

# Solving Concurrent Multiagent Planning using Classical Planning

Daniel Furelos-Blanco and Anders Jonsson

Department of Information and Communication Technologies

Universitat Pompeu Fabra

Roc Boronat 138, 08018 Barcelona, Spain

{daniel.furelos, anders.jonsson}@upf.edu

## Abstract

In this work we present a novel approach to solving concurrent multiagent planning problems in which several agents act in parallel. Our approach relies on a compilation from concurrent multiagent planning to classical planning, allowing us to use an off-the-shelf classical planner to solve the original multiagent problem. The solution can be directly interpreted as a concurrent plan that satisfies a given set of concurrency constraints, while avoiding the exponential blowup associated with concurrent actions. Theoretically, we show that the compilation is sound and complete. Empirically, we show that our compilation can solve challenging multiagent planning problems that require concurrent actions.

## Introduction

Research in multiagent planning has seen a lot of progress in recent years, in part due to the first competition of distributed and multiagent planners, CoDMAP-15 (Komenda, Stolba, and Kovacs 2016). Many recent multiagent planners are based on the MA-STRIPS formalism (Brafman and Domshlak 2008), and can be loosely classified into one of two categories: centralized, in which agents have full information and share the goal, and distributed, in which agents have partial information and individual goals. In CoDMAP-15, the most successful centralized planners were ADP (Crosby, Rovatsos, and Petrick 2013), MAP-LAPKT (Muise, Lipovetzky, and Ramirez 2015) and CMAP (Borrajo 2013), while prominent distributed planners included PSM (Tozicka, Jakubuv, and Komenda 2014), MAPlan (Stolba, Fiser, and Komenda 2016) and MH-FMAP (Torreño, Onaindia, and Sapena 2014).

Although establishing a common set of multiagent benchmark domains was a major step forward, the domains from the centralized track of CoDMAP-15 can all be solved using *sequential* plans in which one agent acts at a time. In contrast, there are many applications that require agents to act in *parallel*; examples include shared journey planning, robot coordination tasks such as RoboCupSoccer (Nardi et al. 2014) and RoboCupRescue (Sheh, Schwertfeger, and Visser 2016), and real-time strategy games, to name a few.

In this paper we consider the problem of *concurrent* centralized multi-agent planning in which agents can act in par-

allel at each time step. This problem is challenging for different reasons: the number of concurrent actions is worst-case exponential in the number of agents, and restrictions are needed to ensure that concurrent actions are well-formed. Usually, these restrictions take the form of *concurrency constraints* (Boutilier and Brafman 2001; Crosby 2013), which can model both the case when two actions *must* occur in parallel, and the case when two actions *cannot* occur in parallel.

Although some planners from CoDMAP-15 (CMAP, MAPlan and MH-FMAP) can produce concurrent plans, there are few that can reliably handle more complex concurrency constraints. A notable exception is the work of Crosby, Jonsson, and Rovatsos (2014), who associate concurrency constraints with the *objects* of a multiagent planning problem and transform the problem into a sequential, single-agent problem that can be solved using a classical planner.

Brafman and Zoran (2014) extended the distributed forward-search planner MAFS (Nissim and Brafman 2014) to support concurrency constraints while preserving privacy. In MAFS, each agent maintains its own search space, and has a queue for expanded states (closed list) and another for states to be expanded (open list). When a state  $s$  is expanded, the agent uses its own operators only; thus, two agents expanding the same state will generate different successors. Messages are exchanged between agents in order to inform each other about the expansion of relevant states. Consequently, agents explore the search space together while preserving privacy. Maliah, Brafman, and Shani (2017) proposed MAFBS, which extended MAFS to use forward and backward messages. This new approach reduced the number of messages required and also resulted in an increase in the privacy of agents.

In this paper we describe a planner that can handle arbitrary concurrency constraints, as originally proposed by Boutilier and Brafman (2001). Our approach is similar to that of Crosby, Jonsson, and Rovatsos (2014) in that we transform a multiagent planning problem into a single-agent problem that is significantly easier to solve, while avoiding the exponential blowup associated with concurrent actions. However, the concurrency constraints of Boutilier and Brafman are significantly more expressive than those of Crosby, Jonsson, and Rovatsos, enabling us to solve multiagent problems with more complex interactions (e.g. effects that depend on the concurrent actions of other agents). We show

that our planner is sound and complete, and perform experiments in several concurrent multiagent planning domains to evaluate its performance.

The remainder of this paper is structured as follows. We first introduce the different planning formalisms that are needed to describe our planner. Next, we describe the compilation from multiagent planning to single-agent planning that our planner employs. We then present the results of experiments with our planner in several domains that require agents to act in parallel. Finally, we relate our planner to existing work in the literature, and conclude with a discussion.

## Background

In this section we describe the planning formalisms that we use in the paper: classical planning, concurrent planning, and concurrent multiagent planning.

### Classical Planning

We consider the fragment of classical planning with conditional effects and negative conditions and goals. Given a set of fluents  $F$ , a *literal*  $l$  is a valuation of a fluent in  $F$ , where  $l = f$  denotes that  $l$  assigns true to  $f \in F$ , and  $l = \neg f$  that  $l$  assigns false to  $f$ . A set of literals  $L$  is *well-defined* if it does not assign conflicting values to any fluent, i.e. if  $L$  does not contain both  $f$  and  $\neg f$  for some fluent  $f \in F$ . Let  $\mathcal{L}(F)$  be the set of well-defined literal sets on  $F$ , i.e. the set of all partial assignments of values to fluents. Given a literal set  $L \in \mathcal{L}(F)$ , let  $\neg L = \{\neg l : l \in L\}$  be the *complement* of  $L$ . We also define the *projection*  $L|_X$  of a literal set  $L$  onto a subset of fluents  $X \subseteq F$ .

A *state*  $s \in \mathcal{L}(F)$  is a well-defined literal set such that  $|s| = |F|$ , i.e. a total assignment of values to fluents. Explicitly including negative literals  $\neg f$  in states simplifies subsequent definitions, but we often abuse notation by defining a state  $s$  only in terms of the fluents that are true in  $s$ , as is common in classical planning.

A classical planning problem is a tuple  $\Pi = \langle F, A, I, G \rangle$ , where  $F$  is a set of fluents,  $A$  is a set of actions,  $I \in \mathcal{L}(F)$  is an initial state, and  $G \in \mathcal{L}(F)$  is a goal condition (usually satisfied by multiple states). Each action  $a \in A$  has a precondition  $\text{pre}(a) \in \mathcal{L}(F)$  and a set of conditional effects  $\text{cond}(a)$ . Each conditional effect  $C \triangleright E \in \text{cond}(a)$  is composed of two literal sets  $C \in \mathcal{L}(F)$  (the condition) and  $E \in \mathcal{L}(F)$  (the effect).

