

Multi-Agent Systems

Multi-Agent Path Finding

Albert-Ludwigs-Universität Freiburg



Bernhard Nebel, Felix Lindner, and Thorsten Engesser

Winter Term 2018/19

1 Motivation



Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

Nebel, Lindner, Engesser – MAS

3 / 81

Agents moving in a spatial environment




A central problem in many applications is the coordinated movement of agents/robots/vehicles in a given spatial environment.

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature



Logistic robots (KARIS)



Airport ground traffic control (atrics)

Nebel, Lindner, Engesser – MAS

4 / 81

2 Multi-agent path finding (MAPF)



- Definition and example
- MAPF Variations
- MAPF Algorithms
- Computational Complexity of MAPF

Motivation
MAPF
Definition and
example
MAPF Variations
MAPF Algorithms
Computational
Complexity of
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

Nebel, Lindner, Engesser – MAS

6 / 81

Multi-agent path finding

Definition (Multi-agent path finding (MAPF) problem)

Given a set of *agents* A , an undirected, simple *graph* $G = (V, E)$, an *initial state* modelled by an injective function $\alpha_0 : A \rightarrow V$, and a *goal state* modelled by another injective function α_* , can α_0 be *transformed* into α_* by *movements of single agents* without collisions?

- **Existence problem**: Does there exist a successful sequence of movements (= *plan*)?
- **Bounded existence problem**: Does there exist a plan of a given *length* k or less?
- **Plan generation problem**: Generate a plan.
- **Optimal plan generation problem**: Generate a shortest plan.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

Computational Complexity of MAPF

Distributed MAPF


MAPF/DU

Summary & Outlook

Literature

Example

Can we find a (central) plan to move the square robot S to v_3 and the circle robot C to v_2 ?



$G = (V, E)$ with $V = \{v_1, v_2, v_3, v_4\}$ and $E = \{\{v_1, v_2\}, \{v_2, v_3\}, \{v_2, v_4\}\}$
 $A = \{S, C\}$ and $\alpha_0(S) = v_1, \alpha_0(C) = v_3, \alpha_*(S) = v_3, \alpha_*(C) = v_2$

Plan: $(C, v_3, v_2), (C, v_2, v_4), (S, v_1, v_2), (S, v_2, v_3), (C, v_4, v_2)$.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

Computational Complexity of MAPF


Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

A special case: 15-puzzle



Pictures from Wikipedia article on 15-Puzzle

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Lecture plan

- MAPF: variations, algorithms, complexity
- Distributed MAPF (each agent plans on it own): DMAPF
- Distributed MAPF with destination uncertainty: MAPF/DU

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Sequential MAPF



- **Sequential MAPF** (or pebble motion on a graph) allows only one agent to move per time step.
- An agent $a \in A$ can move in one step from $s \in V$ to $t \in V$ transforming α to α' , if
 - $\alpha(a) = s$,
 - $\langle s, t \rangle \in E$,
 - there is no agent b such that $\alpha(b) = t$.
- In this case, α' is determined as follows:
 - $\alpha'(a) = t$,
 - for all agents $b \neq a : \alpha(b) = \alpha'(b)$,
- One usually wants to minimize the number of single movements (= **sum-of-cost** over all agents)

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms
Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Parallel MAPF



- **Parallel MAPF** allows many agents to move in parallel, provided they do not collide.
- Two models:
 - **Parallel**: A chain of agents can move provided the first agent can move on an unoccupied vertex.
 - **Parallel with rotations**: A closed cycle in move synchronously.
- In both cases, one is usually interested in the number of parallel steps (= **make-span**).
- However, also the **sum-of-cost** is sometimes considered.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms
Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Anonymous MAPF



- There is a set of **agents** and a set of **targets** (of the same cardinality as the agent set).
- Each target must be reached by one agent.
- This means one first has to assign a target and then to solve the original MAPF problem.
- Interestingly, the problem as a whole is easier to solve (using flow-based techniques).

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms
Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Types of MAPF algorithms



- **A*-based** algorithm (optimal)
- **Conflict-based search** (optimal)
- **Reduction-based approaches**: Translate MAPF to SAT, ASP or to a CSP (usually optimal)
- **Suboptimal search-based algorithms** (may even be incomplete): **Cooperative A*** (CA*), **Hierarchical Cooperative A*** (HCA*) and **Windowed HCA*** (WHCA*).
- **Rule-based algorithms**: **Kornhauser's algorithm**, **Push-and-Rotate**, **BIBOX**, ... (complete on a given class of graphs, but suboptimal)

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms
A*-based algorithm
Cooperative A*

BIBOX
Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

A*-based algorithm



- Define state space:
 - A **state** is an assignment of agents to vertices (modelled by a function α)
 - There is a **transition** from one state α to α' iff there is a legal move from α to α' according to the appropriate semantics (sequential, parallel, or parallel with rotations)
- Search in this state space using the **A*** algorithm.
- Possible **heuristic estimator**: Sum or maximum over the length of the individual movement plans (ignoring other agents).
- Problem**: Large *branching factor* because of many agents that can move.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF


Distributed MAPF

MAPF/DU

Summary & Outlook


Literature

Example: State space for A* algorithm



Convention: Function α is represented by $\langle \alpha(S), \alpha(C) \rangle$

Question: How many states?



