future Work

Collaborative Design Network (CDN) [Xiang, Chen and Havens, AAMAS 05]

- A multiagent component-based design paradigm.
- CDN gives optimal design based on preferences of all agents.
- Scales linearly with the addition of agents.
- Efficient when the overall dependency structure is sparse.
- We use CDN in this work as a *collaborative decision network.*

Design Network (DN)

- is a DAG, $G = (V, E)$ where $V = D \cup T \cup M \cup U$
- D the set of design nodes.
  - design decisions
- T the set of environmental nodes.
  - uncertainty over working environment of the product under design
- M the set of performance nodes.
  - refers to **objective** measures of functionality of the design.
- U the set of utility nodes.
  - **subjective** measures dependent strictly on performance nodes.

DN Continued...

- Syntactically each node is associated with a conditional probability distribution.
- Semantically, the nodes differ. E.g $P(d|\pi(d))$ encodes a design constraint.
- The goal is to find a design $\mathbf{d}^*$ which $EU(\mathbf{d}^*)$ is maximal.

Collaborative Design Network (CDN)

- Collaborative design network extends multiply sectioned Bayesian networks to multiagent decision making.
  - DAG domain structuring.
  - Hypertree agent organization.
  - Belief over private and shared variables.
  - Partial evaluation of partial design communicated over small set of shared variables btw agents.
  - Design is globally optimal.
  - Local design at each agent remains private.

CDN for MAE

- Each time-step an agent:
  - Utilizes a dynamic graphical model.
  - Updates domains for movement and position nodes.
  - Updates utility distributions from locally observed rewards.
  - Communicates with other agents to find globally optimal joint action.

Position Nodes

- Encode probability of uncertain location $ps^{x,1}$ given agent movement $mv^{x,i}$.

$mv^{x,i}$	$ps^{x,1} = (0,0)$	$ps^{x,1} = (1,0)$	$ps^{x,1} = (-1,0)$	$ps^{x,1} = (0,1)$	$ps^{x,1} = (0,-1)$
$north$	0.025	0.025	0.025	0.9	0.025
$south$	0.025	0.025	0.025	0.025	0.9
$east$	0.025	0.9	0.025	0.025	0.025
$west$	0.025	0.025	0.9	0.025	0.025
$halt$	0.9	0.025	0.025	0.025	0.025

Utility Nodes

$ps_A^{x,1}$	$ps_B^{x,1}$	$ps_C^{x,1}$	$P(rw^{ABC,1} = y \mid ps_A^{x,1}, ps_B^{x,1}, ps_C^{x,1})$
(0,0)	(0,0)	(0,0)	0.4
(0,0)	(0,0)	(1,0)	0.2
(0,0)	(0,0)	(2,0)	0.1
	$\vdots$		$\vdots$
(-2,-2)	(-2,-2)	(-2,0)	0
(-2,-2)	(-2,-2)	(-2,-1)	0.3
(-2,-2)	(-2,-2)	(-2,-2)	1.0

Applying CDN to MAE

- Encode movements {*N*, *S*, *E*, *W*, *H*} in **design nodes**.
- Encode uncertain locations given agent movements through **performance nodes.**
- Encode reward in **utility nodes**.
- Communicate *EU* over design nodes btw agents to find maximal utility design which corresponds to the globally optimal joint plan.