An action  $a \in A$  is applicable in state  $s$  if and only if  $\text{pre}(a) \subseteq s$ , and the resulting (triggered) *effect* is given by

$$\text{eff}(s, a) = \bigcup_{C \triangleright E \in \text{cond}(a), C \subseteq s} E,$$

i.e. effects whose conditions hold in  $s$ . We assume that  $\text{eff}(s, a)$  is a well-defined literal set in  $\mathcal{L}(F)$  for each state-action pair  $(s, a)$ . The result of applying  $a$  in  $s$  is a new state  $\theta(s, a) = (s \setminus \neg \text{eff}(s, a)) \cup \text{eff}(s, a)$ . It is straightforward to show that if  $s$  and  $\text{eff}(s, a)$  are in  $\mathcal{L}(F)$ , then so is  $\theta(s, a)$ .

Given a planning problem  $\Pi$ , a *plan* is an action sequence  $\pi = \langle a_1, \dots, a_n \rangle$  that induces a state sequence  $\langle s_0, s_1, \dots, s_n \rangle$  such that  $s_0 = I$  and, for each  $i$  such that  $1 \leq i \leq n$ , action  $a_i$  is applicable in  $s_{i-1}$  and generates

the successor state  $s_i = \theta(s_{i-1}, a_i)$ . The plan  $\pi$  *solves*  $\Pi$  if and only if  $G \subseteq s_n$ , i.e. if the goal condition is satisfied following the application of  $\pi$  in  $I$ .

### Concurrent Planning

*Concurrent planning* is the extension of classical planning that allows actions to be applied in parallel, forming *concurrent* or *joint* actions. Given a classical planning problem  $\Pi = \langle F, A, I, G \rangle$ , an unconstrained concurrent planning problem is given by  $\Pi_{\text{conc}} = \langle F, 2^A, I, G \rangle$ , where the action set  $2^A$  is given by the power set of the original action set  $A$ . The aim is to find a *concurrent plan*  $\pi_{\text{conc}}$  that solves  $\Pi_{\text{conc}}$  by applying a sequence of concurrent actions in  $2^A$ .

To ensure that joint actions have well-defined effects, researchers often impose *concurrency constraints* that model whether two atomic actions *must* or *cannot* be done concurrently. For example, in PDDL 2.1 (Fox and Long 2003), two actions  $a^1$  and  $a^2$  cannot be applied concurrently if  $a^1$  has an effect on a fluent  $f$  and  $a^2$  has a precondition or effect on  $f$ . Crosby (2013) defines concurrency constraints on actions that have the same *object* in their preconditions or effects.

We adopt a formulation of concurrency constraints due to Boutilier and Brafman (2001), later extended by Kovacs (2012). The idea is to extend preconditions and conditional effects with *actions* in addition to fluents. We overload notation and let each  $a \in A$  denote a propositional variable. If  $a^1$  has a precondition  $a^2$ , then  $a^1$  is only applicable if  $a^2$  is concurrent with  $a^1$ . If  $a^1$  has a precondition  $\neg a^2$ , then  $a^1$  is only applicable if  $a^2$  is *not* concurrent with  $a^1$ .

We extend the notation for classical planning as follows. By viewing  $A$  as a set of propositional variables, we can model a subset of atomic actions in  $2^A$  as a well-defined literal set  $b \in \mathcal{L}(A)$  such that  $|b| = |A|$ , analogous to how states are formed from fluents. The subset includes those actions that appear as positive literals in  $b$ , while negative literals denote those actions that are *not* part of the subset.

For a given atomic action  $a \in A$ , the precondition  $\text{pre}(a) \in \mathcal{L}(F \cup A)$  and any condition  $C \in \mathcal{L}(F \cup A)$  of a conditional effect  $C \triangleright E \in \text{cond}(a)$  are extended to include actions in addition to fluents. Each effect  $E \in \mathcal{L}(F)$  of a conditional effect  $C \triangleright E \in \text{cond}(a)$  is exclusively on fluents as before. Given a state  $s$  and a set of actions that are concurrent with  $a$ , represented by a literal set  $b \in \mathcal{L}(A)$ ,  $a$  is applicable if and only if  $\text{pre}(a) \subseteq s \cup b$ , and the resulting effect is given by

$$\text{eff}(s \cup b, a) = \bigcup_{C \triangleright E \in \text{cond}(a), C \subseteq s \cup b} E,$$

i.e. effects whose conditions hold in  $s \cup b$ .

We can now define the semantics of a joint action  $b \in 2^A$ . Concretely,  $b$  satisfies the concurrency constraints if and only if  $\text{pre}(a)|_A \subseteq b \setminus \{a\}$  for each  $a \in b$ . If  $b$  is applicable, its precondition and effect are defined as the union of the preconditions and effects of the constituent atomic actions:

$$\text{pre}(b) = \bigcup_{a \in b} \text{pre}(a)|_F,$$

$$\text{eff}(s, b) = \bigcup_{a \in b} \text{eff}(s \cup b \setminus \{a\}, a), \quad \forall s.$$

As before, we assume that  $\text{eff}(s, b)$  is a well-defined literal set in  $\mathcal{L}(F)$  for each pair of state  $s$  and applicable joint action  $b$ .

### Concurrent Multiagent Planning

In concurrent multiagent planning, each atomic action belongs to an *agent*. We consider the problem of *centralized* multiagent planning in which agents share the goal.

A *concurrent multiagent planning problem* (CMAP) is a tuple  $\Pi = \langle N, F, \{A^i\}_{i=1}^n, I, G \rangle$ , where  $N = \{1, \dots, n\}$  is a set of agents and  $A^i$  is the set of atomic actions of agent  $i \in N$ . This is identical to the standard definition of multiagent planning problems (Brafman and Domshlak 2008), with the only difference being that the actions include concurrency constraints as defined in the previous section. The fluent set  $F$ , initial state  $I$  and goal condition  $G$  are defined as before.

A CMAP implicitly defines a negative concurrency constraint on each pair of atomic actions  $(a^1, a^2) \subseteq A^i$  that belong to the same agent  $i$ . Consequently, each agent can contribute at most one atomic action to each joint action. These concurrency constraints are not included in action definitions.

To illustrate CMAPs, we use the TABLEMOVER domain (Boutilier and Brafman 2001), in which two agents move blocks between rooms. There are two possible strategies:

1. Pick up blocks and carry them using their arms.
2. Put blocks on a table, carry the table together to another room, and tip the table to make the blocks fall down.