Question: Heuristic value for states $\langle v_1, v_2 \rangle$ and $\langle v_2, v_3 \rangle$ under the sum-aggregation?

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

CA*



- Problems with **A*** on MAPF state space:
 - super-exponential** state space, i.e., $m!/(m-n)!$ with m nodes and n agents;
 - huge** branching factor: $n \times d$ for sequential and d^n for parallel MAPF for graphs with maximal degree d .
- CA***: Decoupled planning in space & time
 - Order agents linearly and then plan for each agent **separately** a (shortest) path.
 - Store each path in a **reservation table**, which stores for each node at which time point it is occupied.
 - When planning, take the **reservation table** into account and **avoid** nodes at time points, when they are reserved for other agents; **wait** action is possible.
 - Solvability** depends on chosen order.
 - Our small example is not solvable with this method!

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF


Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Example CA* run



- Linear order: $\langle C, S \rangle$
- Plan for C : (C, v_3, v_2) , (C, v_2, v_4)
- Reservation table: $(0 : v_1)$, $(0 : v_3)$, $(1 : v_2)$, $(2 - n : v_4)$
- Plan for S : **wait**, (S, v_1, v_2) , (S, v_2, v_3)
- Reservation table: $(0 : v_1)$, $(0 : v_3)$, $(1 : v_2)$, $(2 - n : v_5)$, $(1 : v_1)$, $(2 : v_2)$, $(3 - n : v_3)$
- Not solvable with different order!

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

BIBOX is a **rule-based** algorithm that is complete on all **bi-connected** graphs with at least two unoccupied nodes in the graph.

Definition

A graph $G = (V, E)$ is **connected** iff $|V| \geq 2$ and there is *path* between each pair of nodes $s, t \in V$. A graph is **bi-connected** iff $|V| \geq 3$ and for each $v \in V$, the graph $(V - \{v\}, E')$ with $E' = \{\{x, y\} \in E \mid x, y \neq v\}$ is connected.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF


MAPF/DU

Summary & Outlook

Literature

Loop decomposition

Every bi-connected graph can be constructed from a **cycle** by adding **loops** iteratively.



A **loop decomposition** into a basic cycle and additional loops can be done in time $O(|V|^2)$.

Let us name them C_0, L_1, L_2, \dots , where the index depends on the time when the loop is added.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF


Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Moving unoccupied nodes and agents around



- An unoccupied place can be sent to any node.
- Any agent can be sent to any node by rotating the agents in a cycle or in the loop.
- This can be done without disturbing loops with a higher index than the one the agent starts and finishes in.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF


MAPF/DU

Summary & Outlook

Literature

Filling loops

- Starting with highest-index loop: Move agents to destination loop, then shift agents to their destinations.
- Special case: When agents are already in the destination loop, they have to be rotated out of the loop.



- When done with one loop, repeat for next one with next lower index.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF


MAPF/DU

Summary & Outlook

Literature

Reordering agents in the cycle

- Assumption: The destinations for the empty places are in the cycle C_0 (can be relaxed).
- If the agents are in the right order, just rotate them to their destinations.
- Otherwise reorder by successively take one out and re-insert.



Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Runtime and plan length estimation

- Moving an empty place around is in $O(|V|)$ steps.
 - Moving one agent to an arbitrary position can be done in $O(|V|^2)$ steps.
 - Moving one agent to its final destination in a loop needs $O(|V|^2)$.
 - Since this has to be done $O(|V|)$ times, we need overall $O(|V|^3)$ steps.
 - Reordering in the final cycle is also bounded by $O(|V|^3)$.
- Runtime and number of steps is bounded by $O(|V|^3)$.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

A*-based algorithm

Cooperative A*

BIBOX

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

Computational Complexity of MAPF

- **Existence:** For arbitrary graphs with at least one empty place, the problem is polynomial ($O(|V|^3)$ using Kornhauser's algorithm). For BIBOX on bi-connected with at least two empty places also cubic, but smaller constant.
- **Generation:** $O(|V|^3)$, generating the same number of steps, again using Kornhauser's algorithm or BIBOX (on a smaller instance set).
- **Bounded existence:** Is definitely in NP
 - If there exists a solution, then it is polynomially bounded.
 - A solution candidate can be checked in polynomial time for satisfying the conditions of being a movement plan with k of steps or less.
- **Question:** Is the problem also NP-hard?

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

Computational Complexity of MAPF

Distributed MAPF

MAPF/DU

Summary & Outlook

Literature

The Exact Cover By 3-Sets (X3C) Problem

Definition (Exact Cover By 3-Sets (X3C) Problem)

Given a set of elements U and a collection of subsets $C = \{s_j\}$ with $s_j \subseteq U$ and $|s_j| = 3$. Is there a sub-collection of subsets $C' \subseteq C$ such that $\bigcup_{s \in C'} s = U$ and all subsets in C' are pairwise disjoint, i.e., $s_a \cap s_b = \emptyset$ for each $s_a, s_b \in C'$ with $s_a \neq s_b$?

X3C is **NP-complete**.

Example

$U = \{1, 2, 3, 4, 5, 6\}$

$C = \{\{1, 2, 3\}, \{2, 3, 4\}, \{2, 5, 6\}, \{1, 5, 6\}\}$

$C'_1 = \{\{1, 2, 3\}, \{2, 3, 4\}\}$ is not a cover.