Graphical Model

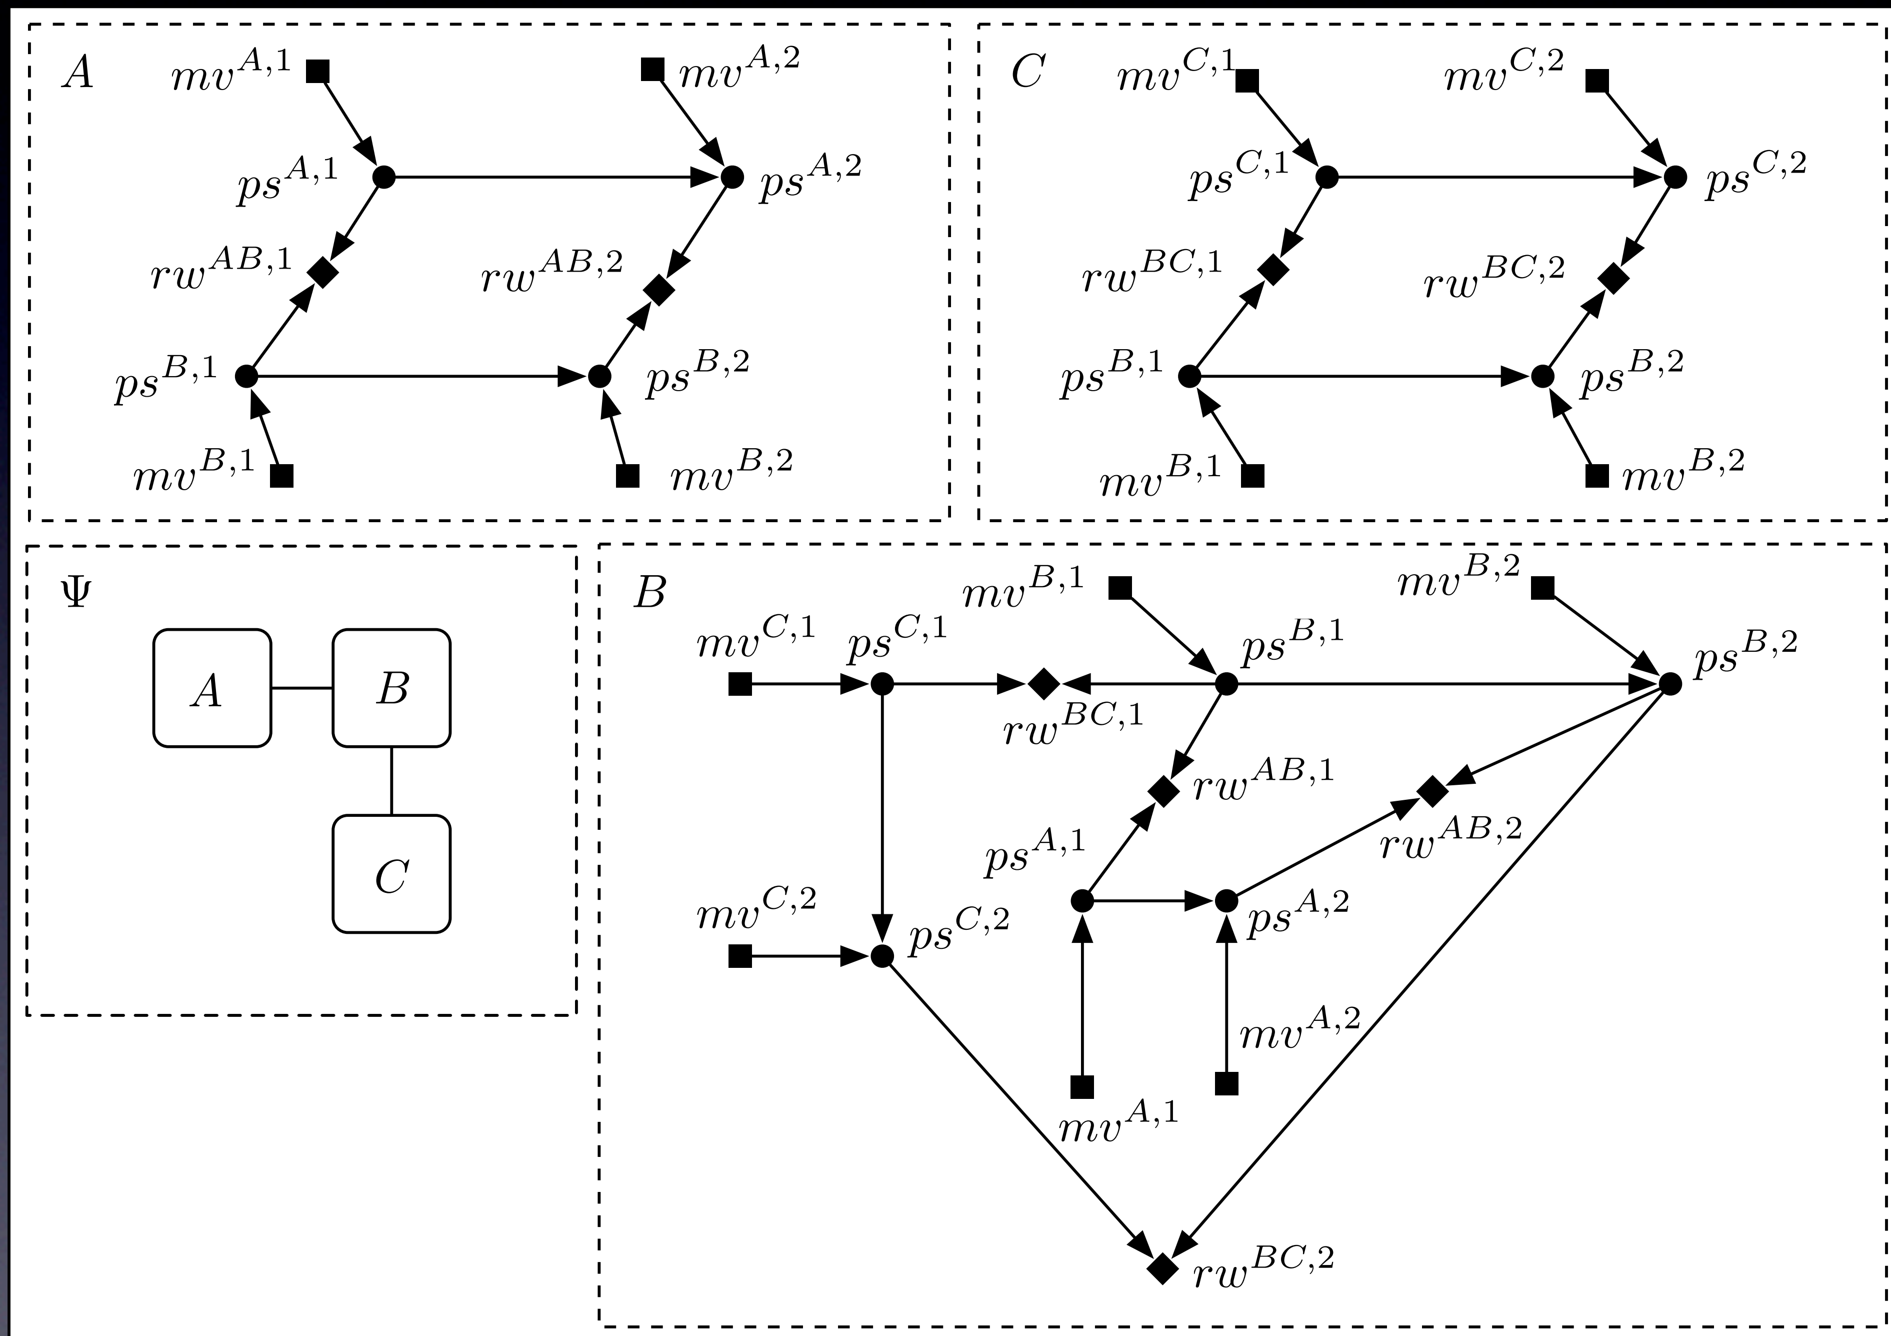

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Graphical Model

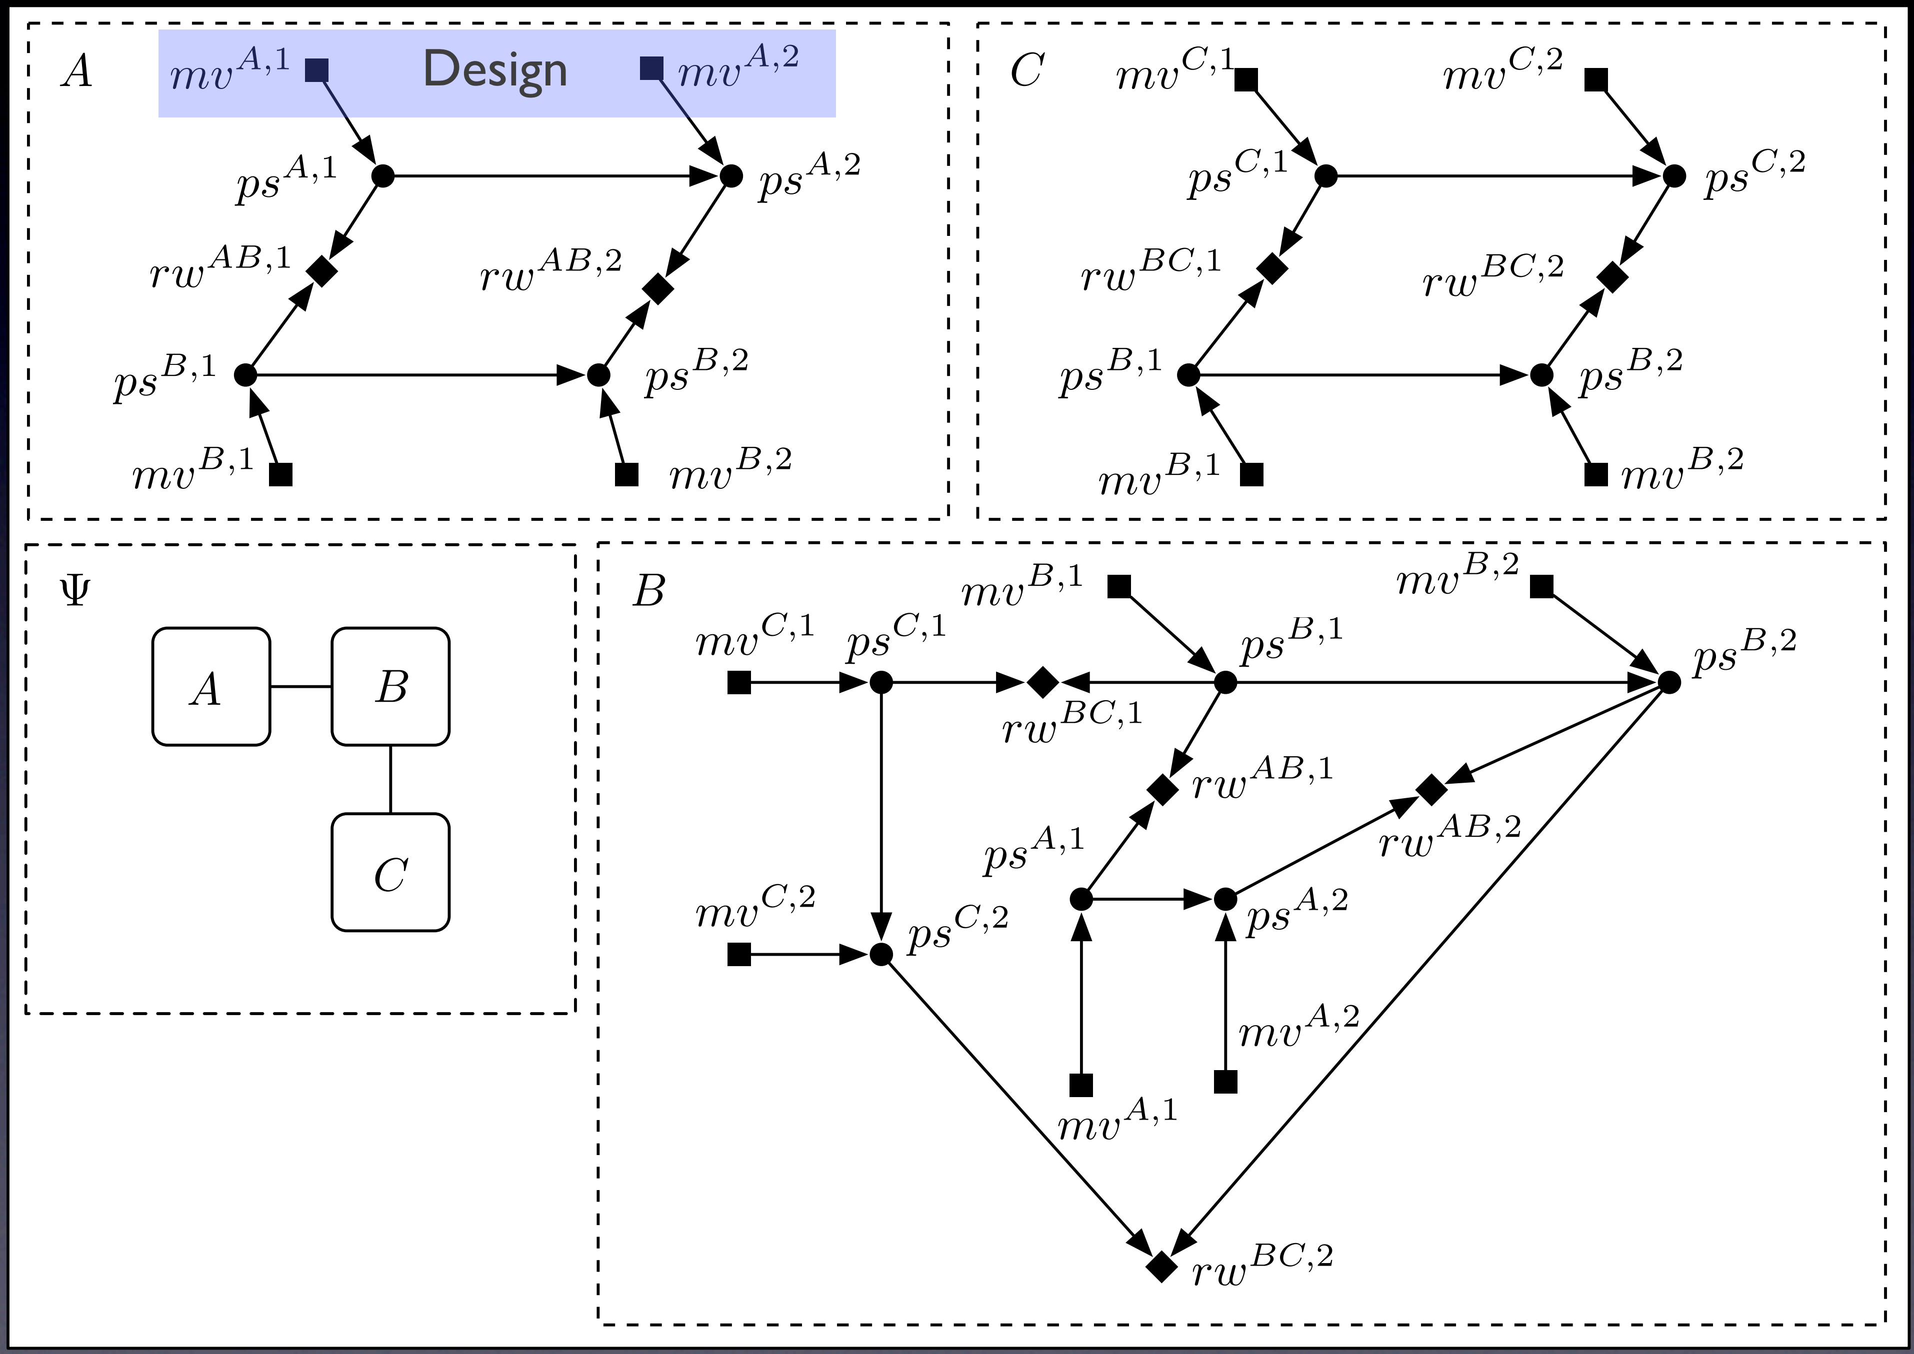

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Graphical Model

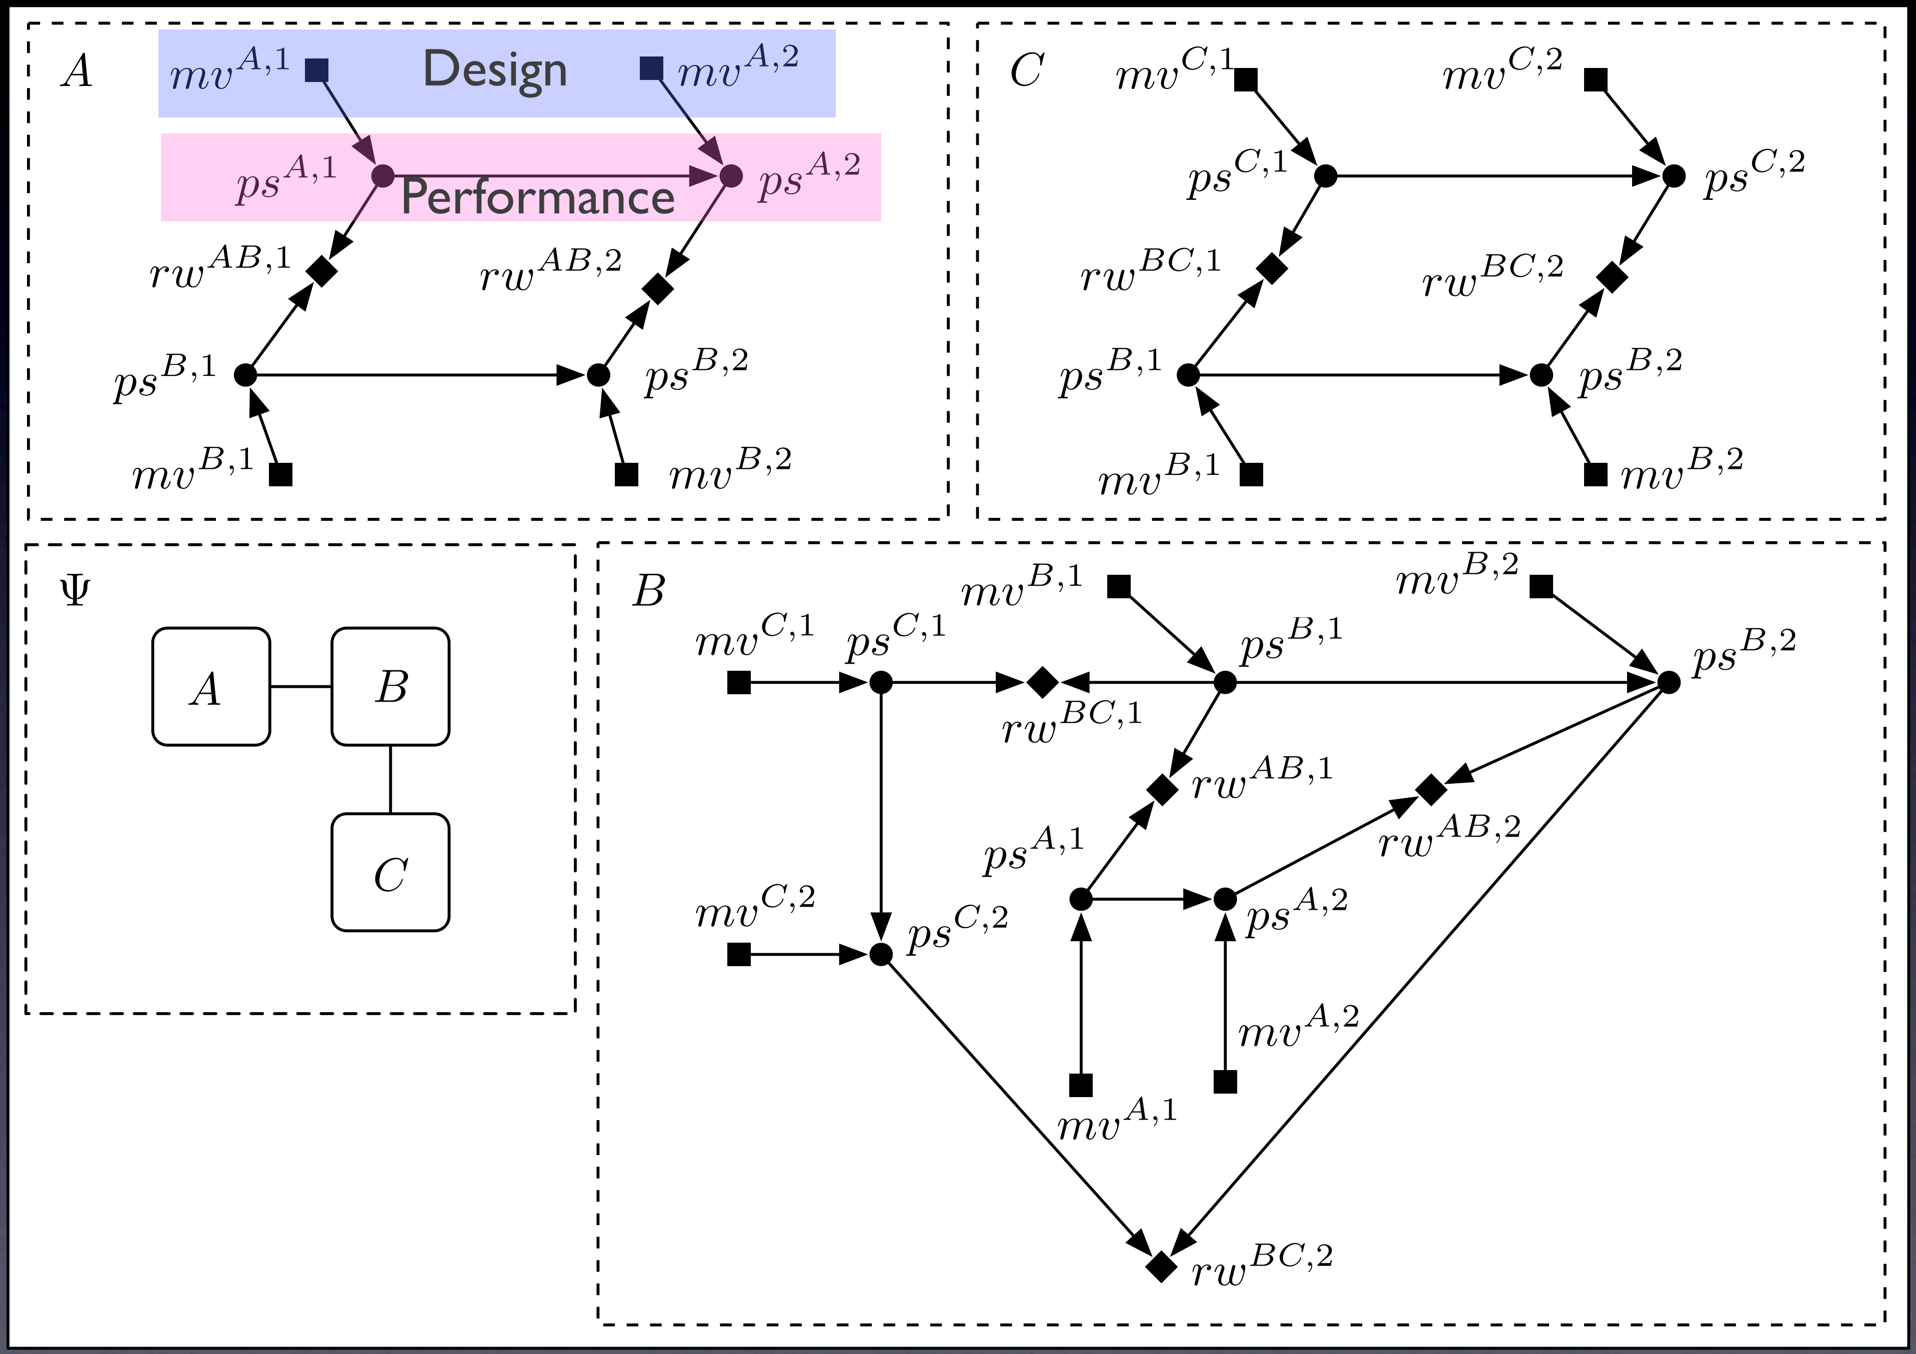

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Graphical Model

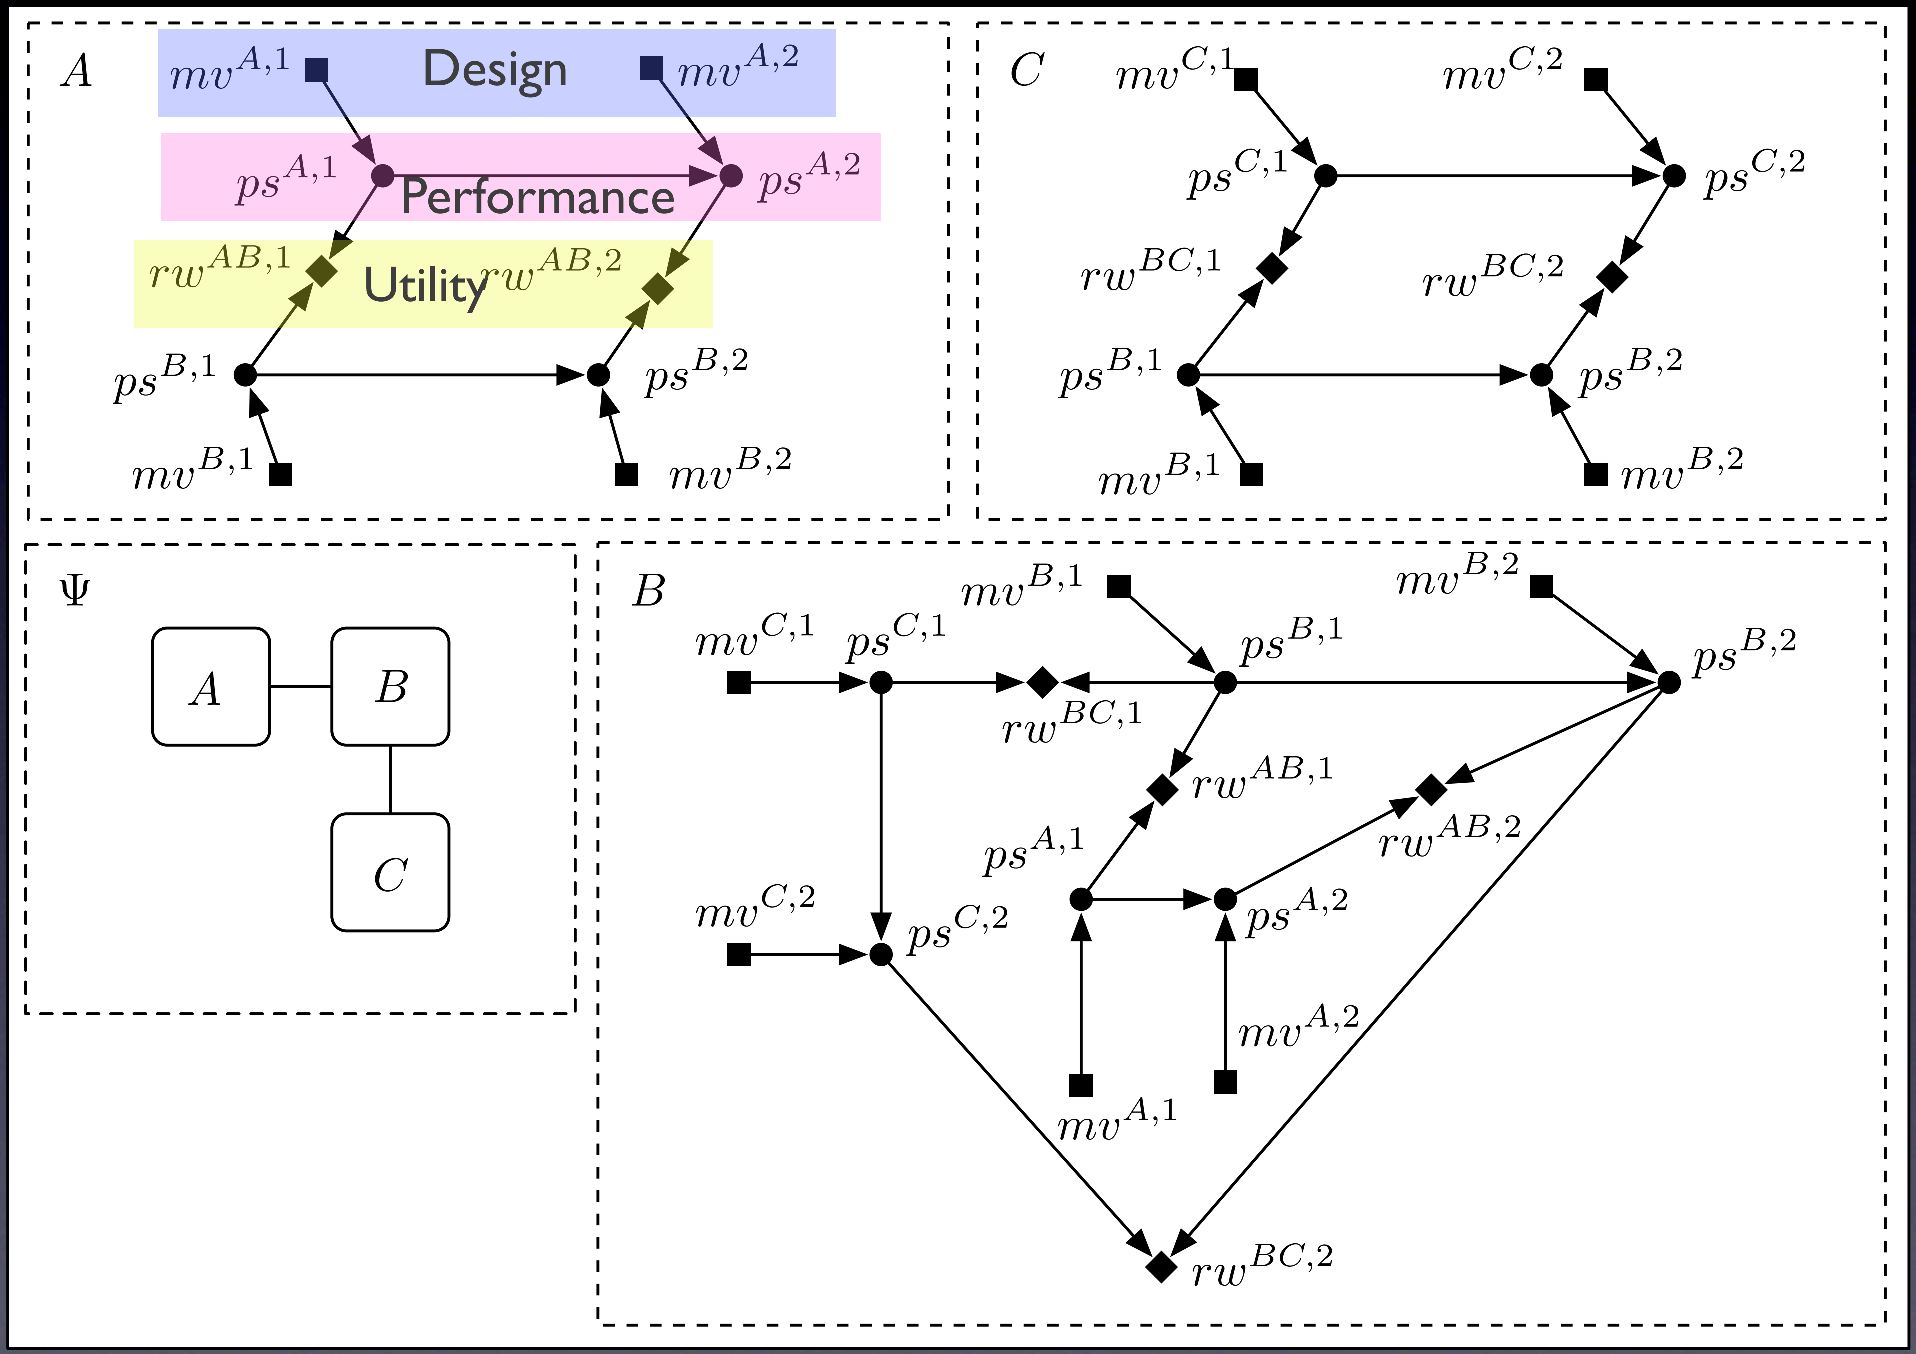

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Graphical Model