Figure 1 shows the definition of the lift-side action in the notation of Kovacs (2012), which is used by agent ?a to lift side ?s of the table. The precondition is that the side must be down (i.e. on the floor) and the agent cannot be holding anything. Moreover, the precondition also states that no other agent ?a2 can lower side ?s2 at the same time. When the action is applied, ?s is no longer down but up, and ?a is busy lifting ?s. The action also has a conditional effect (represented by the when clause): if some side ?s2 is not lifted by any agent ?a2, then all blocks on the table fall to the floor.

Note that the action lift-side is defined using forall quantifiers. In practice, such quantifiers are compiled away, such that the resulting actions have quantifier-free preconditions and effects, as in our definition of actions.

### Compilations for CMAPs

Let  $\Pi = \langle N, F, \{A^i\}_{i=1}^n, I, G \rangle$  be a CMAP, and let  $A = A^1 \cup \dots \cup A^n$  be the set of atomic actions of  $\Pi$ . A straightforward approach to solving  $\Pi$  is to define a concurrent planning problem  $\Pi_{conc} = \langle F, B, I, G \rangle$  and apply a classical planner to solve  $\Pi_{conc}$ . If  $B \subseteq 2^A$  is exactly the subset of joint actions that satisfy the concurrency constraints of the actions in  $\Pi$ , this approach is both sound and complete.

However, even though the joint action set  $B$  might be much smaller than  $2^A$ , it is still worst-case exponential in the number of agents. Most classical planners ground the actions, and if  $B$  is too large, they often do not make it past pre-processing. Moreover, to generate the set  $B$  we would typ-

```
(:action lift-side
:agent ?a - agent
:parameters (?s - side)
:precondition
  (and (at-side ?a ?s)
        (down ?s) (handempty ?a)
        (forall (?a2 - agent ?s2 - side)
          (not (lower-side ?a2 ?s2))))
:effect
  (and (not (down ?s)) (lifting ?a ?s)
        (up ?s) (not (handempty ?a ?s))
        (forall
          (?b - block ?r - room ?s2 - side)
          (when
            (and (inroom Table ?r)
                  (on-table ?b) (down ?s2))
              (forall (?a2 - agent)
                (not (lift-side ?a2 ?s2))))))
        (and (on-floor ?b) (inroom ?b ?r)
              (not (on-table ?b))))))
```

Figure 1: Definition of the TABLEMOVER action lift-side using the notation of Kovacs (2012) (concurrency constraints in bold).

ically have to iterate over *all* joint actions, and test whether each action satisfies the concurrency constraints of atomic actions.

Here we describe an alternative approach to solving a CMAP  $\Pi$ . The idea is to model each joint action  $b = \{a^1, \dots, a^k\}$  using multiple atomic actions: one set of actions for *selecting*  $a^1, \dots, a^k$ , one set of actions for *applying*  $a^1, \dots, a^k$ , and one set of actions for *resetting*  $a^1, \dots, a^k$ . The result is a classical planning problem  $\Pi' = \langle F', A', I', G' \rangle$  such that the size of the action set  $A'$  is *linear* in  $|A|$ , the number of atomic actions of agents.

Simulating a joint action  $b$  using a sequence of atomic actions  $\langle a^1, \dots, a^k \rangle$  is problematic for the following reason: when applying an atomic action  $a^i$ , we may not yet know which atomic actions will be applied by other agents. Since those other actions may be part of the precondition and conditional effects of  $a^i$ , it becomes difficult to ensure that the concurrency constraints of  $a^i$  are correctly enforced.

Our approach is to divide the simulation of a joint action  $b$  into three phases: selection, application, and reset. In the selection phase, we use an auxiliary fluent *active- $a^i$*  to model that the atomic action  $a^i$  has been selected. In the application phase, since the selection of atomic actions is known, we can substitute each action  $a^i$  in preconditions and conditional effects with the auxiliary fluent *active- $a^i$* . In the reset phase, various auxiliary fluents are reset.

Note that this compilation takes into account the multiagent nature of the problem. Each agent can apply at most one atomic action per time step, and agents collaborate to form joint actions whose constituent atomic actions are compatible and/or inapplicable on their own.

We proceed to define the components of the compilation.

### Fluents

We describe the fluents in PDDL format, i.e. each fluent is instantiated by assigning objects to predicates.

The set of fluents  $F' \supseteq F$  includes all original fluents in  $F$ , plus the following auxiliary fluents:

- Fluents free, select, apply and reset modeling the phase.
- For each agent  $i$ , fluents free-agent( $i$ ), busy-agent( $i$ ) and done-agent( $i$ ) that model the agent state: free to select an action, selected an action, and applied the action.
- For each action  $a^i \in A^i$  in the action set of agent  $i$ , a fluent active- $a^i$  which models that  $a^i$  has been selected. We use  $F_{act}$  to denote the subset of fluents of this type.

By simple inspection, the total number of fluents in  $F'$  is given by  $|F'| = |F| + 4 + 3n + \sum_{i \in N} |A^i| = O(|F| + |A|)$ .

The initial state  $I'$  of the compilation  $\Pi'$  is given by

$$I' = I \cup \{\text{free}\} \cup \{\text{free-agent}(i) : i \in N\},$$

i.e. the initial state on fluents in  $F$  is  $I$ , we are not simulating any joint action, and all agents are free to select actions. The goal condition is given by  $G' = G \cup \{\text{free}\}$ , i.e. the goal condition  $G$  has to hold at the end of a joint action simulation.

### Actions

For a literal set  $L \in \mathcal{L}(F \cup A)$ , let  $L|_A/F_{act}$  denote the projection of  $L$  onto  $A$ , followed by a substitution of the actions in  $A$  with the corresponding fluents in  $F_{act}$ . Note that both  $L|_F$  and  $L|_A/F_{act}$  are literal sets on fluents in  $F'$ , i.e. the dependence on actions in  $A$  is removed.

The first four actions in the set  $A'$  allow us to switch between simulation phases, and are defined as follows:

- select-phase:   pre = {free},  
                   cond =  $\{\emptyset \triangleright \{\neg \text{free}, \text{select}\}\}$ .
- apply-phase:   pre = {select},  
                   cond =  $\{\emptyset \triangleright \{\neg \text{select}, \text{apply}\}\}$ .
- reset-phase:   pre = {apply},  
                   cond =  $\{\emptyset \triangleright \{\neg \text{apply}, \text{reset}\}\}$ .
- finish:        pre = {reset, free-agent( $i$ ) :  $i \in N$ },  
                   cond =  $\{\emptyset \triangleright \{\neg \text{reset}, \text{free}\}\}$ .

For each action  $a^i \in A^i$  in the action set of agent  $i$ , we define three new actions in  $A'$ : select- $a^i$ , do- $a^i$  and end- $a^i$ . These actions represent the three steps that an agent must perform during the simulation of a joint action.