$C'_2 = \{\{1, 2, 3\}, \{2, 3, 4\}, \{1, 5, 6\}\}$ is not an exact cover.

$C'_3 = \{\{2, 3, 4\}, \{1, 5, 6\}\}$ is an **exact cover**.

Motivation

MAPF

Definition and example

MAPF Variations

MAPF Algorithms

Computational Complexity of MAPF

Distributed MAPF

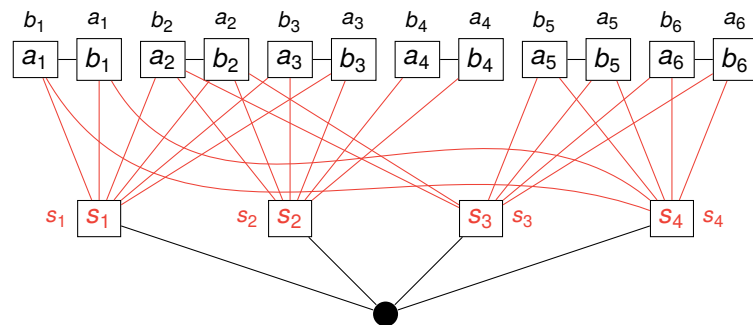
MAPF/DU

Summary & Outlook

Literature

NP-hardness of MAPF: Reduction from X3C

$$C = \{\{1, 2, 3\}, \{2, 3, 4\}, \{2, 5, 6\}, \{1, 5, 6\}\}$$



Claim: There is an exact cover by 3-sets iff the constructed MAPF instance can be solved in at most $k = 11/3|U|$ moves.

3 Distributed MAPF

- Implicit coordination
- Joint execution
- Agent types
- Conservative replanning


Going beyond MAPF

- In MAPF, planning is performed **centrally**, then the plan is communicated to all agents and execution is done **decentrally**.
- What if there is no central instance and communication of plans is impossible?
- In this setting, which we call **DMAPF**, we assume that everybody wants to achieve the common goal of reaching all destinations.
 - ⇒ Each agent needs to **plan decentrally**.
 - ⇒ What **kind of plans** do we need to generate?
 - ⇒ How do we define the **joint execution** of such plans?

Implicitly coordinated plans (in a cooperative setting)

- An agent plans its **own actions** ...
- ... in a way to **empower** the other agents to reach the **common goal**.
- This implies to **plan for the other agents**.
- We consider **one possibility** for the other agent to continue the plan, i.e., the plan will be a **linear plan**.
- We assume that plans are non-redundant, i.e., that they are **cycle-free**.
- Executing such a plan will thus never lead to a **dead end**, i.e., a state from which the other agents **cannot reach the common goal**.
- However, almost certainly, agents will come up with different (perhaps **conflicting**) plans.
- How do we define **joint execution** of such conflicting plans?

Example: Two implicitly coordinated plans



How to solve the problem?

$$\begin{aligned}\pi_C &= \langle (C, v_3, v_2), (C, v_2, v_4), (S, v_1, v_2), (S, v_2, v_3), (C, v_4, v_2) \rangle \\ \pi_S &= \langle (S, v_1, v_2), (S, v_2, v_4), (C, v_3, v_2), (C, v_2, v_1), (S, v_4, v_2), \\ &\quad (S, v_2, v_3), (C, v_1, v_2) \rangle\end{aligned}$$


Motivation
MAPF
Distributed
MAPF
Implicit coordination
Joint execution
Agent types
Conservative
replanning
MAPF/DU
Summary &
Outlook
Literature

Joint execution

- Let us assume, all agents have planned and a subset of them came up with a *family of plans* $(\pi_i)_{i \in A}$.
- Among the agents that have a plan with their *own action as the next action to execute*, one is chosen.
- The action of the chosen agent is *executed*.
- Agents, which have anticipated the action, *track* that in their plans.
- All other agents have to *replan* from the new state.
- Since everybody has a successful plan, no acting agent will ever execute an action that leads to a *dead end*.

Motivation
MAPF
Distributed
MAPF
Implicit coordination
Joint execution
Agent types
Conservative
replanning
MAPF/DU
Summary &
Outlook
Literature

Example execution



Planning, executing, and replanning:

1. $C : \langle (C, v_3, v_2), (C, v_2, v_4), (S, v_1, v_2), (S, v_2, v_3), (C, v_4, v_2) \rangle$
2. $S : \langle (S, v_1, v_2), (S, v_2, v_4), (C, v_3, v_2), (C, v_2, v_1), (S, v_4, v_2), (S, v_2, v_3), (C, v_1, v_2) \rangle$
3. $C : \langle (C, v_2, v_4), (S, v_1, v_2), (S, v_2, v_3), (C, v_4, v_2) \rangle$

Done!

Motivation
MAPF
Distributed
MAPF
Implicit coordination
Joint execution
Agent types
Conservative
replanning
MAPF/DU
Summary &
Outlook
Literature


Lazy and eager agents

What can go wrong?

- Agents could be *lazy*: Sometimes they choose a plan where they expect that another agent should act, although they could act.
- Agents may *wait forever* for each other to act (dish washing dilemma).
- Agents could be *eager*: If agents could act (without creating a cycle or a dead end), they choose to act.
- Agents might create cyclic executions (without creating plans that are cyclic), leading to *infinite executions*.