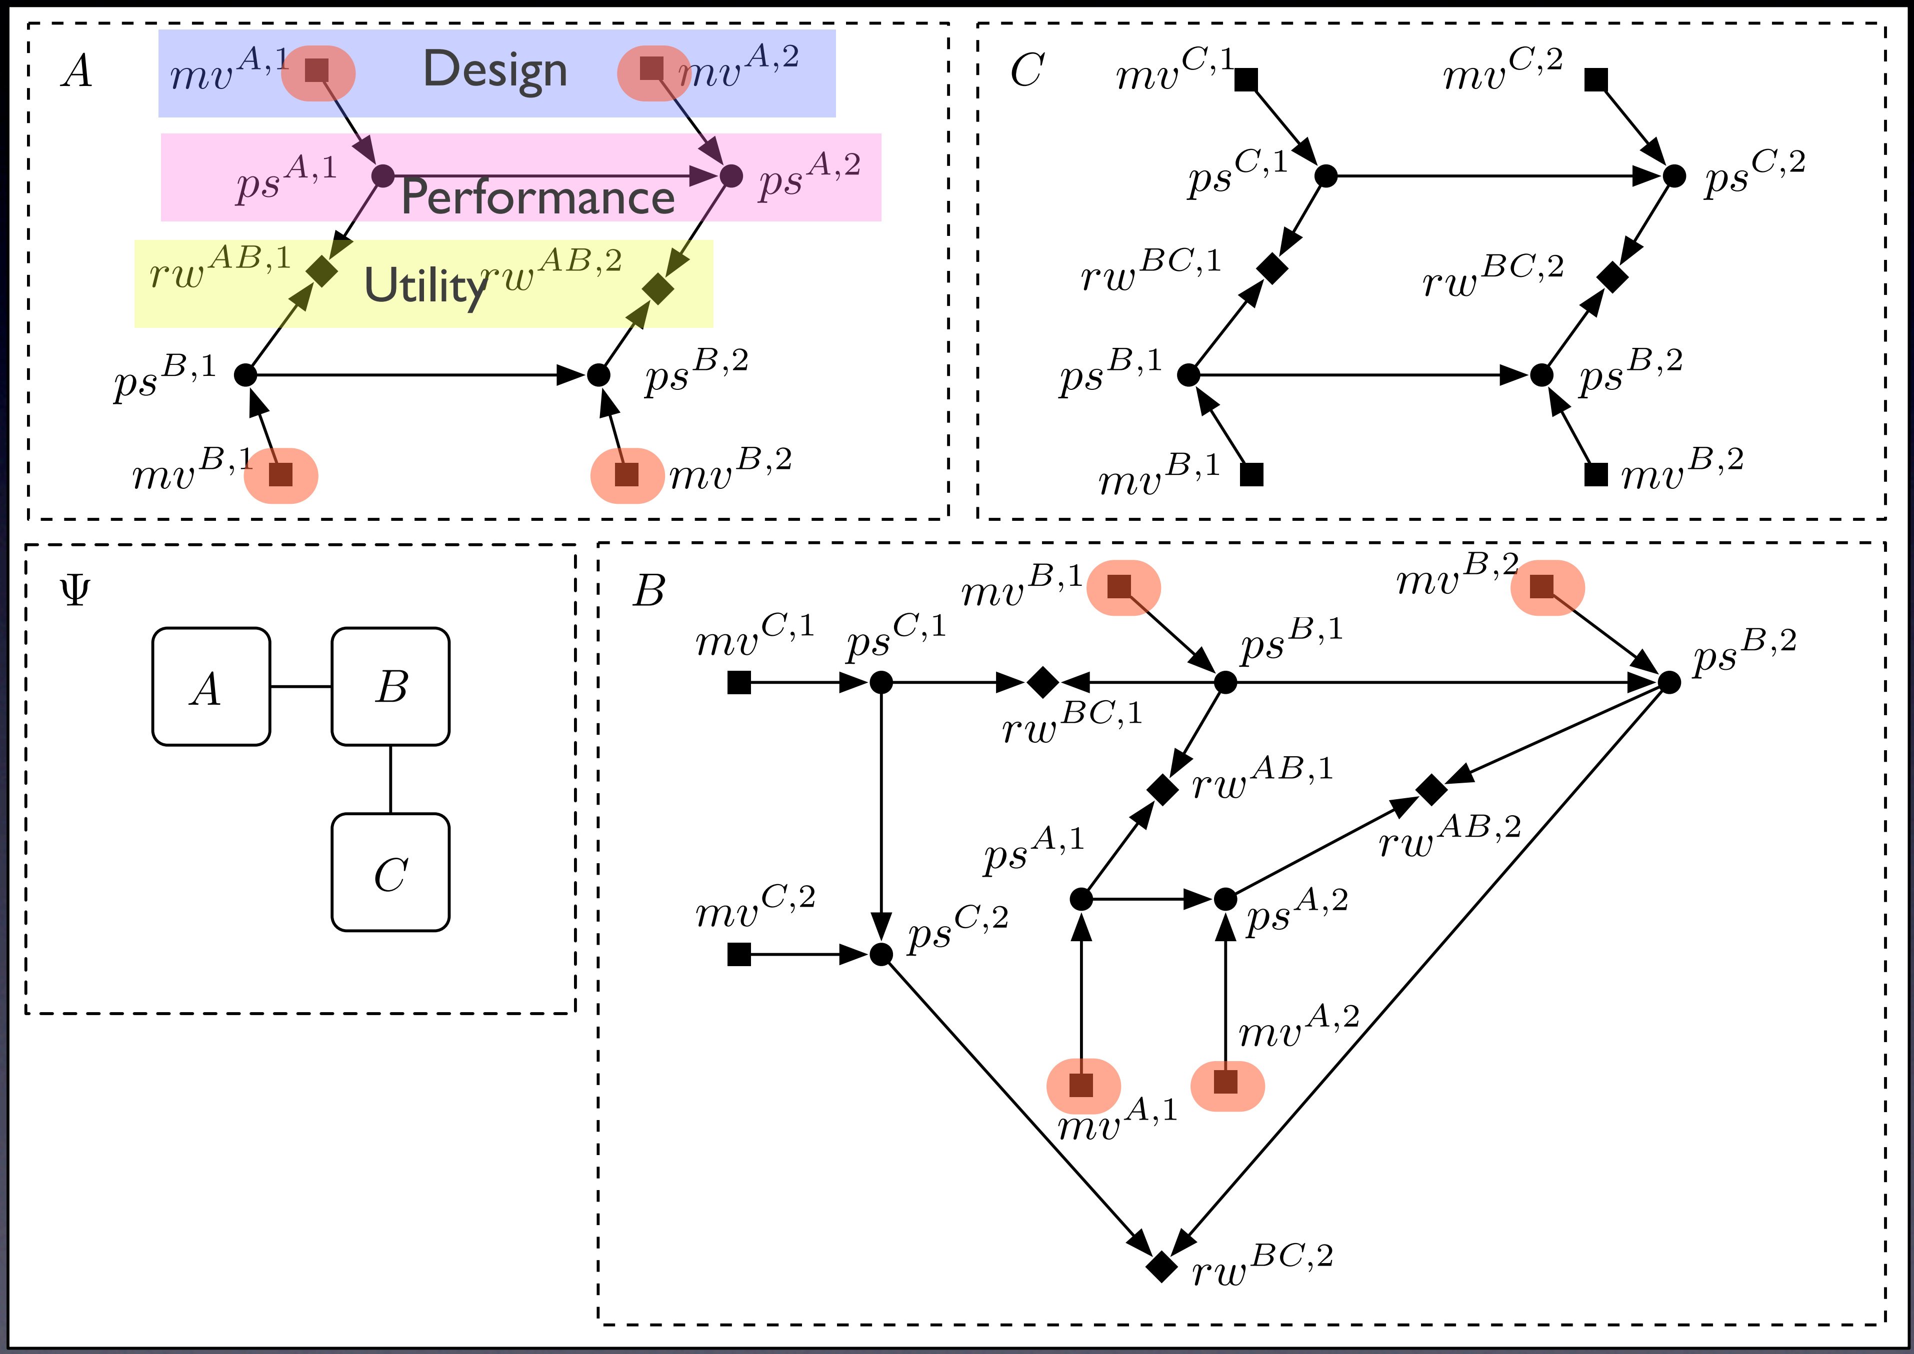

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Graphical Model

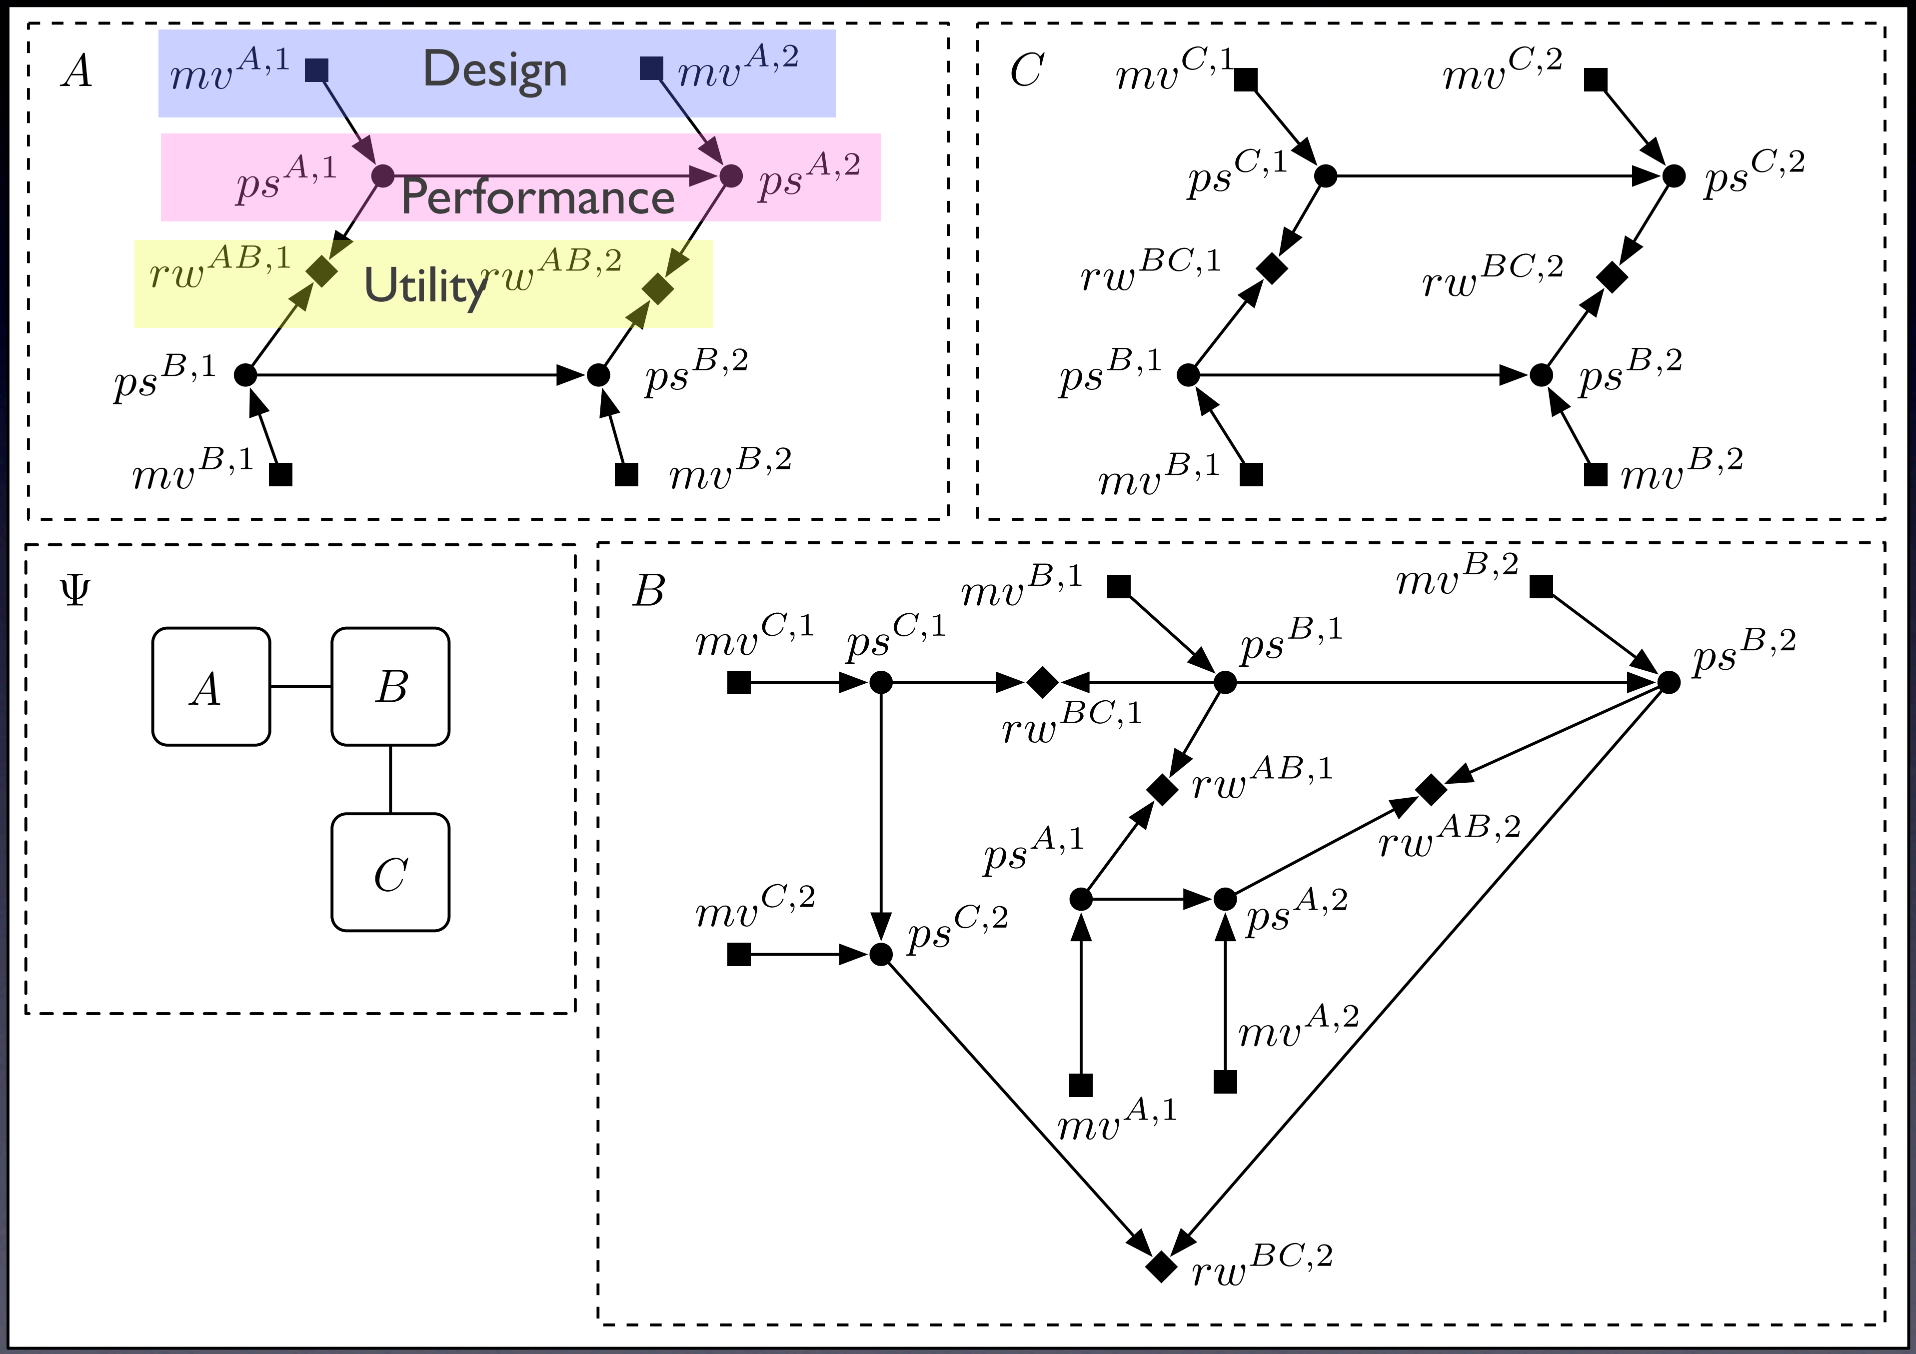

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Graphical Model

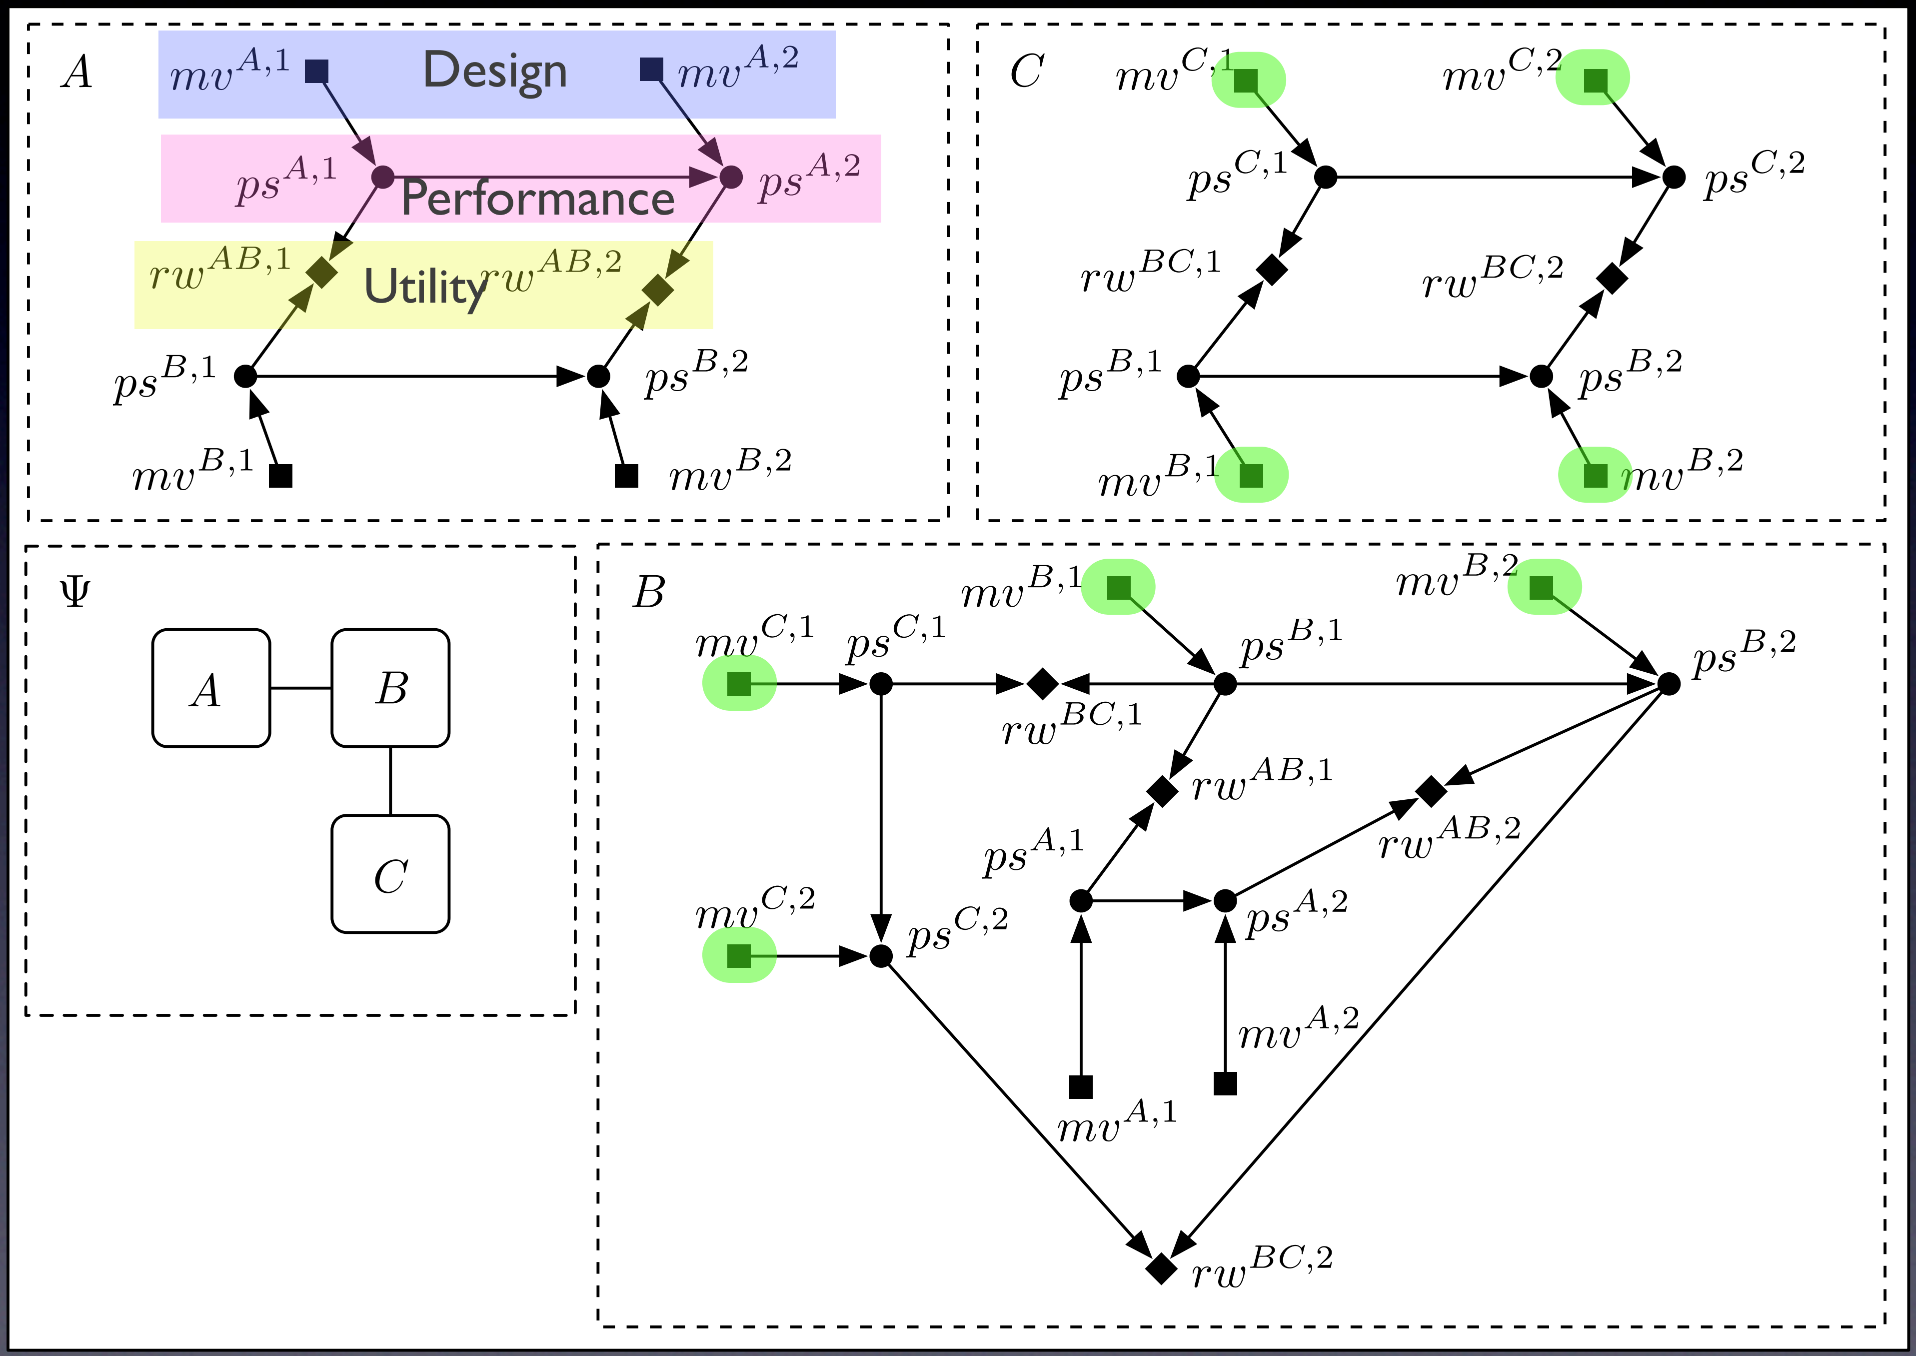

CDN of a 3 agent group (A,B,C) for expedition/planning, where Ψ is the hypertree.

Outline

- Introduction
- What is Multiagent Expedition?
- Collaborative Design Network
- Graphical Model for Multiagent Expedition
- **Recursive Model for Multiagent Expedition**
- Experimental Results & Discussion
- Conclusion & Future Work

Recursive Modeling Method (RMM) [Gmytrasiewicz et al. IEEE 98]

- Loosely-coupled multiagent decision making paradigm.
  - No explicit communication btw agents.
- Matrix-based agent representation.
- Agents model other agents in order to coordinate actions.
- Agents have probability distributions over other agent's models.