The action select- $a^i$  causes  $i$  to select action  $a^i$  during the selection phase, and is defined as follows:

$$\text{pre} = \{\text{select}, \text{free-agent}(i)\} \cup \text{pre}(a^i)|_F,$$

$$\text{cond} = \{\emptyset \triangleright \{\text{busy-agent}(i), \neg \text{free-agent}(i), \text{active-}a^i\}\}.$$

The precondition ensures that we are in the selection phase, that  $i$  is free to select an action, and that the precondition of  $a^i$  holds on fluents in  $F$ . The effect prevents  $i$  from selecting another action, and marks  $a^i$  as selected.

The action do- $a^i$  applies the effect of  $a^i$  in the application phase, and is defined as follows:

$$\text{pre} = \{\text{apply}, \text{busy-agent}(i), \text{active-}a^i\} \cup \text{pre}(a^i)|_A/F_{act},$$

$$\text{cond} = \{\emptyset \triangleright \{\text{done-agent}(i), \neg \text{busy-agent}(i)\}\} \\ \cup \{C|_F \cup C|_A/F_{act} \triangleright E : C \triangleright E \in \text{cond}(a^i)\}.$$

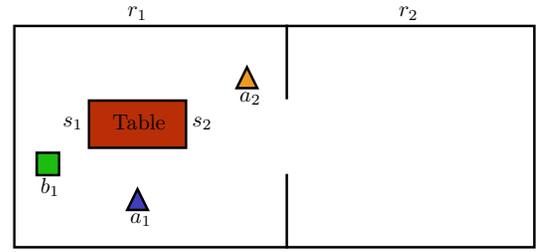


Figure 2: Initial state of a simple TABLEMOVE instance.

The precondition ensures that we are in the application phase, that  $a^i$  was previously selected, and that all concurrency constraints in the precondition of  $a^i$  hold. The effect is to apply all conditional effects of  $a^i$ , where each condition  $C|_F \cup C|_A/F_{act}$  is generated from  $C$  by substituting each action  $b^j \in A$  with active- $b^j$ . Agent  $i$  is also marked as done to prevent  $a^i$  from being applied a second time.

The action end- $a^i$  resets auxiliary fluents to their original value, and is defined as follows:

$$\text{pre} = \{\text{reset}, \text{done-agent}(i), \text{active-}a^i\},$$

$$\text{cond} = \{\emptyset \triangleright \{\text{free-agent}(i), \neg \text{done-agent}(i), \neg \text{active-}a^i\}\}.$$

The precondition ensures that we are in the reset phase and that  $a^i$  was previously selected and applied (due to done-agent( $i$ )). The effect is to make agent  $i$  free to select actions again, and to mark  $a^i$  as no longer selected.

Again, by inspection we can see that the total number of actions in  $A'$  is given by  $|A'| = 4 + 3 \sum_i |A^i| = O(|A|)$ .

### Properties

Figure 2 shows an example instance of TABLEMOVE in which the goal is for agents  $a_1$  and  $a_2$  to move block  $b_1$  from room  $r_1$  to room  $r_2$ . An example concurrent plan that solves this instance is defined as follows:

```

1 (to-table a1 r1 s2) (pickup-floor a2 b1 r1)
2 (putdown-table a2 b1 r1)
3 (to-table a2 r1 s1)
4 (lift-side a1 s2) (lift-side a2 s1)
5 (move-table a1 r1 r2 s2) (move-table a2 r1 r2 s1)
6 (lower-side a1 s2)

```

In this plan, agent  $a_2$  first puts the block on the table, and then  $a_1$  and  $a_2$  concurrently lift each side of the table and move the table to room  $r_2$ . Finally,  $a_1$  lowers its side of the table, causing the table to tip and the block to fall to the floor.

The following sequence of classical actions in  $A'$  can be used to simulate the first joint action of the concurrent plan:

```

1 (select-phase )
2 (select-to-table a1 r1 s2)
3 (select-pickup-floor a2 b1 r1)
4 (apply-phase )
5 (do-pickup-floor a2 b1 r1)
6 (do-to-table a1 r1 s2)
7 (reset-phase )
8 (end-to-table a1 r1 s2)
9 (end-pickup-floor a2 b1 r1)
10 (finish )

```

We show that the compilation is both sound and complete.

**Theorem 1 (Soundness).** *A classical plan  $\pi'$  that solves  $\Pi'$  can be transformed into a concurrent plan  $\pi$  that solves  $\Pi$ .*

*Proof.* When fluent *free* is true, the only applicable action is *select-phase*. The only way to make *free* true again is to cycle through the three phases and end with the *finish* action.

During the selection phase, a subset of actions  $a^1, \dots, a^k$  are selected, causing the corresponding agents to be busy. Because of the precondition  $\text{free-agent}(i)$  of the *finish* action, each selected action  $a^i$  has to be applied in the application phase, and reset in the reset phase. The resulting simulated joint action is given by  $b = \{a^1, \dots, a^k\}$ .

The precondition of  $b$  holds since the precondition of each  $a^i$  on fluents in  $F$  is checked in the selection phase, during which no fluents in  $F$  change values. The concurrency constraints of  $a^i$  are checked in the application phase when all actions have already been selected. This also ensures that the conditional effects of  $a^i$  are correctly applied. Finally, auxiliary fluents are cleaned in the reset phase. Hence the joint action  $b$  satisfies all concurrency constraints and is correctly simulated by the corresponding action subsequence of  $\pi'$ .

Let  $\pi$  be the concurrent plan composed of the sequence of joint actions simulated by the plan  $\pi'$ . Since  $\pi'$  solves  $\Pi'$ , the goal condition  $G$  holds at the end of  $\pi'$ , implying that  $G$  also holds at the end of  $\pi$ . This implies that  $\pi$  solves  $\Pi$ .  $\square$

**Theorem 2 (Completeness).** *A concurrent plan  $\pi$  that solves  $\Pi$  corresponds to a classical plan  $\pi'$  that solves  $\Pi'$ .*

*Proof.* Let  $b = \{a^1, \dots, a^k\}$  be a joint action of the concurrent plan  $\pi$ . We can use a sequence of actions in  $A'$  to simulate  $b$  by selecting, applying and resetting each action among  $a^1, \dots, a^k$ . Since  $b$  is part of  $\pi$ , its precondition and concurrency constraints have to hold, implying that the precondition and concurrency constraints of each atomic action hold. Hence the action sequence is applicable and results in the same effect as  $b$ . By concatenating such action sequences for each joint action of  $\pi$ , we obtain a plan  $\pi'$ . Since  $\pi$  solves  $\Pi$ , the goal condition  $G$  holds at the end of  $\pi$ , implying that  $G$  holds at the end of  $\pi'$ . This implies that  $\pi'$  solves  $\Pi'$ .  $\square$

## Extensions

The basic compilation checks concurrency constraints in the application phase. Here we describe an extension that checks negative concurrency constraints in the selection phase, allowing a classical planner to identify inadmissible joint actions as early as possible, reducing the branching factor.