Motivation
MAPF
Distributed
MAPF
Implicit coordination
Joint execution
Agent types
Conservative
replanning
MAPF/DU
Summary &
Outlook
Literature

Example for infinite execution



Motivation
MAPF
Distributed MAPF
Implicit coordination
Joint execution
Agent types
Conservative replanning
MAPF/DU
Summary & Outlook
Literature

π_1 (S initially): $\langle (S, v_2, v_3), (S, v_3, v_4), (S, v_4, v_5), (C, v_6, v_7), \dots \rangle$
 π_2 (C initially): $\langle (C, v_6, v_5), (C, v_5, v_4), (C, v_4, v_3), (S, v_2, v_1), \dots \rangle$
 π_3 (C after (S, v₂, v₃)): $\langle (C, v_6, v_5), (C, v_5, v_4), (S, v_3, v_2), (C, v_4, v_3), \dots \rangle$
 π_4 (S after (C, v₆, v₅)): $\langle (S, v_3, v_2), (S, v_2, v_1), (S, v_1, v_8), (S, v_8, v_7), \dots \rangle$
 π_5 (C after (S, v₃, v₂)): $\langle (C, v_5, v_6), (C, v_6, v_7), (C, v_7, v_8), (C, v_8, v_1), \dots \rangle$
 π_5 (S after (C, v₅, v₆)): $\langle (S, v_2, v_3), \dots \rangle$

Optimally eager agents

Motivation
MAPF
Distributed MAPF
Implicit coordination
Joint execution
Agent types
Conservative replanning
MAPF/DU
Summary & Outlook
Literature

- Eager agents avoid **deadlocks**, however they are **hyper-active**.
 - They might even **move away** from their destination!
 - So, let force them to be smart: They should generate only **optimal plans** ... and among those optimal plans they should also be **eager**.
 - In our previous example: After the square agent moved right, the circle agent will choose to move left!
- Does it always work out?

Optimally eager agents are always successful

Motivation
MAPF
Distributed MAPF
Implicit coordination
Joint execution
Agent types
Conservative replanning
MAPF/DU
Summary & Outlook
Literature

Theorem

Optimally eager agents are always successful on all solvable DMAPF instances.

Proof.

By induction over the length of a shortest plan k .

$k=0$: Obviously true.

Assume the claim is true for k . Consider a DMAPF instance such that there exists a shortest plan of length $k + 1$. Because the agents are eager, at least one agent wants to move. One agent will move (according to an optimal plan) and by this reduce the necessary number of steps by one. Hence, we have now an instance with plan length k and the induction hypothesis applies. \square

Conservative replanning

Motivation
MAPF
Distributed MAPF
Implicit coordination
Joint execution
Agent types
Conservative replanning
MAPF/DU
Summary & Outlook
Literature

- Optimally eager agents have to solve a sequence of **NP-hard** problems.
 - Is it possible to solve the problem more efficiently?
 - **Conservative replanning**: Always start at the initial state and consider the already executed movements as a prefix of the new plan.
- Avoids infinite executions because plans have to be **cycle-free**.
- ⇒ The agents might visit the **entire state space**



- Assume agents are selected for execution following a pattern similar to a **Gray counter**.

Other ways to coordinate?



- One way to avoid NP-hardness or exponentially longer plans might be to use polynomial-time *approximation algorithms*. However, if different such algorithm are used, also an exponential blowup could result.
- Is it possible to use the *rule-based* algorithms (which are polynomial)?
- Assume that everybody uses the same algorithm: Of course, the agents would act in coordinated way, but this more like *central* planning.
- If the agents may use different algorithms, then it is not clear how to avoid *cyclic executions*.
- Conservative replanning is not helpful in this context, because the executed actions might not be a prefix of a valid plan!

Motivation
MAPF
Distributed MAPF
Implicit coordination
Joint execution
Agent types
Conservative replanning
MAPF/DU
Summary & Outlook
Literature

4 MAPF/DU: MAPF under destination uncertainty



- Implicitly Coordinated Branching Plans
- Strong plans
- Stepping Stones
- Execution cost
- Execution guarantees
- Computational Complexity: Reminder
- Computational Complexity of MAPF/DU

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

MAPF/DU: MAPF under destination uncertainty



MAPF under *destination uncertainty* (MAPF/DU):

- The *common goal* of all agents is that everybody reaches its destination.
- All agents know their own destinations, but these are *not common knowledge* any longer.
- For each agent, there exists a *set of possible destinations*, which are *common knowledge*.
- All agents plan and re-plan *without communicating* with their peers.
- A *success announcement action* becomes necessary, which the agents may use to announce that they have reached their destination (and after that they are not allowed to move anymore).

→ Models multi-robot interactions without communication

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

MAPF/DU: Conceptual problems



- We need a *solution concept* for the agents: *implicitly coordinated branching plans*.
 - We need to find *conditions* that guarantee *success* of joint execution.
 - We have to determine the *computational complexity* for finding plans and deciding solvability.
- Since MAPF/DU is a special case of *epistemic planning* (initial state uncertainty which is monotonically decreasing), we can use concepts and results from this area.


Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

MAPF/DU representation & state space

- In addition to the sets of agents A , the graph $G = (V, E)$, and the assignment of agents to nodes α , we need a function to represent the **possible destinations** $\beta : A \rightarrow 2^V$.
- We assume that the set of possible destinations are **pairwise disjoint** (this can be relaxed, though).
- An **objective state** is given by the pair $s = \langle \alpha, \beta \rangle$ representing the **common knowledge** of all agents.
- A **subjective state** of agent i is given by $s^i \langle \alpha, \beta, i, v \rangle$ with $v \in \beta(i)$, representing the private knowledge of agent i .
- A **MAPF/DU instance** is given by $\langle A, G, s_0, \alpha_* \rangle$, where $s_0 = \langle \alpha_0, \beta_0 \rangle$.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

MAPF/DU: Implicitly coordinated branching plans



- Square agent S **wants to go to** v_3 and **knows** that circle agent C wants to go to v_1 or v_4 .
- C **wants to go to** v_4 and **knows** that S wants to go to v_2 or v_3 .
- Let us assume S forms a plan in which it moves in order to **empower** C to reach their common goal.
- S needs **shifting its perspective** in order to plan for all possible destinations of C (**branching on destinations**).
- Planning for C , S must **forget** about its own destination.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Branching plans: Building blocks

Branching plans consist of:

- Movement actions**: $(\langle agent \rangle, \langle sourcenode \rangle, \langle targetnode \rangle)$, i.e., a movement of an agent
- Success announcement**: $(\langle agent \rangle, S)$, after that all agents know that the agent has reached its destination and it cannot move anymore
- Perspective shift**: $[\langle agent \rangle : \dots]$, i.e., from here on we assume to plan with the knowledge of agent $\langle agent \rangle$. This can be unconditional or conditional on $\langle agent \rangle$'s destinations.
- Branch on all destinations**: $(? \langle dest_1 \rangle \{ \dots \}, \dots, ? \langle dest_n \rangle \{ \dots \})$, where all destinations of the current agent have to be listed. For each case we try to find a successful plan to reach the goal state.

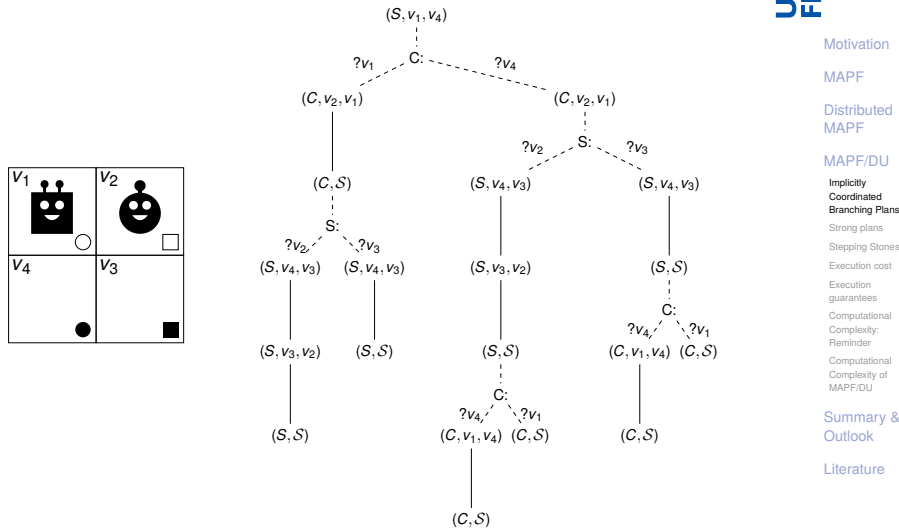
Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Semantics of branching plans

- Movement actions** modify α in the obvious way.
- A **success announcement** of agent i transforms β to β' such that $\beta'(i) = \emptyset$ in order to signal that i cannot move anymore.
- A **perspective shift from i to j with subsequent branching on destinations** transforms the subjective state $s^i = \langle \alpha, \beta, i, v_i \rangle$ to a set of subjective states $s^k = \langle \alpha, \beta, j, v_{jk} \rangle$ with all $v_{jk} \in \beta(j)$.
- A **perspective shift from i to j without subsequent branching on destinations** induces the same transformation, but enforces that the subsequent plans are the same for all states subjective states s^k .

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Branching plan: Example



Nebel, Lindner, Engesser – MAS

49 / 81

Strong plans

Similar to the notion of **strong plans** in non-deterministic single-agent planning, we define ***i*-strong plans** for an agent *i* to be:


- **cycle-free**, i.e., not visiting the same objective state twice;
- **always successful**, i.e. always ending up in a state such that all agents have announced success;
- **covering**, i.e., for all combinations of possible destinations of agents different from *i*, success can be reached.

Nebel, Lindner, Engesser – MAS

50 / 81

Subjectively and objectively strong plans

- A plan is called **subjectively strong** if it is *i*-strong for some agent *i*.
- A plan is called **objectively strong** if it is *i*-strong for each agent *i*.
- An instance is **objectively** or **subjectively solvable** if there exists an objectively or subjectively strong plan, respectively.




- There does not exist a *T*-strong plan, but an *S*- and a *C*-strong plan.
- Difference between **subjective** and **objective** solvability concerns only the first acting agent!