Payoff Matrix Representation

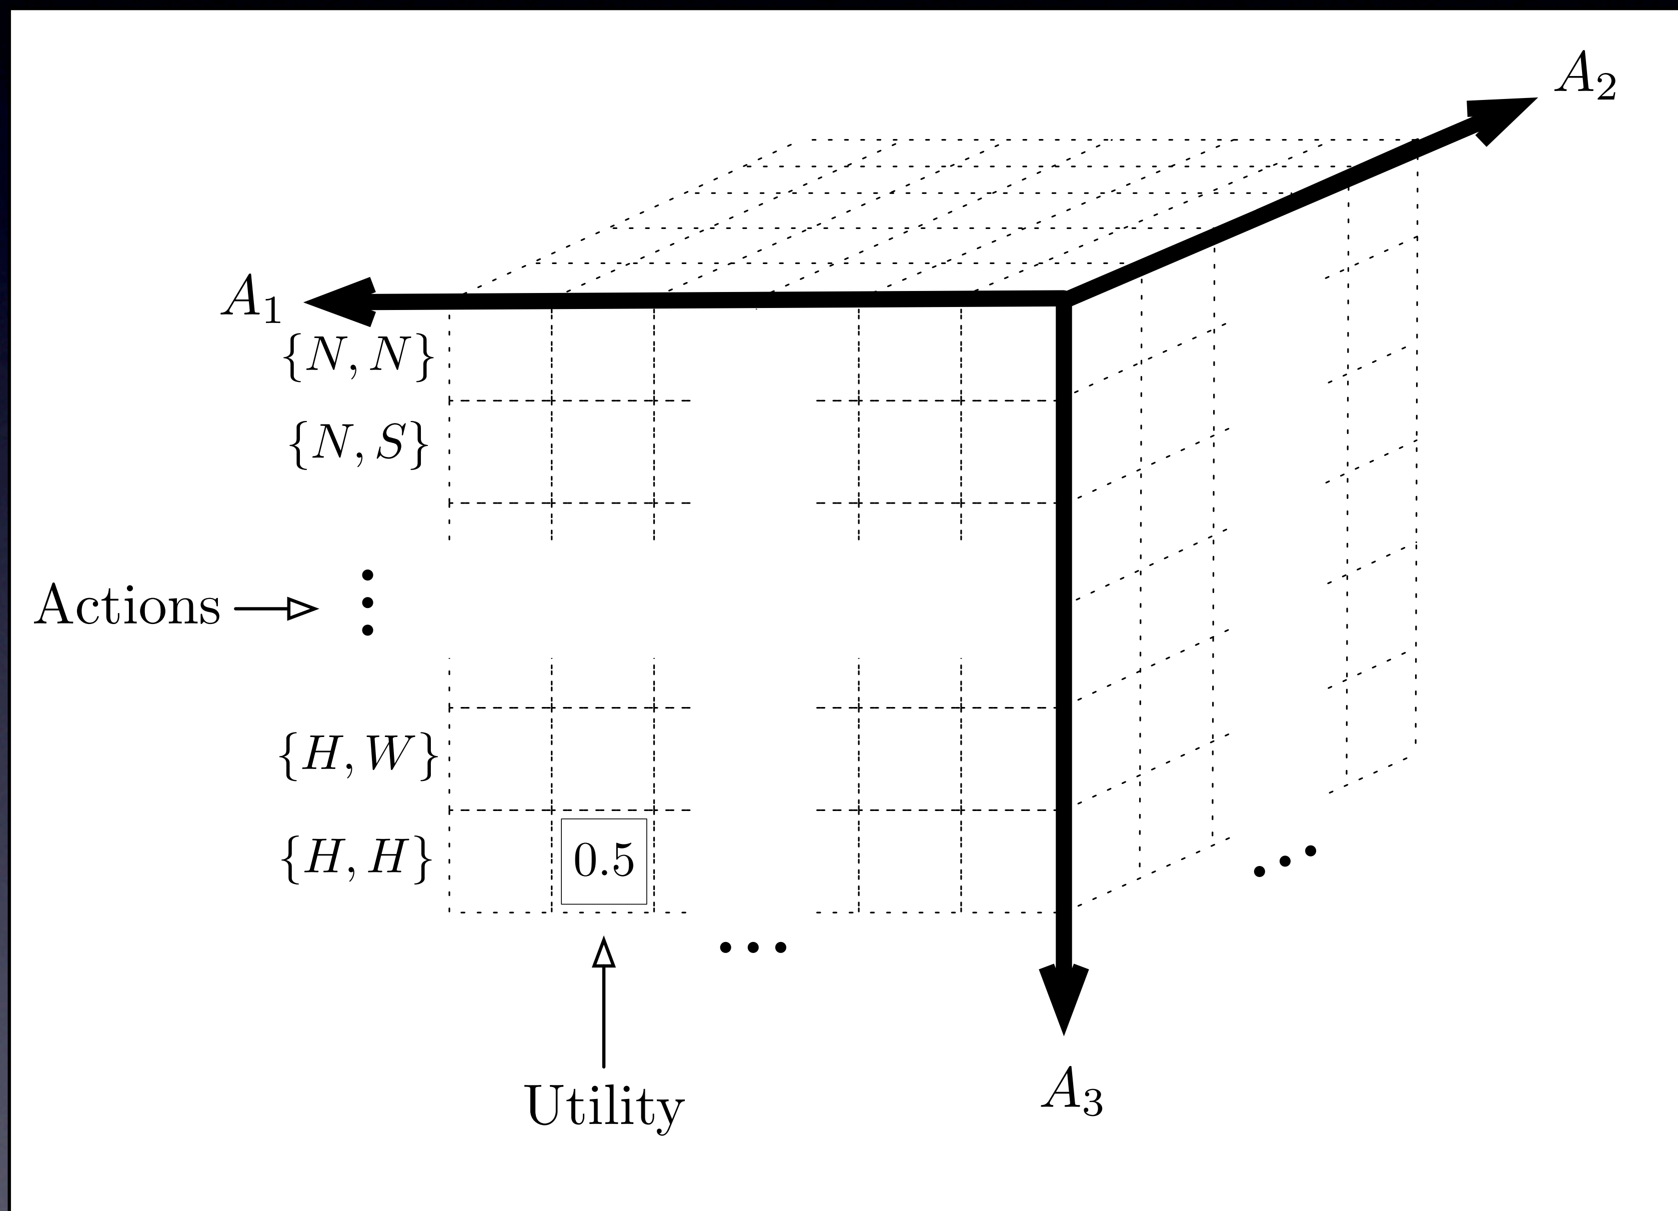

Payoff Matrix Representation

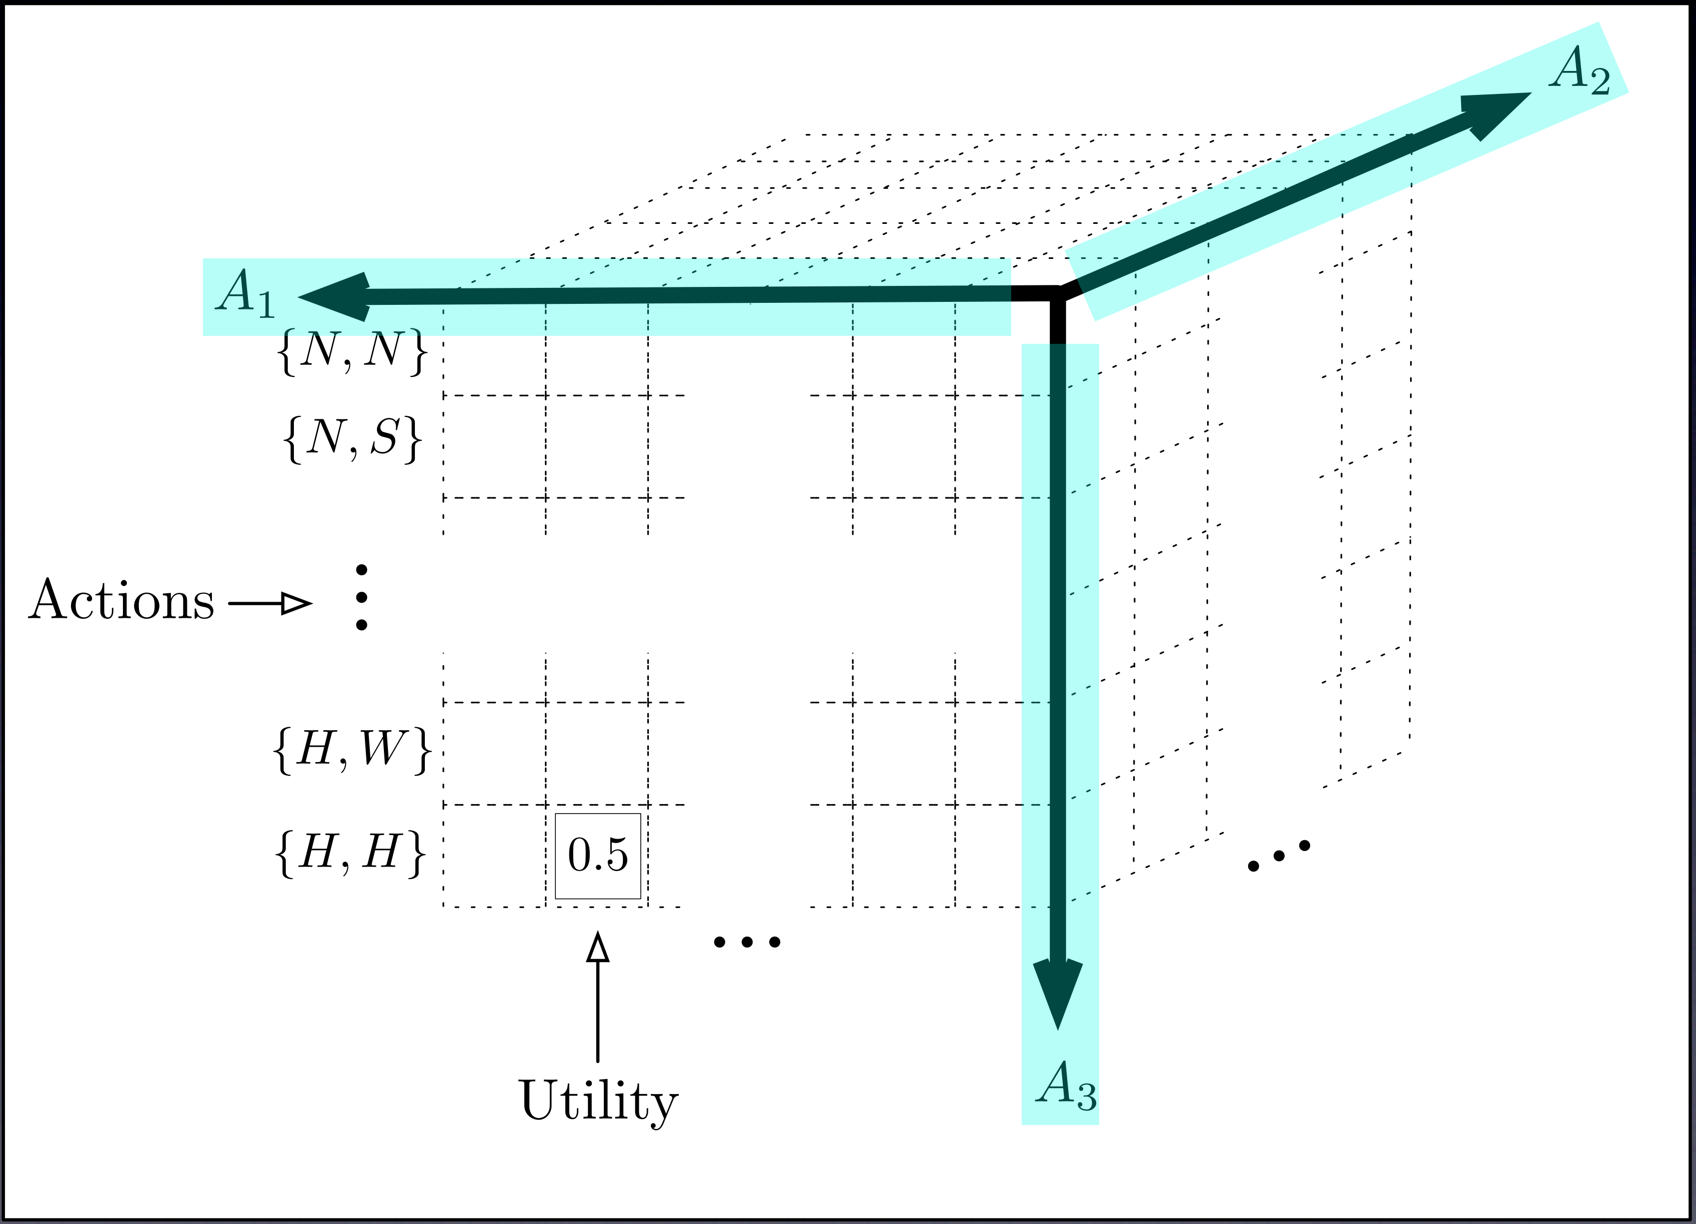

Payoff Matrix Representation

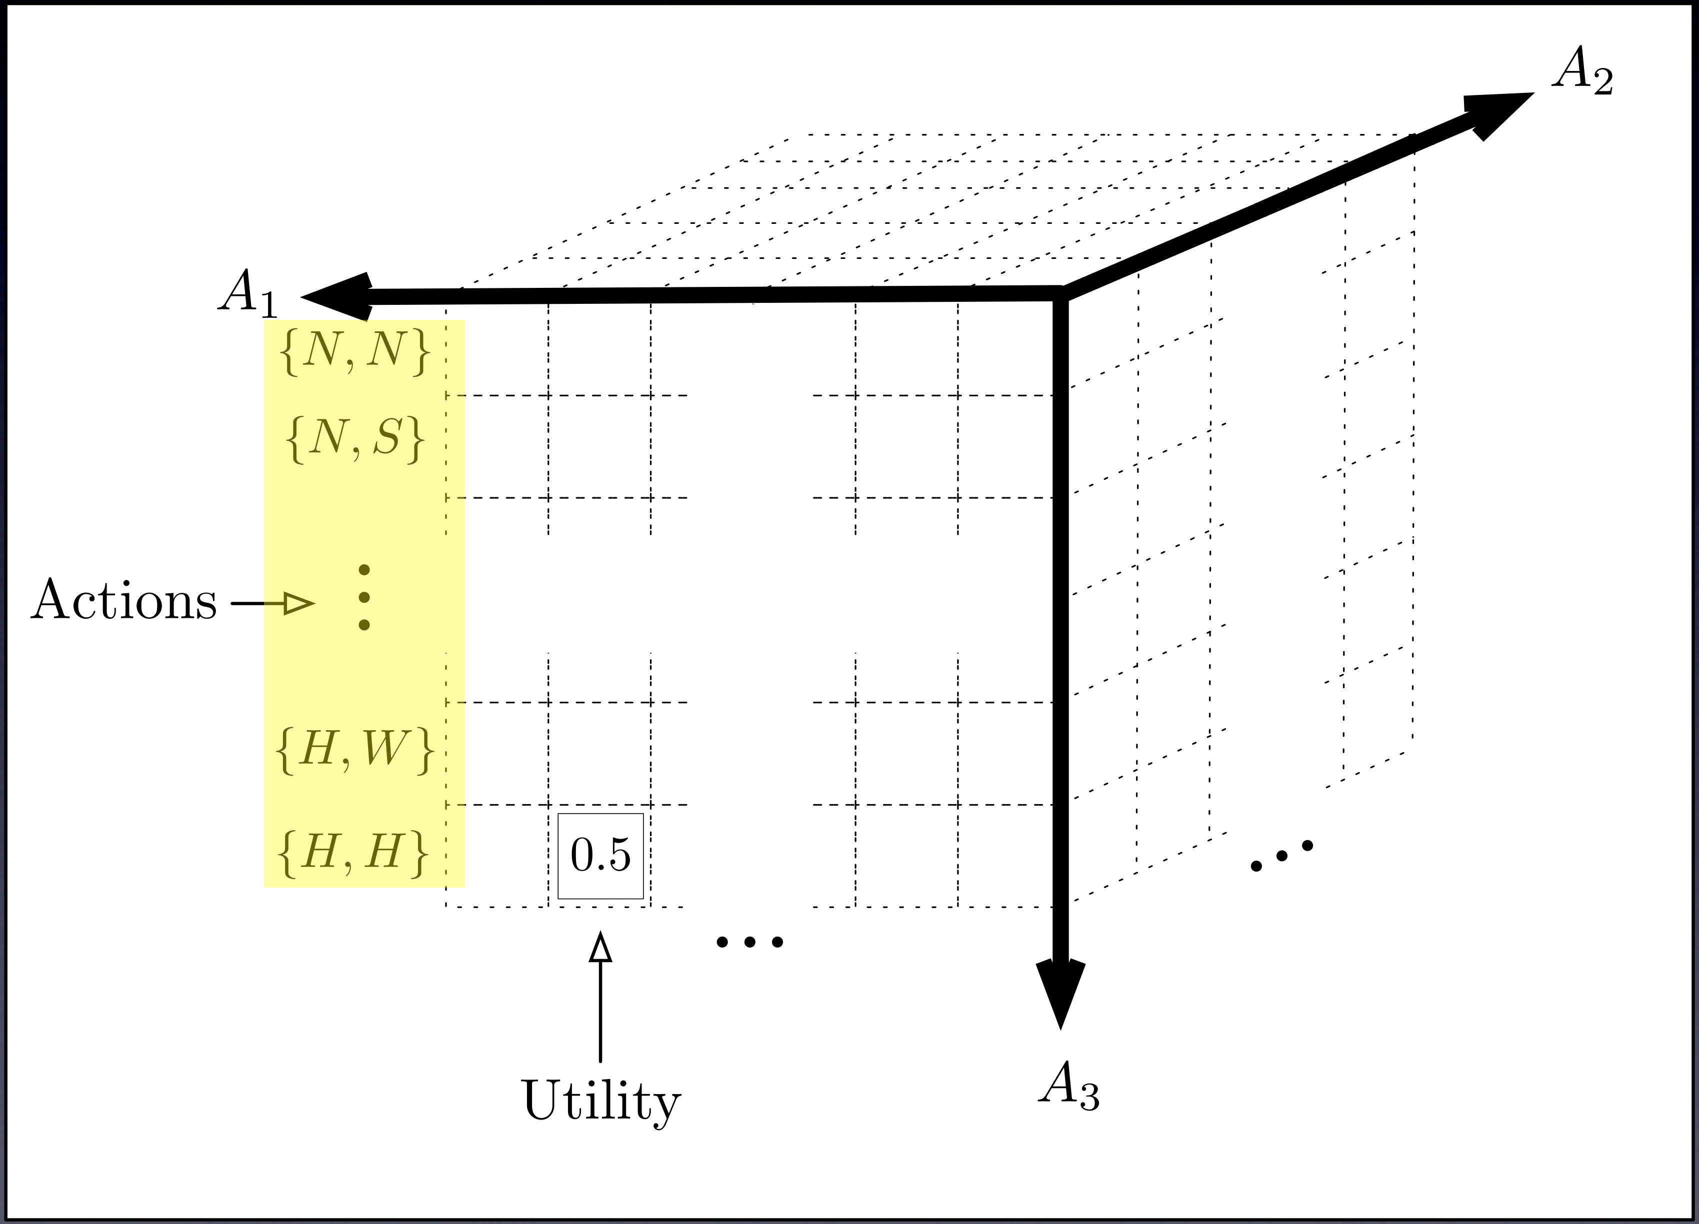

Payoff Matrix Representation

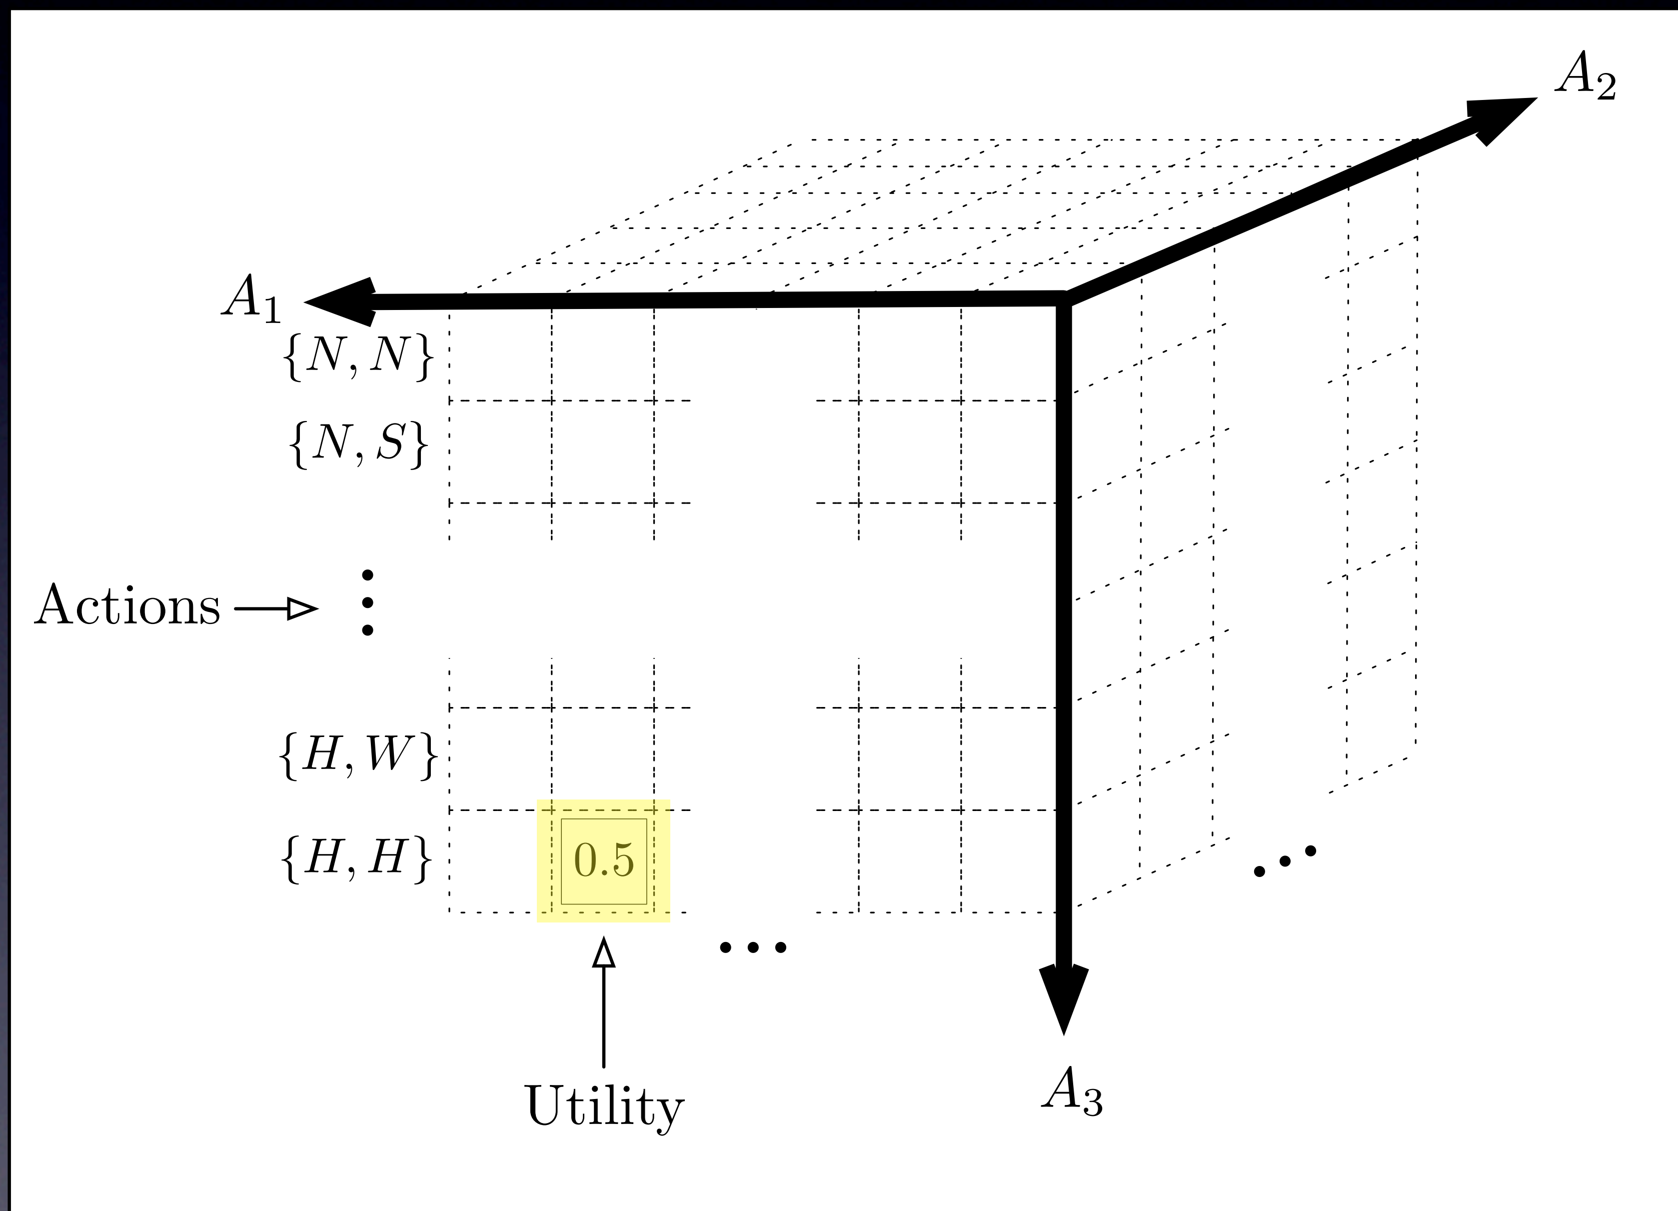

Payoff Matrix Representation

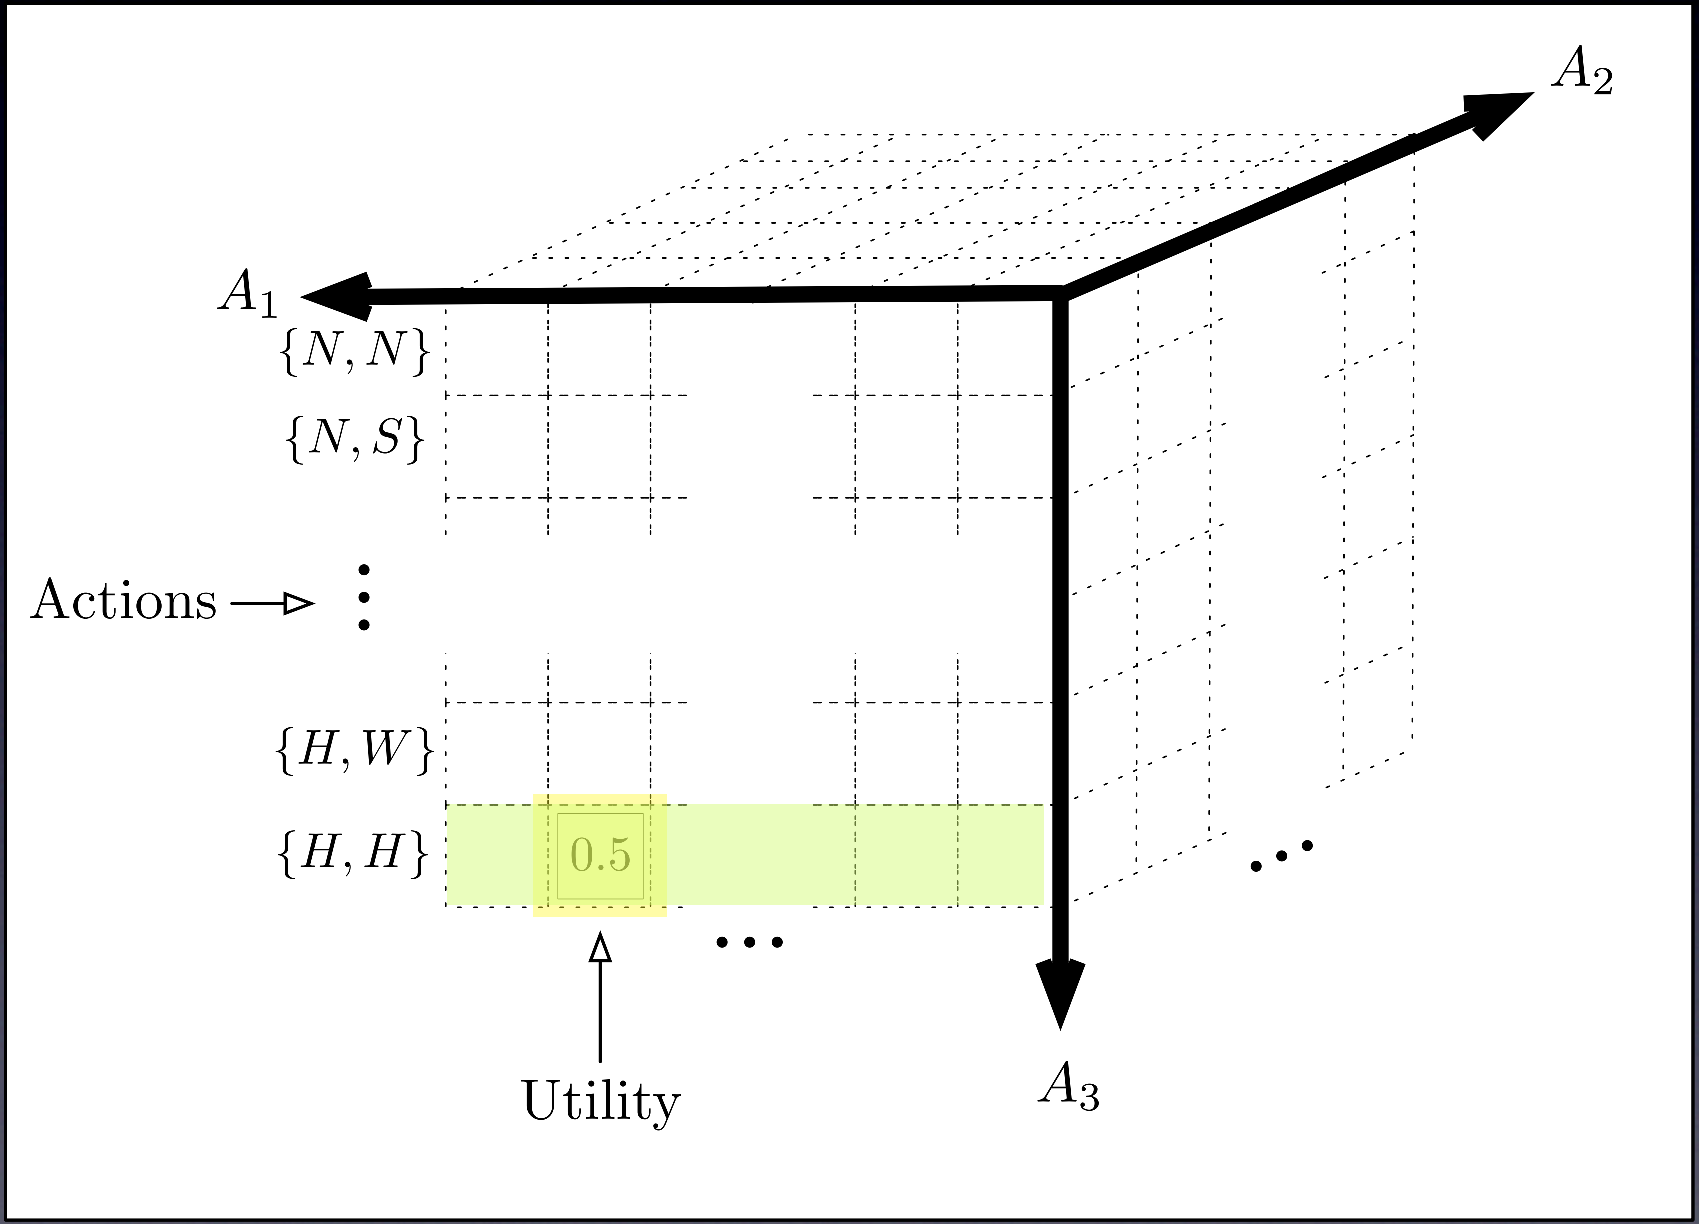

Payoff Matrix Representation

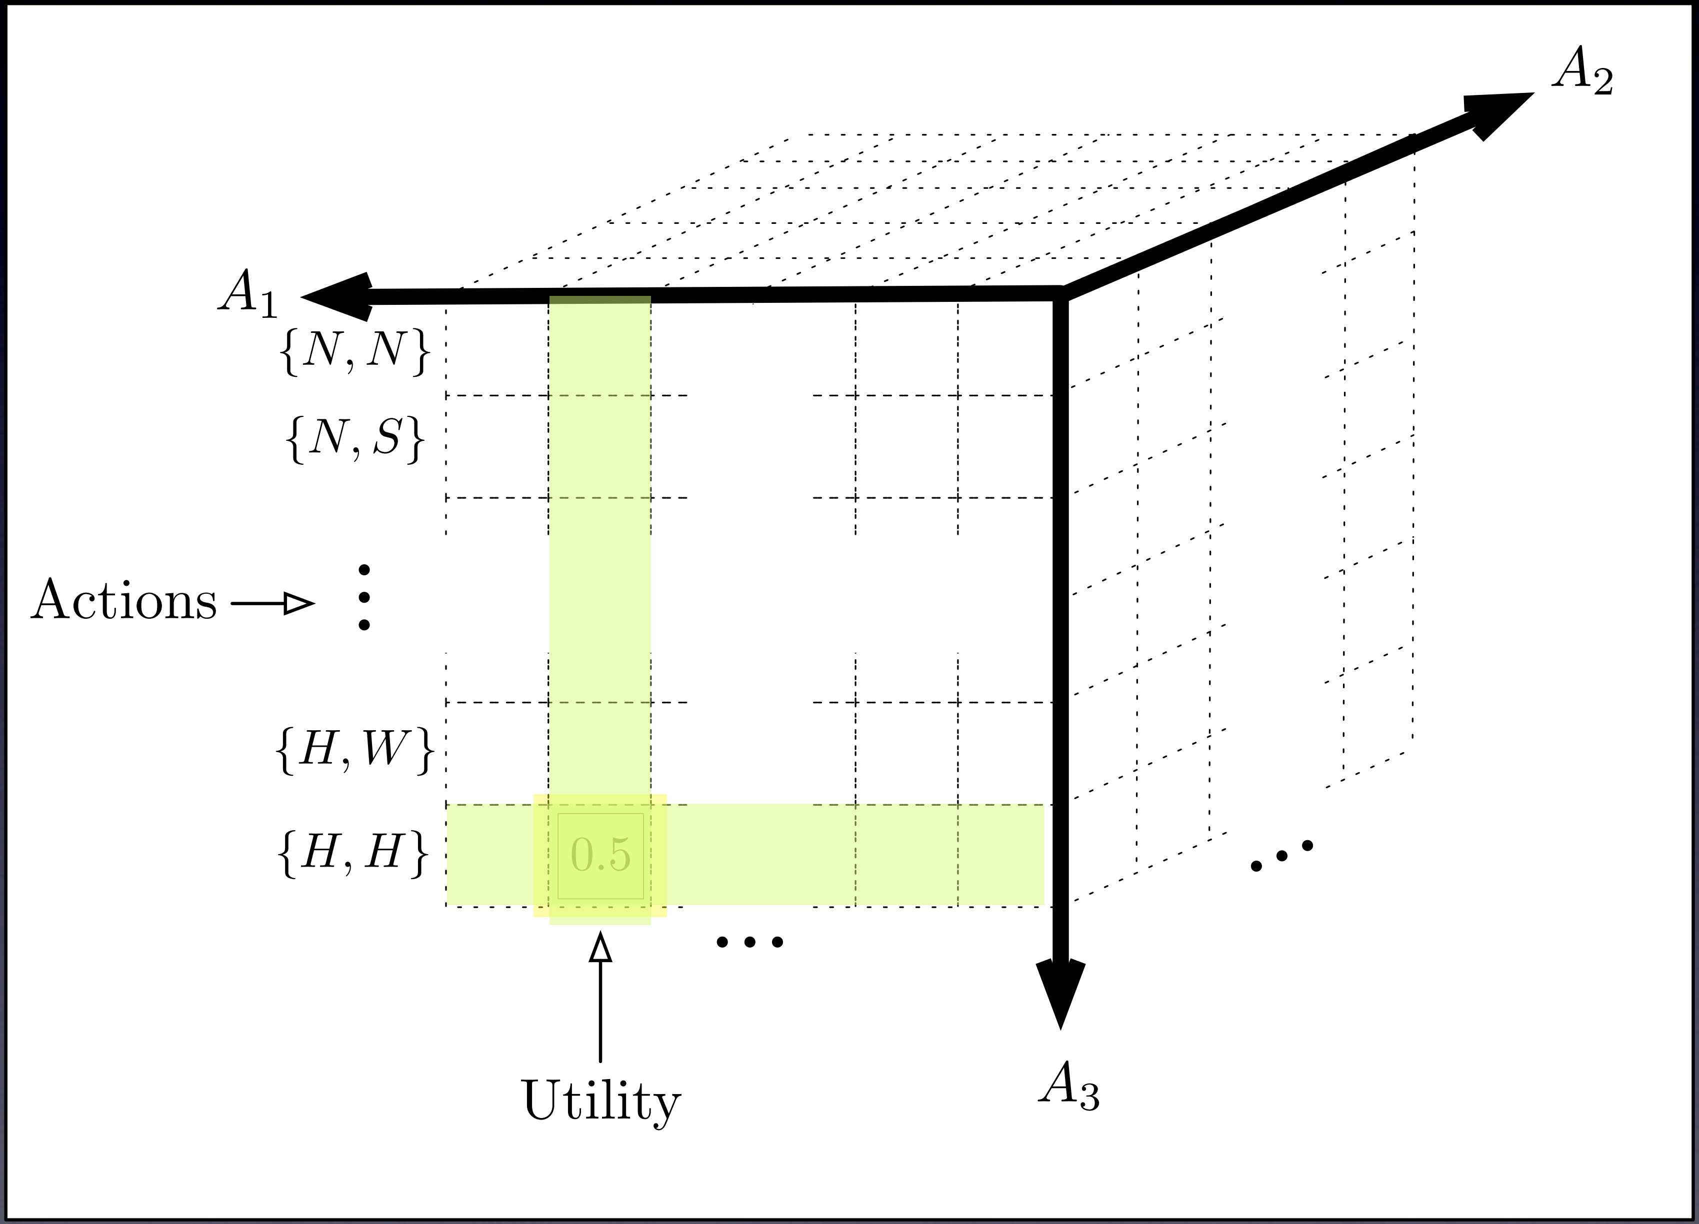

Payoff Matrix Representation

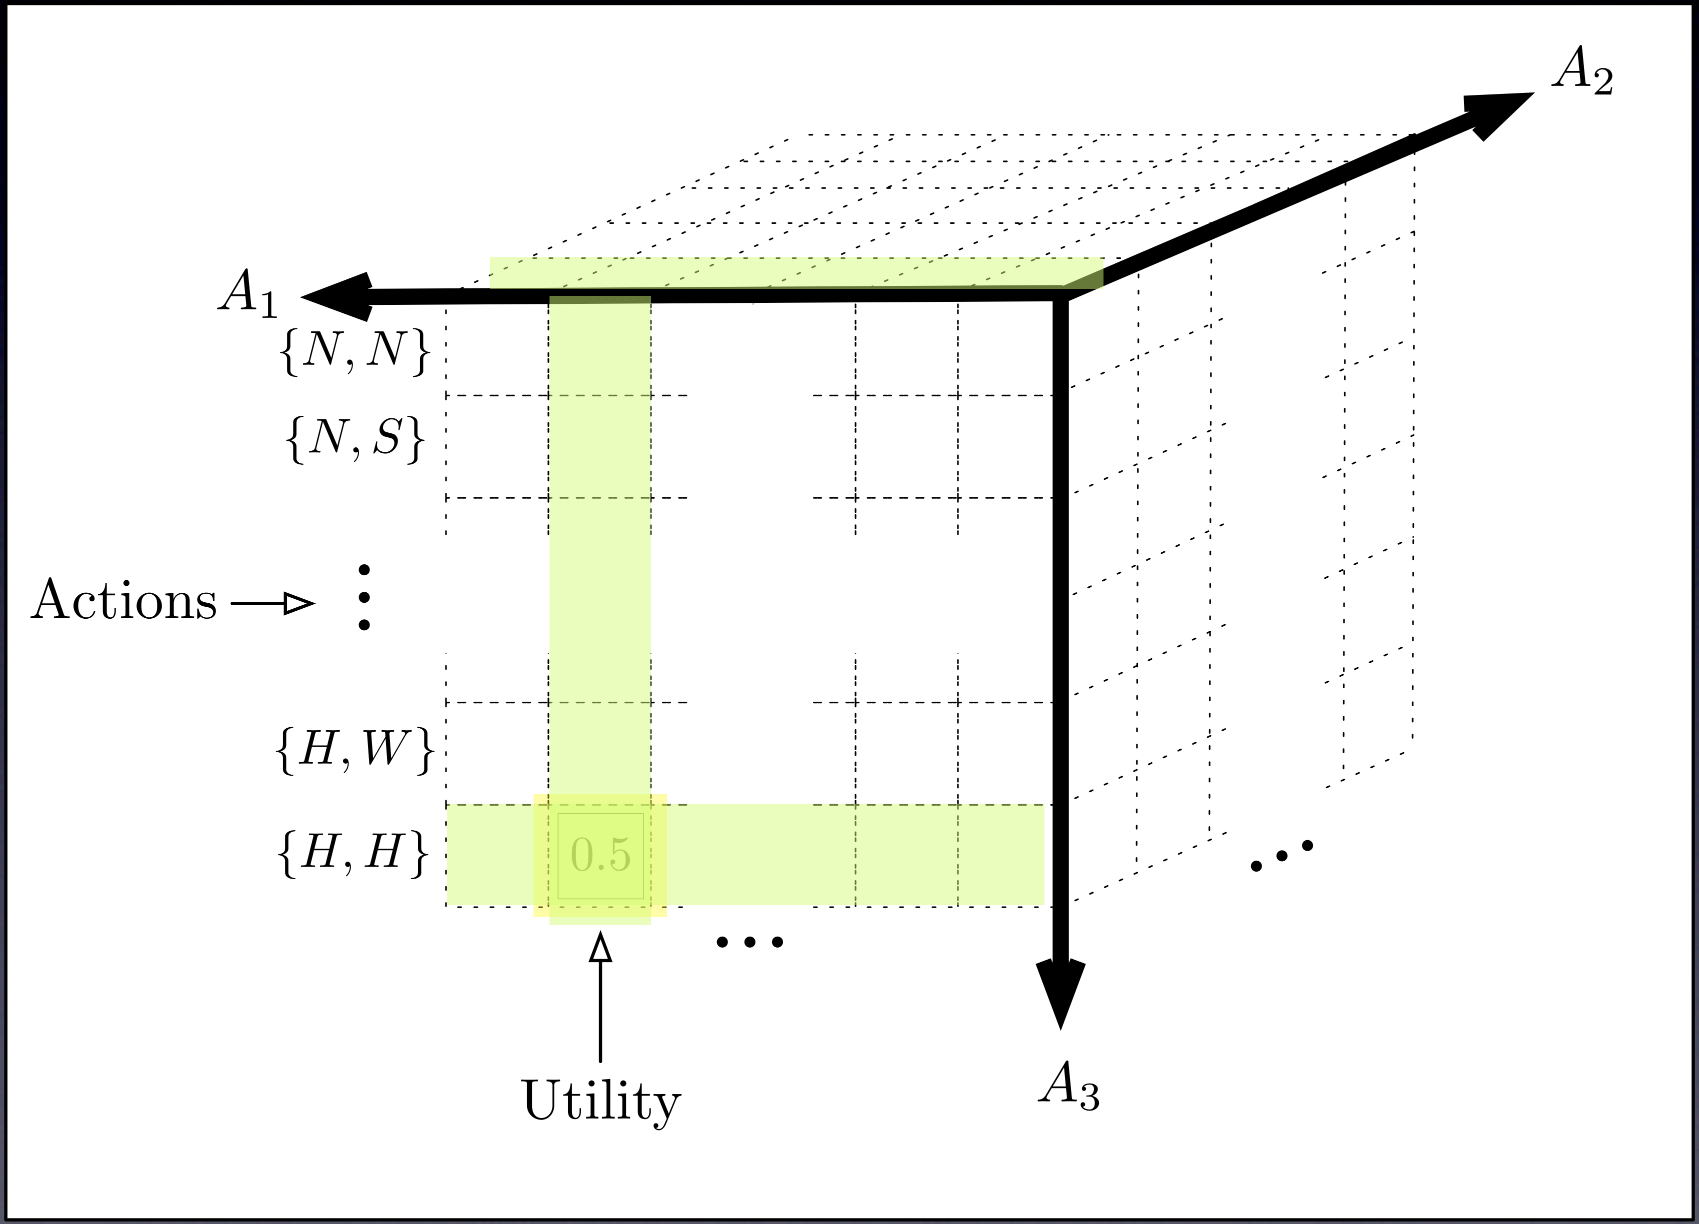

RMM for MAE

- Observations are local to each agent in MAE.
- How does RMM evaluate joint actions of agents when some payoffs of other agents are unknown?

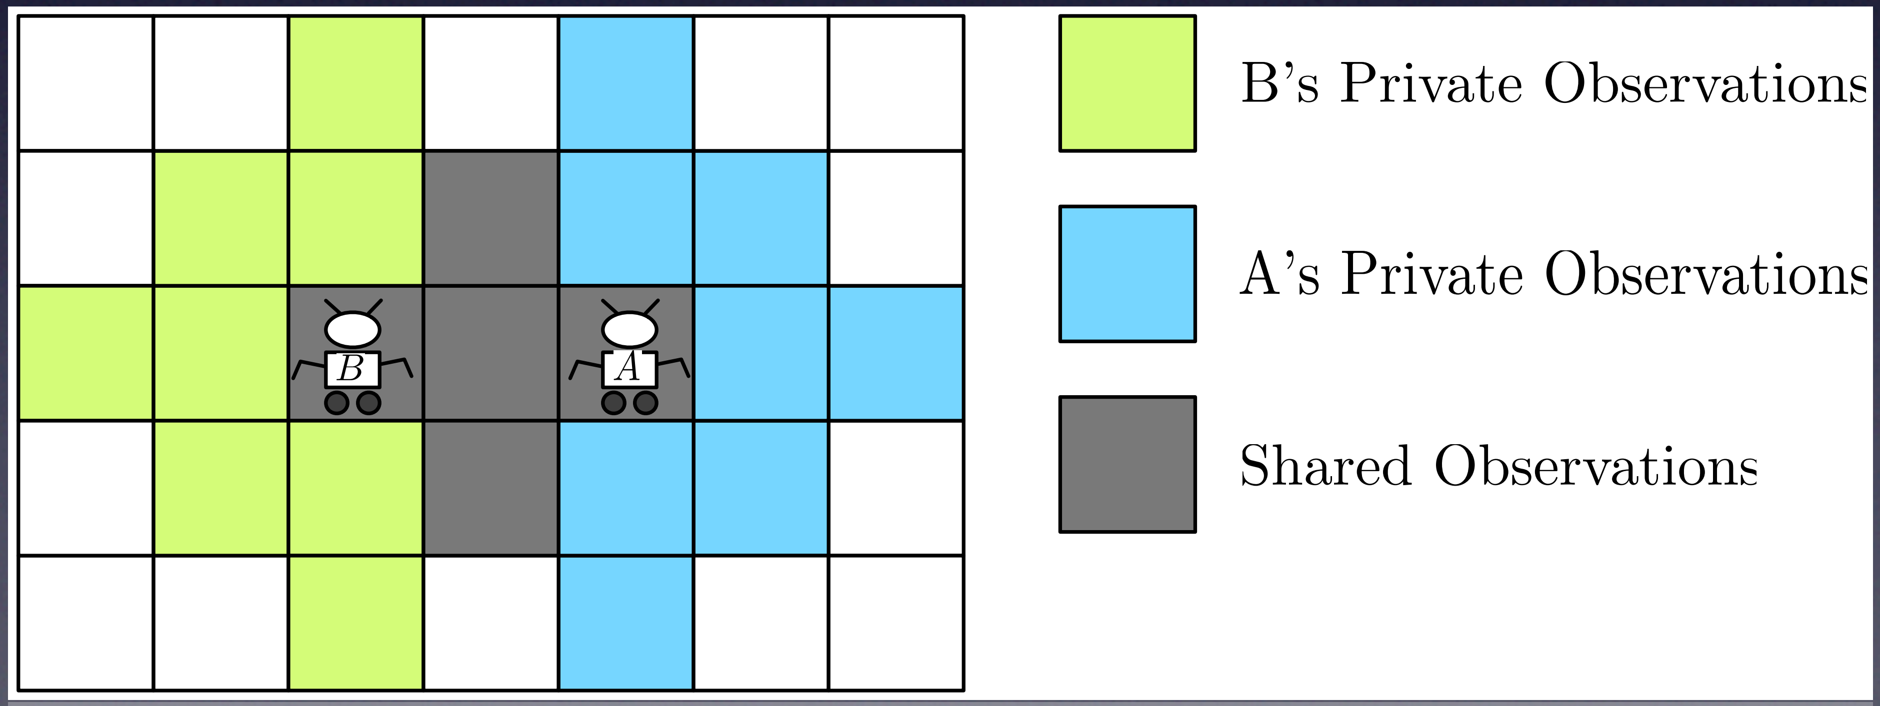

RMM for MAE cont.

- Knowing the correct state that an agent is in allows for successful planning.
- **Idea:** Agents model other agents states in the RMM tree.
- A state categorizes a neighbourhood payoff.
- Based on past observations of agent actions, update belief on the state of neighbouring agents.

Recursive Model Structure

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \mathbf{allLow}, nbp_C^B = \mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \mathbf{allLow}, nbp_C^B = \neg\mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\mathbf{allLow}, nbp_C^B = \mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\mathbf{allLow}, nbp_C^B = \neg\mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.70
(N,N)	(N,S)	(N,N)	0.70