Assume that action  $a^i$  has a negative concurrency constraint  $\neg a^j$ . As before, we can simulate this constraint using the fluent  $\neg \text{active-}a^j$ . However,  $a^j$  may be selected *after*  $a^i$  in the selection phase, in which case  $\neg \text{active-}a^j$  holds when selecting  $a^i$ . To prevent inadmissible joint actions from being selected, we introduce additional fluents in the set  $F'$ :

- For each action  $a^i \in A^i$  in the action set of agent  $i$ , a fluent  $\text{req-neg-}a^i$  which indicates that  $a^i$  cannot be selected.

We now redefine the action  $\text{select-}a^i$  as follows:

$$\begin{aligned} \text{pre} &= \{\text{select}, \text{free-agent}(i), \neg \text{req-neg-}a^i\} \cup \text{pre}(a^i)|_F \\ &\cup \{\neg \text{active-}b^j : \neg b^j \in \text{pre}(a^i)\}, \\ \text{cond} &= \{\emptyset \triangleright \{\text{busy-agent}(i), \neg \text{free-agent}(i), \text{active-}a^i\}\} \\ &\cup \{\emptyset \triangleright \{\text{req-neg-}b^j : \neg b^j \in \text{pre}(a^i)\}\}. \end{aligned}$$

To select  $a^i$ ,  $\text{req-neg-}a^i$  has to be false. For each negative concurrency constraint  $\neg b^j$  of  $a^i$ , action  $\text{select-}a^i$  adds fluent  $\text{req-neg-}b^j$ , preventing  $b^j$  from being selected after  $a^i$ .

With this extension, we only need to check *positive* concurrency constraints (i.e. required concurrency) in the application phase. We also redefine  $\text{end-}a^i$  such that fluents of type  $\text{req-neg-}a^i$  are reset in the cleanup phase, using the opposite effect of  $\text{select-}a^i$ . The initial state and goal condition do not change since the new fluents are always false while no joint action is simulated.

The second extension is to impose a bound  $C$  on the number of atomic actions selected in the selection phase, resulting in a classical planning problem  $\Pi'_C = \langle F'_C, A'_C, I'_C, G'_C \rangle$ . The fluent set  $F'_C \supseteq F'$  extends  $F'$  with fluents  $\text{count}(j)$ ,  $0 \leq j \leq C$ . Counter parameters are added to the *select* and *reset* actions so that they can respectively increment and decrement the value of the counter. Crucially, no *select* action is applicable when  $j = C$ , preventing us from selecting more than  $C$  actions. The benefit is to reduce the branching factor by disallowing joint actions with more than  $C$  atomic actions.

We leave the following proposition without proof:

**Proposition 3.** *The compilation  $\Pi'_C$  that includes both proposed extensions is sound.*

Note that the compilation  $\Pi'_C$  is not complete. For instance, consider a concurrent multiagent plan that contains a joint action involving 4 atomic actions. If  $C < 4$ , then the concurrent multiagent plan cannot be converted into an equivalent classical plan without exceeding the bound  $C$ .

## Experimental Results

We tested our compilations in two different sets of domains: centralized multiagent domains from the CoDMAP-15 competition, and four domains (MAZE, TABLEMOVER, WORKSHOP and BOXPUSHING) that require concurrency<sup>1</sup>.

In each domain, we used three variants of our compilations: unbounded joint action size, and joint action size bounded by  $C = 2$  and  $C = 4$ . In all variants, we used the extension that identifies negative concurrency constraints in the selection phase. The resulting classical planning problems were then solved using Fast Downward (Helmert 2006) in the LAMA setting (Richter and Westphal 2010).

All experiments ran on Intel Xeon E5-2673 v4 @ 2.3GHz processors. They had a time limit of 30 minutes and a memory limit of 8 GB.

<sup>1</sup>The code of the compilation and the domains are available at <https://github.com/aig-upf/universal-pddl-parser-multiagent>.

## CoDMAP Domains

Although the centralized multiagent benchmark domains of CoDMAP involve multiple agents, none of these domains require concurrency between agents, i.e. their instances can be solved by sequences of atomic actions. Instead, the purpose of CoDMAP was to solve problems involving privacy over predicates, constants and objects.

As our approach and CoDMAP planners differ in their purpose, we decided to compare the former against the classical planner it uses to solve the compilation (Fast Downward). Since CoDMAP domains can be sequentially solved, Fast Downward can be directly applied on them. Furthermore, we will be able to see how the results of solving a problem sequentially compare to the results if the same problem is solved when concurrency is allowed. In addition, the coverage of Fast Downward is almost as high as the winner of the centralized track at CoDMAP-15, ADP (Crosby, Rovatsos, and Petrick 2013), which solved 222 instances.

We compare the results of our compilation with those obtained by running Fast Downward (FD) directly on the given instances. We wanted to test the following hypotheses:

1. Using our compilations it is possible to solve approximately the same number of instances as a classical planner that ignores the multiagent nature of the problem.
2. The plans resulting from solving the compiled problems are implicitly more compressed since atomic actions are grouped into joint actions.

The domains forming this set of experiments were manually modified to add the appropriate negative concurrency constraints. Otherwise, the instances would violate our assumption that the triggered effect  $\text{eff}(s, b)$  of a joint action is well-defined, producing invalid plans. For instance, an atomic action adding a fluent  $f$  and another action deleting  $f$  could have been part of the same joint action. Therefore, note that the complexity of the tasks increases with respect to the original tasks that had no concurrency constraints.

Table 1 shows the results for this set of experiments. There are four different metrics:

- Coverage: number of instances solved.
- Time: average number of seconds taken to find a solution.
- Plan length: number of actions in the plan, corresponding to the number of joint actions for our compilations.
- #Actions: total number of actions instantiated by FD.

From the table, we observe that the planner with the largest coverage is FD (219, 91.25%). The compilation with the largest coverage is the variant whose joint action size is bounded by 2 (158, 65.83%), followed by the unbounded variant (143, 59.58%) and the variant whose bound is set to 4 (140, 58.3%). Thus, although all variants solve less instances, they are not far from the results obtained by FD for many domains. As expected, when concurrency is not required to solve a MAP, the auxiliary fluents and additional copies of actions (corresponding to the three phases of a joint action simulation) introduced by our compilations do not pay off in solution efficiency, resulting in a smaller coverage and a larger time to solve the compiled instances.