Nebel, Lindner, Engesser – MAS

51 / 81

Structure of strong plans: Stepping stones

- A **stepping stone** for agent *i* is a state in which *i* can move to each of its possible destinations, announcing success, and afterwards, for each possible destination, there exists an *i*-strong plan to solve the resulting states.
- *S* can create a **stepping stone for C** by moving from *v*₁ via *v*₄ to *v*₃.
- *C* can now move to *v*₁ or *v*₄ and announce success.
- In each case, *S* can move afterwards to its destination (or stay) and announce success.



Nebel, Lindner, Engesser – MAS

52 / 81

Stepping Stone Theorem




Theorem

Given an *i*-solvable MAPF/DU instance, there exists an *i*-strong branching plan such that the only branching points are those utilizing stepping stones.

Proof sketch.

Remove non-stepping stone branching points by picking one branch without success announcement. \square



Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Execution cost



The **execution cost** of a branching plan is the number of atomic actions of the longest execution trace.

Theorem

Given an *i*-solvable MAPF/DU instance over a graph $G = (V, E)$, then there exists an *i*-strong branching plan with execution cost bounded by $O(|V|^4)$.

Proof sketch.

Direct consequence of the stepping stone theorem and the maximal number of movements in the MAPF problem. \square

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature


Joint execution and execution guarantees



- Joint **execution** is defined similarly to the fully observable case: One agent is chosen; afterwards the plan is tracked or the agent has to replan.
- In the MAPF/DU framework not all agents might have a plan initially!
- One might hope that **optimally eager agents** are **always successful**.
- In **epistemic planning** this was proven to be true only in the **uniform knowledge** case.
- We do not have uniform knowledge ... and indeed, **execution cycles** cannot be excluded.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

A counter example



A number on an edge means that there are as many nodes on a line.

- Agent 2 has a shortest eager plan moving first to v_6 .
- Agent 1 has then a shortest eager plan moving first to v_4 .
- Agent 2 has then a shortest eager plan moving first to v_5 .
- Agent 1 has then a shortest eager plan moving first to v_2 .

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

- Perhaps **conservatism** can help!
- Similarly to DMAPF, conservative replanning means that the already executed actions are used as a **prefix** in the plan to be generated.
- Differently from DMAPF, we assume that after a **success announcement**, the initial state is modified so that the **real destination** of the agent is known in the initial state.
- Otherwise we could not solve instances that are only **subjectively solvable**.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

- **Conservative, eager agents** are **always successful**, but might **visit the entire state space** before terminating.
- Adding **optimal eagerness** can help to reduce the execution length.


Theorem

For solvable MAPF/DU instances, joint execution and replanning by conservative, optimally eager agents is always successful and the execution length is polynomial.

Proof idea.

After the second agent starts to act, all agents have an **identical perspective** and for this reason produce objectively strong plans with the same execution costs, which can be shown to be bounded polynomially using the stepping stone theorem. □

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature



- Assume *S* moves first to v_4 .
- Assume *C* re-plans. From now on, in replanning from the beginning, it has to do a perspective shift to *S*, because it now has to extend the partial plan starting with (S, v_4, v_1) , i.e., it has to create an **objectively strong plan**.
- Assume that *C* moves now to v_1 .
- From now on, also *S* has to make a perspective shift to *C*, effectively “forgetting” its own destination, i.e., it also has to create a **objectively strong plan**.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

- We use **Turing machines** as formal models of algorithms
- This is justified, because:
 - we assume that Turing machines can compute all computable functions
 - the resource requirements (in term of time and memory) of a Turing machine are only polynomially worse than other models
- The regular type of Turing machine is the **deterministic** one: **DTM** (or simply **TM**)
- Often, however, we use the notion of **nondeterministic** TMs: **NDTM**

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Complexity classes P and NP
NP-completeness
The class co-NP
The class PSPACE
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Computational Complexity: Problems, solutions, and complexity



UNI
FREIBURG

- A **problem** is a set of pairs (I, A) of strings in $\{0, 1\}^*$.
 I : instance; A : answer
 If all answers $A \in \{0, 1\}$: **decision problem**
- A **decision problem** is the same as a **formal language**:
 the set of strings formed by the instances with answer 1
- An algorithm **solves** (or **decides**) a problem if it computes
 the right answer for all instances.
- **Complexity of an algorithm**: function

$$T: \mathbb{N} \rightarrow \mathbb{N},$$

measuring the **number of basic steps** (or memory
 requirement) the algorithm needs to compute an answer
 depending on the **size** of the instance

- **Complexity of a problem**: complexity of the most efficient
 algorithm that solves this problem.