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

Recursive Model Structure

Payoff Matrix

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \textbf{allLow}, nbp_C^B = \textbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \textbf{allLow}, nbp_C^B = \neg\textbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\textbf{allLow}, nbp_C^B = \textbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\textbf{allLow}, nbp_C^B = \neg\textbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.70
(N,N)	(N,S)	(N,N)	0.70
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

Recursive Model Structure

Payoff Matrix

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \mathbf{allLow}, nbp_C^B = \mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \mathbf{allLow}, nbp_C^B = \neg\mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\mathbf{allLow}, nbp_C^B = \mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\mathbf{allLow}, nbp_C^B = \neg\mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.70
(N,N)	(N,S)	(N,N)	0.70
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

2^n models for missing observations

Recursive Model Structure

Payoff Matrix

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

Model probabilities

$P(nbp_A^B = \mathbf{allLow}, nbp_C^B = \mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.50
(N,N)	(N,S)	(N,N)	0.50
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \mathbf{allLow}, nbp_C^B = \neg\mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\mathbf{allLow}, nbp_C^B = \mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.40
(N,N)	(N,S)	(N,N)	0.40
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

$P(nbp_A^B = \neg\mathbf{allLow}, nbp_C^B = \neg\mathbf{allLow})$

A	B	C	u
(N,N)	(N,N)	(N,N)	0.85
(N,N)	(N,N)	(N,S)	0.70
(N,N)	(N,S)	(N,N)	0.70
	⋮		⋮
(H,H)	(H,H)	(H,E)	0.30
(H,H)	(H,H)	(H,W)	0.41
(H,H)	(H,H)	(H,H)	0.72

2^n models for missing observations

RMM Update Issues

Need to compute: $P(area_A, area_C | move_A, move_C)$ (1)

Can compute: $P(area_A | move_A) \cdot P(area | move_C)$ (2)

- Specifying joint probabilities is difficult in RMM.
- More difficult as the number of agents increases.
  - Each joint probability distribution is larger.
  - Number of distributions is exponential m^n.
- Strong independence assumptions are needed to equate (1) & (2).

Outline

- Introduction
- What is Multiagent Expedition?
- Collaborative Design Network
- Graphical Model for Multiagent Expedition
- Recursive Model for Multiagent Expedition
- **Experimental Results & Discussion**
- Conclusion & Future Work

Agent teams

- CDN
  - Agents communicate over agent interfaces
- RMM (No communication)
  - Agents update belief about state of other agents
- Greedy (No communication)
  - GRDU: agent maximizes unilateral utility
  - GRDB: agent maximizes bilateral + unilateral utility

Experimental Instances

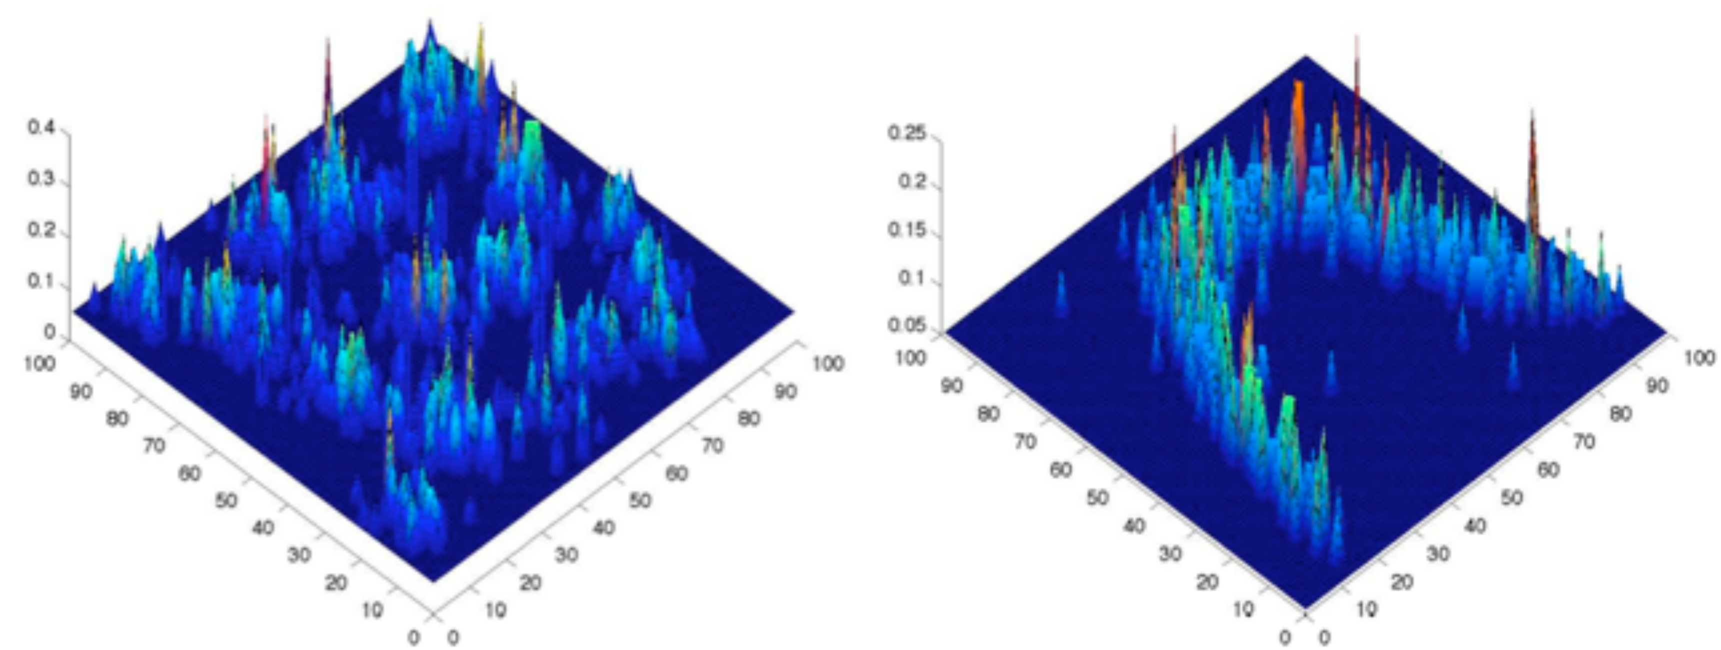

(a) (b)