Regarding plan length, the compilations always result in shorter solutions than FD. Since FD outputs a sequential plan with no joint actions, it has no way of compressing the plans, unlike the compilations in which atomic actions are grouped into joint actions.

As stated in the previous section, the number of actions of the compilation with unbounded joint action size is approximately  $3|A|$ , and the number of actions of the compilation with joint action size bounded by  $C$  is approximately  $(C + 1)|A|$ , which is reflected quite well in the number of instantiated actions (results vary somewhat because of the reachability tests performed by FD during grounding).

## Domains with Required Concurrency

In this section, we first give a brief description of the domains that were included in the experiments, and the other algorithms that were used for comparison.

The MAZE domain (Crosby 2014) consists of a grid of interconnected locations. Each agent in the maze must move from an initial location to a target location. The connection between two adjacent locations can be one of the following:

- Door: can only be traversed by one agent at a time. Some are initially locked, and are unlocked by pushing a specific switch, placed anywhere in the maze.
- Bridge: can be crossed by multiple agents at once, but is destroyed after the first crossing.
- Boat: can only be used by two or more agents in the same direction.

The WORKSHOP domain is a new domain in which the objective is to perform inventory in a high-security storage facility. It has the following characteristics:

- To open a door, one agent has to press a switch while another agent simultaneously turns a key.
- To do inventory on a pallet, one agent has to use a forklift to lift the pallet while another agent examines it (for security reasons, labels are located underneath pallets).
- There are also actions for picking up a key, entering or exiting a forklift, moving an agent, and driving a forklift.

The BOXPUSHING domain (Brafman and Zoran 2014) consists in a grid of interconnected locations. Agents must push the boxes from one location to another. A different number of agents is needed depending on the box size: one for small boxes, two for mediums, and three for large boxes.

The algorithm that we use for comparison is that of Crosby, Jonsson, and Rovatsos (2014) (which we refer to as CJR), who define concurrency constraints in the form of affordances on subsets of objects. For example, the affordance on the subset of objects  $\{\text{location, boat}\}$  in the MAZE domain is  $[2, \infty]$ , representing that at least two agents have to row the boat between the same two locations at once. Even though their algorithm cannot handle the concurrency constraints of CMAPs, the MAZE and WORKSHOP domains can be reformulated using their concurrency constraints.

The concurrency constraints of CJR are not as expressive as those of Kovacs because:

Domain	$N$	Coverage				Time (s.)				Plan length				# Actions			
		2	4	$\infty$	FD	2	4	$\infty$	FD	2	4	$\infty$	FD	2	4	$\infty$	FD
BLOCKSWORLD	20	7	2	4	<b>20</b>	759.5	-	-	<b>0.2</b>	<b>32.1</b>	-	-	32.8	6848	12323	4110	<b>1270</b>
DEPOT	20	13	10	9	<b>17</b>	202.9	246.4	223.9	<b>58.3</b>	30.6	15.7	<b>14.9</b>	44.0	10100	18176	6061	<b>2007</b>
DRIVERLOG	20	18	17	18	<b>20</b>	67.3	58.8	73.7	<b>26.1</b>	21.1	<b>20.5</b>	25.2	35.6	38416	69145	23051	<b>7386</b>
ELEVATORS08	20	9	8	10	<b>20</b>	13.8	12.5	9.5	<b>0.2</b>	31.0	<b>30.3</b>	36.4	65.1	10779	19399	6469	<b>2155</b>
LOGISTICS00	20	<b>20</b>	<b>20</b>	<b>20</b>	<b>20</b>	1.9	2.9	212.1	<b>0.0</b>	30.3	<b>28.0</b>	30.1	50.2	1781	3202	1070	<b>318</b>
ROVERS	20	<b>20</b>	<b>20</b>	19	<b>20</b>	45.2	75.2	20.9	<b>0.1</b>	46.5	47.4	<b>42.9</b>	56.8	18314	32962	10990	<b>2609</b>
SATELLITES	20	19	17	19	<b>20</b>	82.8	128.0	32.2	<b>1.0</b>	<b>32.6</b>	35.5	34.2	55.9	45106	81188	27065	<b>8122</b>
SOKOBAN	20	0	0	0	<b>18</b>	-	-	-	<b>32.3</b>	-	-	-	<b>54.1</b>	3319	5970	1993	<b>663</b>
TAXI	20	<b>20</b>	<b>20</b>	<b>20</b>	<b>20</b>	1.3	2.6	0.7	<b>0.0</b>	14.8	<b>14.7</b>	<b>14.7</b>	18.7	544	975	328	<b>108</b>
WIRELESS	20	2	2	2	<b>4</b>	-	-	-	-	-	-	-	-	15644	28156	9388	<b>3128</b>
WOODWORKING08	20	14	8	4	<b>20</b>	290.0	256.0	-	<b>0.9</b>	22.4	<b>11.4</b>	-	46.1	17406	31327	10445	<b>3447</b>
ZENOTRAVEL	20	16	16	18	<b>20</b>	87.2	125.6	164.8	<b>1.5</b>	<b>23.6</b>	24.1	34.2	46.9	67586	121652	40553	<b>13502</b>
Total	240	158	140	143	<b>219</b>												

Table 1: Summary of results for CoDMAP domains. Planners “2” and “4” are compilations having joint action size bounded by 2 and 4 respectively, while “ $\infty$ ” is the variant with unbounded joint action size. “FD” is Fast Downward directly applied to the given instances (no compilation involved).  $N$  is number of instances; time and length are averages over all instances solved, for all planners that solved at least 5 instances. The total number of actions is an average over all instances.

- Actions cannot be used in conditional effects, so their algorithm cannot solve instances of TABLEMOVER.
- To represent concurrency constraints on multiple action templates in PDDL, they have to be defined on the same subset of objects. In contrast, the constraints of Kovacs can be defined on arbitrary pairs of actions.

Furthermore, CJR does not separate the atomic action selection from the atomic action application. This is a big problem since one of the atomic actions can delete the precondition of other atomic actions, thus canceling the formation of the joint action. For example, in the MAZE domain, the action for crossing a bridge requires that the bridge exists, and destroys the bridge as an effect. Therefore, as this approach does not separate the selection from the application, this action can be done just by one agent at a time (and not by infinite agents as the problem states). The same occurs in the BOXPUSHING domain. Instances where a medium or a large box must be moved cannot be solved with this approach because the first agent to “push” the box will move it. Thus, the box location precondition for the other agent(s) does not hold, so the box is not moved in the end.

As for our compilations, we used Fast Downward in the LAMA 2011 setting to solve the instances produced by CJR.