Computational Complexity: Complexity classes P and NP



UNI
FREIBURG

Problems are categorized into **complexity classes** according to
 the requirements of computational resources:

- The class of problems decidable on **deterministic Turing
 machines** in **polynomial time**: **P**
 - Problems in P are assumed to be **efficiently solvable**
 (although this might not be true if the exponent is very large)
 - In practice, a reasonable definition
- The class of problems decidable on **non-deterministic
 Turing machines** in **polynomial time**, i.e., having a poly.
 length accepting computation for all positive instances: **NP**
- More classes are definable using other resource bounds on
 time and memory

Computational Complexity: Upper and lower bounds



UNI
FREIBURG

- **Upper bounds** (**membership** in a class) are usually easy to
 prove:
 - provide an **algorithm**
 - show that the resource bounds are respected
- **Lower bounds** (**hardness** for a class) are usually difficult to
 show:
 - the technical tool here is the **polynomial reduction** (or any
 other appropriate reduction)
 - show that some hard problem can be reduced to the
 problem at hand

Computational Complexity: Polynomial reduction



UNI
FREIBURG

- Given languages L_1 and L_2 , L_1 can be **polynomially
 reduced to** L_2 , written $L_1 \leq_p L_2$, if there exists a polynomial
 time-computable function f such that

$$x \in L_1 \iff f(x) \in L_2.$$


Rationale: it cannot be harder to decide L_1 than L_2

- L is **hard** for a class C (**C-hard**) if all languages of this class
 can be reduced to L .
- L is **complete** for C (**C-complete**) if L is C -hard and $L \in C$.

Computational Complexity: NP-complete problems



- A problem is **NP-complete** iff it is **NP-hard** and **in NP**.
- Example: **SAT** (the satisfiability problem for propositional logic) is NP-complete (Cook/Karp)
 - Membership is obvious, hardness follows because computations on a NDTM correspond to satisfying truth assignments of certain formulae



Nebel, Lindner, Engesser – MAS

65 / 81

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Implicitly
Coordinated
Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution
guarantees
Computational
Complexity:
Reminder
Complexity
classes P and NP
NP-completeness
The class co-NP
The class
PSPACE
Computational
Complexity of
MAPF/DU
Summary &
Outlook

Literature

Computational Complexity: The complexity class co-NP



- Note that there is some **asymmetry** in the definition of NP:
 - It is clear that we can decide **SAT** by using a **NDTM** with polynomially bounded computation
 - There exists an accepting computation of polynomial length iff the formula is satisfiable
 - In other words: Checking a proposed solution (of poly size) is easy.
 - What if we want to decide **UNSAT**, the complementary problem?
 - It seems necessary to check **all** possible truth-assignments!
- Define **co-C** = $\{L \subseteq \Sigma^* : \Sigma^* \setminus L \in C\}$ (provided Σ is our alphabet)
- **co-NP** = $\{L \subseteq \Sigma^* : \Sigma^* \setminus L \in NP\}$
- Examples: **UNSAT**, **TAUT** \in co-NP!
- **Note:** P is closed under complement, in particular,

$$P \subseteq NP \cap \text{co-NP}$$

Nebel, Lindner, Engesser – MAS

66 / 81

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Implicitly
Coordinated
Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution
guarantees
Computational
Complexity:
Reminder
Complexity
classes P and NP
NP-completeness
The class co-NP
The class
PSPACE
Computational
Complexity of
MAPF/DU
Summary &
Outlook

Literature

Computational Complexity: PSPACE



There are problems even more difficult than NP and co-NP...

Definition ((N)PSPACE)

PSPACE (**NPSPACE**) is the class of decision problems that can be decided on deterministic (non-deterministic) Turing machines using only **polynomially many tape cells**.

Some facts about PSPACE:

- PSPACE is **closed under complements** (... as all other deterministic classes)
- PSPACE is **identical** to NPSPACE (because non-deterministic Turing machines can be simulated on deterministic TMs using only quadratic space: Savitch's Theorem)
- $NP \subseteq PSPACE$ (because in polynomial time one can "visit" only polynomial space, i.e., $NP \subseteq NPSPACE$)

Nebel, Lindner, Engesser – MAS

67 / 81

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Implicitly
Coordinated
Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution
guarantees
Computational
Complexity:
Reminder
Complexity
classes P and NP
NP-completeness
The class co-NP
The class
PSPACE
Computational
Complexity of
MAPF/DU
Summary &
Outlook

Literature

Computational Complexity: PSPACE-completeness



Definition (PSPACE-completeness)

A decision problem (or language) is **PSPACE-complete** if it is in PSPACE and all other problems in PSPACE can be polynomially reduced to it.

Intuitively, **PSPACE-complete** problems are the "hardest" problems in PSPACE (similar to NP-completeness). They appear to be "harder" than **NP-complete** problems from a **practical point of view**.

An example for a PSPACE-complete problem is the **NDFA equivalence problem**:

Instance: Two non-deterministic finite state automata A_1 and A_2 .

Question: Are the languages accepted by A_1 and A_2 identical?

Nebel, Lindner, Engesser – MAS

68 / 81

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Implicitly
Coordinated
Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution
guarantees
Computational
Complexity:
Reminder
Complexity
classes P and NP
NP-completeness
The class co-NP
The class
PSPACE
Computational
Complexity of
MAPF/DU
Summary &
Outlook

Literature

this is true.

Computational complexity of MAPF/DU bounded plan existence




Theorem

Deciding whether there exists an eager MAPF/DU i -strong or objectively strong plan with execution cost k or less is **PSPACE-complete**.

Proof sketch.

Since plans have polynomial depth, all execution traces can be generated non-deterministically and tested using only polynomial space, i.e., PSPACE-membership. For hardness, reduction from QBF. Example construction for

$$\forall x_1 \exists x_2 \forall x_3 : (x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee x_3)$$




Nebel, Lindner, Engesser – MAS

69 / 81

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

The reduction enlarged



Nebel, Lindner, Engesser – MAS

70 / 81

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Complexity with a fixed number of agents



These results probably imply that the technique could not be used online.