Experimental Instances

Dense

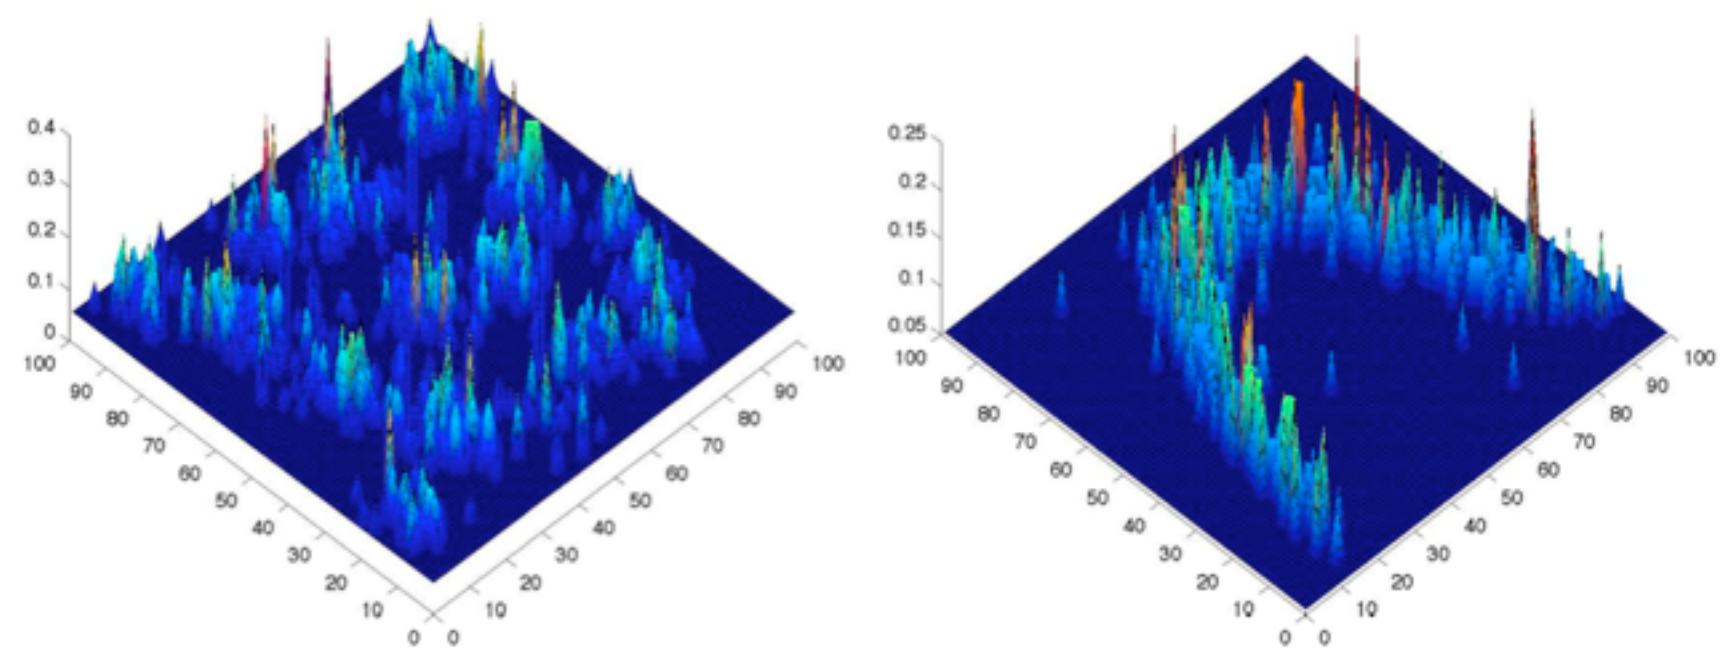

(a) (b)

Experimental Instances

Dense

Path

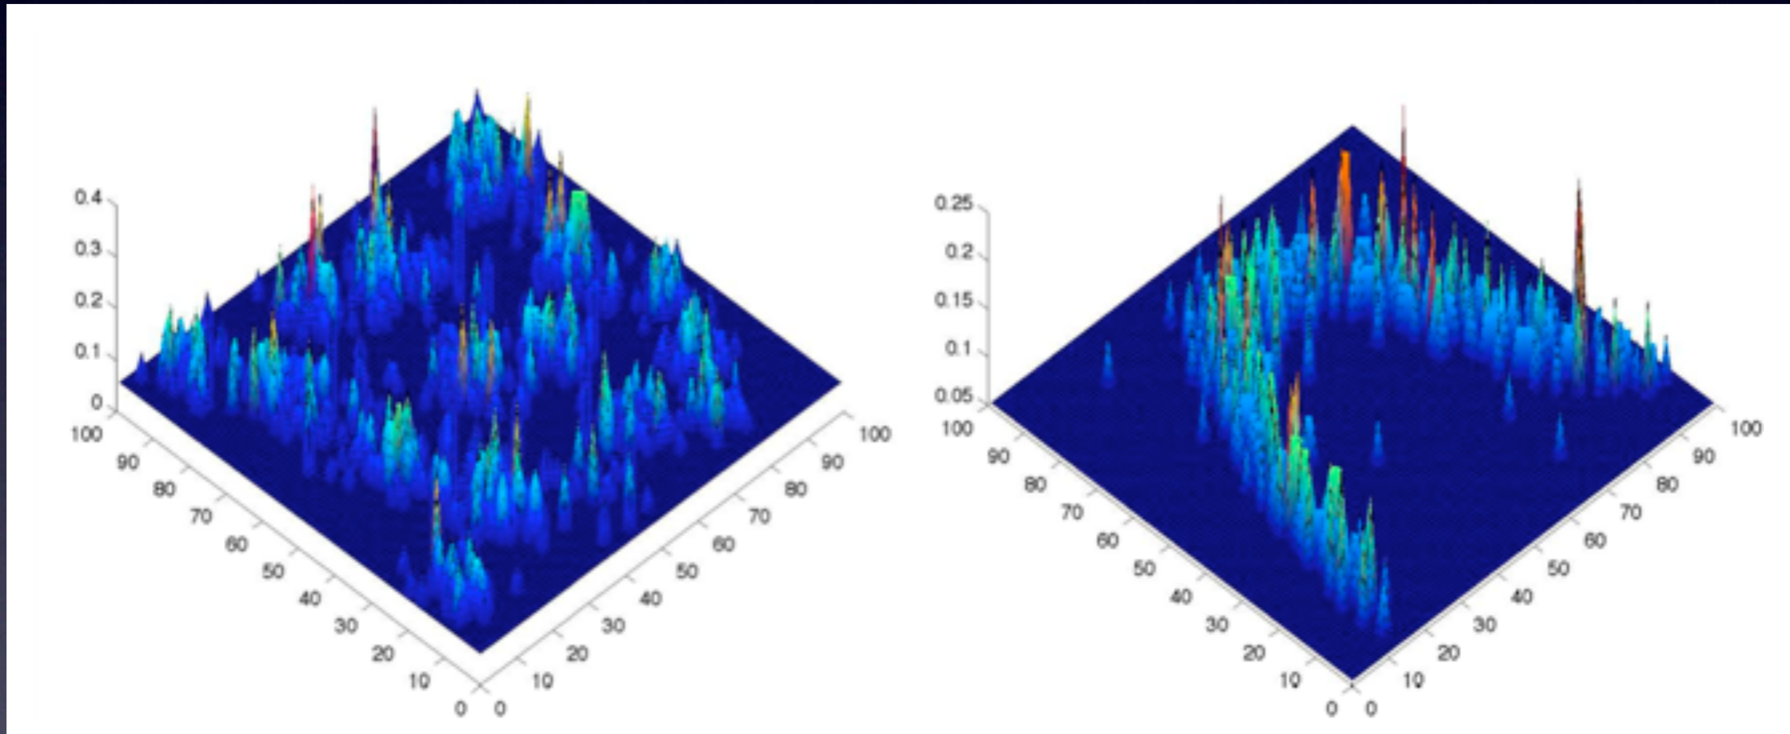

(a) (b)

Experimental Results

- 30 runs for RMM,CDN,GRDU & GRDB
- 40 time-steps per run

Reward collected

Table 1: Experimental results. Highest means bolded.

	Barren		Dense		Path	
	μ	σ	μ	σ	μ	σ
CDN	**55.84**	4.21	**25.14**	3.27	**20.41**	3.39
GRDU	48.56	0.56	12.32	0.20	12.20	0.15
GRDB	48.64	0.62	18.57	1.10	16.80	2.39
RMM	50.35	5.95	18.50	3.39	18.71	2.79

Experimental Results

- 30 runs for RMM,CDN,GRDU & GRDB
- 40 time-steps per run

Reward collected

Table 1: Experimental results. Highest means bolded.

	Barren		Dense		Path	
	μ	σ	μ	σ	μ	σ
CDN	**55.84**	4.21	**25.14**	3.27	**20.41**	3.39
GRDU	48.56	0.56	12.32	0.20	12.20	0.15
GRDB	48.64	0.62	18.57	1.10	16.80	2.39
RMM	50.35	5.95	18.50	3.39	18.71	2.79

Experimental Results

- 30 runs for RMM,CDN,GRDU & GRDB
- 40 time-steps per run

Reward collected

Table 1: Experimental results. Highest means bolded.

	Barren		Dense		Path	
	μ	σ	μ	σ	μ	σ
CDN	**55.84**	4.21	**25.14**	3.27	**20.41**	3.39
GRDU	48.56	0.56	12.32	0.20	12.20	0.15
GRDB	48.64	0.62	18.57	1.10	16.80	2.39
RMM	50.35	5.95	18.50	3.39	18.71	2.79

Experimental Results

- 30 runs for RMM,CDN,GRDU & GRDB
- 40 time-steps per run

Reward collected

Table 1: Experimental results. Highest means bolded.

	Barren		Dense		Path	
	μ	σ	μ	σ	μ	σ
CDN	**55.84**	4.21	**25.14**	3.27	**20.41**	3.39
GRDU	48.56	0.56	12.32	0.20	12.20	0.15
GRDB	48.64	0.62	18.57	1.10	16.80	2.39
RMM	50.35	5.95	18.50	3.39	18.71	2.79

Significance Testing

Table 1: The t-test results.

CDN	GRDU	GRDB	RMMBU
Barren	$\surd$ 99.99	$\surd$ 99.99	$\surd$ 99.99
Dense	$\surd$ 99.99	$\surd$ 99.99	$\surd$ 99.99
Path	$\surd$ 99.99	$\surd$ 99.99	$\surd$ 96.20

Comparison between CDN and each other method.

Performance Discussion

- CDN has higher mean reward collected on all instances than RMM. Why?
  - RMM has no communication.
  - If multiple local optimal plans exist that involve bilateral action:
  - No way for agents to agree which to take.
  - What about adopting a social convention?

Social Convention

- A social convention defines, for each agent what action to take when multiple optimal actions exist.
- Lexicographic ordering as social convention:
- Assume: *u* < *b*

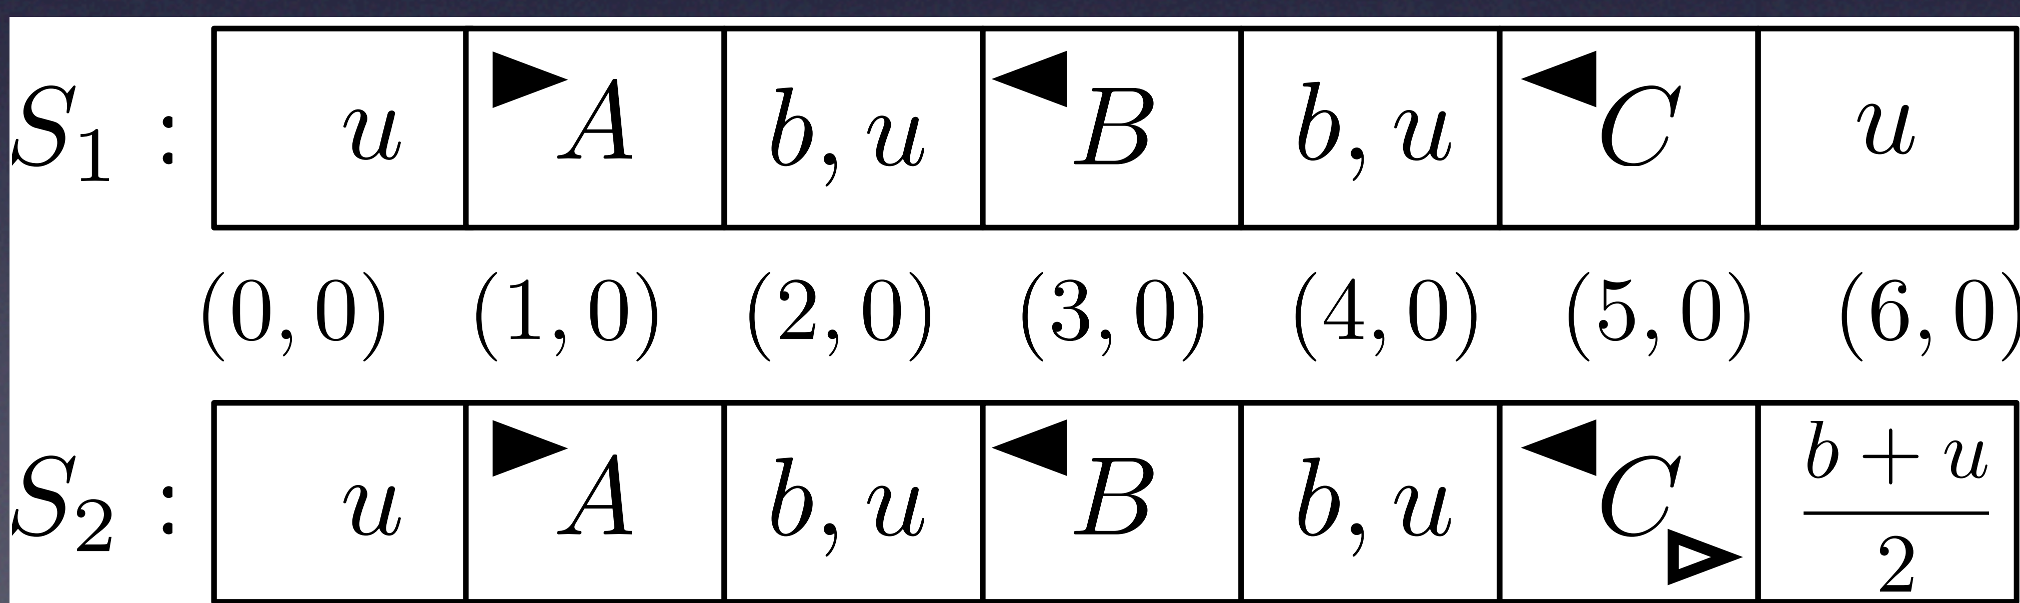

Discussion / Conclusion

- Work motivated by lack of comparative LCF - TCF research
- **Setup level**
  - The agent organization is easier to set-up in LCF.
  - TCF more involved.
- **Modeling level**
  - RMM and LCFs are limited by the need to model agent interactions without sufficient information.
  - In TCF we design agent interfaces such that the agent sub-domains are rendered conditionally independent to take advantage of communication.
  - RMM uses an exponentially complex matrix-based representation. A MAID could be adopted, but the above limitation stands.

Discussion / Conclusion

- **Decision-making level**
  - RMM and LCFs must guess about the states of other agents based on observation.
  - RMM may misjudge states and may misjudge when multiple optimal joint plans exist.
  - Social convention cannot alleviate this difficulty.
  - In TCFs conditional independence rendering interfaces convey sufficient states and decisions and lead to better coordination.

Discussion / Conclusions

- **Generality**
  - Both RMM and CDN are decision-theoretic.
  - The *difference* lies in the agent coupling, and promises that our empirical results can generalize to other domains.

Thanks for listening.

References

- [Pollock90] M. Pollock and M. Ringuette. Introducing the tileworld: experimentally evaluating agent architectures. In T. Dietterich and W. Swartout, editors, Proc. of the 8th National Conf. on Artificial Intelligence, pages 183–189, Menlo Park, CA, 1990. AAAI Press.
- [Noh98] S. Noh and P. J. Gmytrasiewicz. Coordination and belief update in a distributed anti-air environment. In Proceedings of the 31st Hawaii International Conference on System Sciences, volume 5, pages 142–151, Los Alamitos, CA, January 1998. IEEE Computer Society.
- [Bern02] The Complexity of Decentralized Control of Markov Decision Processes. D.S. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. Mathematics of Operations Research, 27(4):819-840, 2002.
- [Blythe99] Decision-Theoretic Planning, Jim Blythe, AI Magazine, Summer 1999.
- [Lee02] A Multi-agent Q-learning Framework for Optimizing Stock Trading Systems, Jae Won Lee, Jangmin O, Database and Expert Systems Applications : 13th International Conference, DEXA 2002 Aix-en-Provence, France, September 2-6, 2002. Proceedings, Lecture Notes in Computer Science, 2002.
- [Bexker04] R. Becker, S. Zilberstein, V. Lesser, and C. V. Goldman. Solving transition independent decentralized markov decision processes, Journal of Artificial Intelligence Research, 22:423-455, 2004.
- [Xiang04] A Decision-Theoretic Graphical Model for Collaborative Design on Supply Chains, Y. Xiang, J. Chen, and A. Deshmukh. Canadian AI 2004, LNAI 3060, pp. 355-369, 2004.
- [Thrun05] S. Thrun, W. Burgard, and D. Fox. Probabilistic Robotics. MIT Press, 2005.
- [Xiang05] Optimal Design in Collaborative Design Network, Y. Xiang, J. Chen and W.S. Havens, Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, July 2005.