Another algorithm we could have used for comparison is MAFS (Brafman and Zoran 2014). Unlike our approach and Crosby, Jonsson, and Rovatsos’s, it is a distributed approach that preserves privacy and that has been shown to work in the BOXPUSHING domain. However, we have not tested it as we have not had access to the code to perform experiments.

Table 2 shows the results for the four domains. To provide an idea of how each planner behaves as a function of the number of agents, the table shows for each domain the same metrics for different numbers of agents. In the case of BOXPUSHING with three agents, as the version with the bound set to 2 cannot solve instances with large boxes, there are separate results for the cases where the number such boxes is 0, and the cases where it is greater than 0.

In terms of coverage, the unbounded compilation ( $\infty$ ) per-

forms the best (93, 69.92%). The variant bounded to 2 and the variant bounded to 4 have similar coverage: 82 (61.65%) and 89 (66.92%) respectively. The former is clearly affected by the fact it cannot solve the instances requiring the concurrent action of 3 agents in the BOXPUSHING domain. Finally, CJR is the approach with the worst coverage (17, 12.78%). It performs reasonably well in MAZE in spite of its limitation regarding the action for crossing a bridge. On the other hand, its results are not very good in WORKSHOP, and it cannot solve instances from the TABLEMOVER and BOXPUSHING domains. Furthermore, the higher the number of agents, the worse the coverage becomes because problems are harder to solve (the number of available actions grows).

Regarding execution time, the unbounded compilation and the compilation bounded to 2 are the fastest. The higher the number of agents, the longer it takes to compute a plan.

In the case of plan length (i.e. number of joint actions), we observe that the plans obtained with our approach are shorter than the ones obtained with CJR. CJR obtains worse results because it explicitly builds joint actions only if their constituent joint actions are associated with a concurrency constraint. Thus, any action that appears out of a joint action, can be considered as a joint action of size 1. In contrast, our approach allows to combine atomic actions arbitrarily, so it already compresses the solution while planning.

In summary, it is not clear which variant is the best since they highly depend on the problem. In general, the version bounded to 2 works well, but it is useless if concurrency is required between 3 or more agents. The version bounded to 4 has a similar coverage to the unbounded one, but mainly because there are not instances requiring concurrency between more than 4 agents. Thus, we believe that the unbounded option is the most convenient since its performance is not far from the best in all domains, and it is more general.

Finally, we have performed preliminary scalability results in the BOXPUSHING domain. In these experiments we checked the results when  $n \in \{1, \dots, 10\}$  agents are required to move a box from a room  $r_1$  to a neighbor room

Domain	$N$	Coverage				Time (s.)				Plan length				# Actions			
		2	4	$\infty$	CJR	2	4	$\infty$	CJR	2	4	$\infty$	CJR	2	4	$\infty$	CJR
MAZE	20	<b>13</b>	8	6	11	351.9	435.2	<b>144.4</b>	192.8	47.2	22.0	<b>11.7</b>	77.3	41723	69368	<b>27900</b>	156886
$a = 10$	10	<b>8</b>	6	5	7	243.6	564.8	<b>169.1</b>	225.5	48.3	25.0	<b>12.2</b>	79.6	39909	67417	<b>26155</b>	119374
$a = 15$	10	<b>5</b>	2	1	4	<b>525.2</b>	-	-	-	<b>45.4</b>	-	-	-	43989	71807	<b>30080</b>	194397
TABLEMOVER	24	<b>15</b>	12	<b>15</b>	-	<b>263.3</b>	336.5	341.0	-	<b>58.7</b>	59.0	61.5	-	7487	13127	<b>4667</b>	-
$a = 2$	12	10	10	<b>11</b>	-	<b>103.8</b>	226.4	214.6	-	63.5	<b>62.0</b>	64.5	-	3450	6154	<b>2098</b>	-
$a = 4$	12	<b>5</b>	2	4	-	<b>558.2</b>	-	-	-	<b>49.0</b>	-	-	-	11524	20100	<b>7236</b>	-
WORKSHOP	20	<b>15</b>	13	13	6	132.3	298.6	<b>51.8</b>	629.0	35.7	37.0	<b>32.5</b>	63.5	18002	31000	11502	<b>5425</b>
$a = 4$	10	<b>8</b>	<b>8</b>	<b>8</b>	5	42.1	261.6	<b>36.6</b>	587.3	<b>37.3</b>	43.9	<b>37.3</b>	65.8	7772	13621	4847	<b>2351</b>
$a = 8$	10	<b>7</b>	5	5	1	235.5	357.8	<b>76.0</b>	-	33.9	26.0	<b>24.8</b>	-	28231	48378	18157	<b>8499</b>
BOXPUSHING	69	39	56	<b>59</b>	-	<b>26.8</b>	79.9	63.9	-	<b>9.4</b>	11.0	10.2	-	3075	5360	<b>1932</b>	-
$a = 2$	21	<b>21</b>	<b>21</b>	<b>21</b>	-	<b>14.1</b>	15.6	16.2	-	10.5	10.6	<b>10.3</b>	-	2099	3775	<b>1261</b>	-
$a = 3$	48	18	35	<b>38</b>	-	<b>41.5</b>	118.5	90.2	-	<b>8.2</b>	11.2	10.2	-	3502	6054	<b>2226</b>	-
$l = 0$	21	<b>18</b>	17	<b>18</b>	-	<b>41.5</b>	64.7	53.3	-	<b>8.2</b>	8.3	8.3	-	3373	5887	<b>2116</b>	-
$l > 0$	27	-	18	<b>20</b>	-	-	169.2	<b>123.5</b>	-	-	13.9	<b>12.0</b>	-	3602	6184	<b>2312</b>	-

Table 2: Summary of results for domains requiring concurrency. Planners “2” and “4” are compilations having joint action size bounded by 2 and 4 respectively, while “ $\infty$ ” is the variant with unbounded joint action size. CJR is the compilation proposed by Crosby, Jonsson, and Rovatsos (2014).  $a$  is the number of agents,  $N$  is number of instances; time and length are averages over all instances solved, for all planners that solved at least 5 instances. For BOXPUSHING,  $l$  is the number of large boxes. The total number of actions is an average over all instances (solutions are not required to get this metric).

$r_2$ . As  $n$  grows, the time required by FD for grounding increases due to the memory requirements. For  $n = 6$ , FD needs 6 seconds to find a solution; for  $n = 7$ , it needs 110 seconds; and for  $n > 7$  it surpasses the memory limit.

## Related Work

Several other authors consider the problem of concurrent multiagent planning. Boutilier and Brafman (2001) describe a partial-order planning algorithm for solving MAPs with concurrent actions, based on their formulation of concurrency constraints, but do not present any experimental results. CMAP (Borrajo 2013) produces an initial sequential plan for solving a MAP, but performs a post-processing step to compress the sequential plan into a concurrent plan.