For a fixed number of agents, however, the bounded planning problem is polynomial.

Theorem

For a fixed number c of agents, deciding whether there exists a MAPF/DU i -strong or objectively strong plan with execution cost of k or less can be done in time $O(n^{c^2+c})$.

That means, for two agents, it takes “only” $O(n^6)$ time – but in practice it should be faster.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Nebel, Lindner, Engesser – MAS

71 / 81

An algorithm for generating an objective MAPF/DU plan for two agents



- 1 Determine in the **state space of all node assignments** the distance to the initial state using **Dijkstra**: $O(|V|^4)$ time.
- 2 For each of the $O(|V|^2)$ configurations check, whether it is a **potential stepping stone** for one agent, i.e., whether all potential destinations of this agent are reachable using **Dijkstra** on the modified graph, where the other agent blocks the way: $O(|V|^4)$ time.
- 3 For all $O(|V|^2)$ potential stepping stones, check whether for each of the $O(|V|)$ possible destination of the first agent, the second agent can reach its possible destinations and use **Dijkstra** to compute the shortest path: altogether $O(|V|^5)$ time.
- 4 Consider all stepping stones and minimize over the maximum plan depth. Among the minimal plans select those that are **eager** for the planning agent.

Motivation
MAPF
Distributed MAPF
MAPF/DU
Implicitly Coordinated Branching Plans
Strong plans
Stepping Stones
Execution cost
Execution guarantees
Computational Complexity: Reminder
Computational Complexity of MAPF/DU
Summary & Outlook
Literature

Nebel, Lindner, Engesser – MAS

72 / 81

5 Summary & Outlook



Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

Summary



Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

- DMAPF generalizes the MAPF problem by dropping the assumption that plans are **generated centrally** and then communicated.
- MAPF/DU generalizes the MAPF problem further by dropping the assumptions that destinations are **common knowledge**.
- A solution concept for this setting are ***i*-strong branching plans** corresponding to implicitly coordinated policies in the area of epistemic planning.
- The backbone of such plans are **stepping stones**.
- Joint execution can be guaranteed to be **successful and polynomially bounded** if all agents are **conservative** and **optimally eager**.
- While **plan existence** in general is PSPACE-complete, it is polynomial for a fixed number of agents.

Outlook



Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature



- Do the results still hold for **planar graphs**?
- Is MAPF/DU plan existence also PSPACE-complete?
- How would more general forms of describing the common knowledge about destinations affect the results?
- **Overlap** of destinations or **general Boolean combinations**
- Can we get similar results for other execution semantics?
- **Concurrent** executions of actions
- Can we be more aggressive in expectations about possible destinations?
- Use **forward induction**, i.e., assume that actions in the past were rational.
- Are other forms of implicit coordination possible?
- More **communication**? Coordination in **competitive scenarios**?

6 Literature



Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

Literature (1)

-  **D. Kornhauser, G. L. Miller, and P. G. Spirakis.**
Coordinating pebble motion on graphs, the diameter of permutation groups, and applications.
In 25th Annual Symposium on Foundations of Computer Science (FOCS-84), pages 241–250, 1984.
-  **O. Goldreich.**
Finding the shortest move-sequence in the graph-generalized 15-puzzle is NP-hard.
In Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation, pages 1–5. 2011.




Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

Literature (2)

-  **H. Ma, S. Koenig, N. Ayanian, L. Cohen, W. HÃunig, T. K. Satish Kumar, T. Uras, H. Xu, C. A. Tovey, G. Sharon:**
Overview: Generalizations of Multi-Agent Path Finding to Real-World Scenarios.
[CoRR abs/1702.05515, 2017.](#)
-  **A. Felner, R. Stern, S. E. Shimony, E. Boyarski, M. Goldenberg, G. Sharon, N. R. Sturtevant, G. Wagner, and P. Surynek.**
Search-Based Optimal Solvers for the Multi-Agent Pathfinding Problem: Summary and Challenges.
In Proceedings of the Tenth International Symposium on Combinatorial Search (SOCS-17), pages 29–37, 2017.
-  **P Surynek.**
A novel approach to path planning for multiple robots in bi-connected graphs.
In Proc. 2009 IEEE International Conference on Robotics and Automation, ICRA 2009, pages 3613–3619, 2009.

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature

Literature (3)

-  **B. Nebel, T. Bolander, T. Engesser and R. Mattmüller.**
Implicitly Coordinated Multi-Agent Path Finding under Destination Uncertainty.
Accepted to be published in Journal of Artificial Intelligence research.
-  **T. Bolander, T. Engesser, R. Mattmüller and B. Nebel.**
Better Eager Than Lazy? How Agent Types Impact the Successfulness of Implicit Coordination.
In Proceedings of the Sixteenth Conference on Principles of Knowledge Representation and Reasoning (KR-18), pages 445-453. 2018.
-  **T. Engesser, T. Bolander, R. Mattmüller, and B. Nebel.**
Cooperative epistemic multi-agent planning for implicit coordination.
In Proceedings of the Ninth Workshop on Methods for Modalities (M4MICLA-17), pages 75–90, 2017.

Motivation
MAPF
Distributed
MAPF
MAPF/DU
Summary &
Outlook
Literature