Jonsson and Rovatsos (2011) present a best-response approach for MAPs with concurrent actions, where each agent attempts to improve its own part of a concurrent plan while the actions of all other agents are fixed. However, their approach only serves to improve an existing concurrent plan, and is unable to compute an initial concurrent plan. FMAP (Torreño, Onaindia, and Sapena 2014) is a partial-order planner that also allows agents to execute actions in parallel, but the authors do not present experimental results for MAP domains that require concurrency.

The planner of Crosby, Jonsson, and Rovatsos (2014) is similar to ours in that it also converts CMAPs into classical planning problems. The authors only present results from the MAZE domain, and concurrency constraints are defined as affordances on object sets that appear as arguments of actions. These concurrency constraints are not as flexible as those of Boutilier and Brafman (2001), since the latter can model any arbitrary concurrency constraint between pairs of actions, as well as unidirectional constraints that only affect one of the two actions. Moreover, affordances are not used to define concurrency constraints in conditional effects.

Brafman and Zoran (2014) extended the MA-STRIPS modeling language to support concurrency constraints, and the MAFS multiagent distributed planner to solve this kind of problems while preserving privacy. They examined the scalability of their approach in the BOXPUSHING domain.

Compilations from multiagent to classical planning have also been considered by other authors. Muise, Lipovetzky, and Ramirez (2015) proposed a transformation to respect privacy among agents. The resulting classical planning problem was then solved using a centralized classical planner as in our approach. Besides, compilations to classical planning have also been used in temporal planning, obtaining state-of-the-art results in many of the International Planning Competition domains (Jiménez, Jonsson, and Palacios 2015).

## Conclusion

This work makes several contributions to concurrent planning. A common framework is introduced for different planning forms. We focused on the relation between concurrent and multiagent planning. We proposed a sound and complete method for compiling CMAPs into classical planning problems. The method does not need an exponential number of actions to represent the problem; instead, the number of resulting actions is linear in the description of the CMAP while respecting explicit concurrency constraints.

In future work, it would be interesting to explore strategies for encouraging concurrent actions that involve several agents (i.e. strategies for finding shorter concurrent plans). Furthermore, privacy preserving is a central topic on multiagent planning; thus, this approach could be combined with suitable privacy-preserving mechanisms in the future.

## Acknowledgments

This work has been supported by the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502).

## References

- Borrajó, D. 2013. Plan Sharing for Multi-Agent Planning. In *DMAP 2013 - Proceedings of the Distributed and Multi-Agent Planning Workshop at ICAPS*, 57–65.
- Boutillier, C., and Brafman, R. I. 2001. Partial-Order Planning with Concurrent Interacting Actions. *J. Artif. Intell. Res. (JAIR)* 14:105–136.
- Brafman, R. I., and Domshlak, C. 2008. From One to Many: Planning for Loosely Coupled Multi-Agent Systems. In *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling, ICAPS 2008, Sydney, Australia, September 14–18, 2008*, 28–35.
- Brafman, R. I., and Zoran, U. 2014. Distributed Heuristic Forward Search with Interacting Actions. In *Proceedings of the 2nd ICAPS Distributed and Multi-Agent Planning workshop (ICAPS DMAP-2014)*.
- Crosby, M.; Jonsson, A.; and Rovatsos, M. 2014. A Single-Agent Approach to Multiagent Planning. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, 237–242.
- Crosby, M.; Rovatsos, M.; and Petrick, R. P. A. 2013. Automated Agent Decomposition for Classical Planning. In *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling, ICAPS 2013, Rome, Italy, June 10-14, 2013*.
- Crosby, M. 2013. A temporal approach to multiagent planning with concurrent actions. *PlanSIG*.
- Crosby, M. 2014. *Multiagent Classical Planning*. Ph.D. Dissertation, University of Edinburgh.
- Fox, M., and Long, D. 2003. PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains. *J. Artif. Int. Res.* 20(1):61–124.
- Helmert, M. 2006. The Fast Downward Planning System. *J. Artif. Intell. Res. (JAIR)* 26:191–246.
- Jiménez, S.; Jonsson, A.; and Palacios, H. 2015. Temporal Planning With Required Concurrency Using Classical Planning. In *Proceedings of the 25th International Conference on Automated Planning and Scheduling (ICAPS'15)*.
- Jonsson, A., and Rovatsos, M. 2011. Scaling Up Multiagent Planning: A Best-Response Approach. In *Proceedings of the 21st International Conference on Automated Planning and Scheduling, ICAPS 2011, Freiburg, Germany June 11-16, 2011*.
- Komenda, A.; Stolba, M.; and Kovacs, D. L. 2016. The International Competition of Distributed and Multiagent Planners (CoDMAP). *AI Magazine* 37(3):109–115.
- Kovacs, D. L. 2012. A Multi-Agent Extension of PDDL3.1. In *Proceedings of the 3rd Workshop on the International Planning Competition (IPC)*, 19–27.
- Maliah, S.; Brafman, R. I.; and Shani, G. 2017. Increased Privacy with Reduced Communication in Multi-Agent Planning. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS 2017, Pittsburgh, Pennsylvania, USA, June 18-23, 2017*, 209–217.
- Muise, C.; Lipovetzky, N.; and Ramirez, M. 2015. MAP-LAPKT: Omnipotent Multi-Agent Planning via Compilation to Classical Planning. In *Competition of Distributed and Multiagent Planners*.
- Nardi, D.; Noda, I.; Ribeiro, F.; Stone, P.; von Stryk, O.; and Veloso, M. 2014. RoboCup soccer leagues. *AI Magazine* 35(3):77–85.
- Nissim, R., and Brafman, R. I. 2014. Distributed Heuristic Forward Search for Multi-agent Planning. *J. Artif. Intell. Res.* 51:293–332.
- Richter, S., and Westphal, M. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *Journal of Artificial Intelligence Research* 39:127–177.
- Sheh, R.; Schwertfeger, S.; and Visser, A. 2016. 16 years of robocup rescue. *KI - Künstliche Intelligenz* 30(3):267–277.
- Stolba, M.; Fiser, D.; and Komenda, A. 2016. Potential Heuristics for Multi-Agent Planning. In *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling, ICAPS 2016, London, UK, June 12-17, 2016*, 308–316.
- Torreño, A.; Onaindia, E.; and Sapena, O. 2014. FMAP: Distributed cooperative multi-agent planning. *Appl. Intell.* 41(2):606–626.
- Tozicka, J.; Jakubuv, J.; and Komenda, A. 2014. Generating Multi-Agent Plans by Distributed Intersection of Finite State Machines. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, 1111–1112.