Efficient SAT Approach to Multi-Agent Path Finding Under the Sum of Costs Objective

Pavel Surynek[1] and **Ariel Felner**[2] and **Roni Stern**[3] and **Eli Boyarski**[4]

Abstract. In the *multi-agent path finding* (MAPF) the task is to find non-conflicting paths for multiple agents. In this paper we present the first SAT-solver for the *sum-of-costs* variant of MAPF which was previously only solved by search-based methods. Using both a lower bound on the sum-of-costs and an upper bound on the makespan, we are able to have a reasonable number of variables in our SAT encoding. We then further improve the encoding by borrowing ideas from ICTS, a search-based solver. Experimental evaluation on several domains showed that there are many scenarios where the new SAT-based method outperforms the best variants of previous sum-of-costs search solvers - the ICTS and ICBS algorithms.

1 Introduction and Background

The *multi-agent path finding* (MAPF) problem consists of a graph, $G = (V, E)$ and a set $A = \{a_1, a_2, \ldots a_m\}$ of m agents. Time is discretized into time steps. The arrangement of agents at time-step t is denoted as α_t. Each agent a_i has a start position $\alpha_0(a_i) \in V$ and a goal position $\alpha_+(a_i) \in V$. At each time step an agent can either *move* to an adjacent empty location[5] or *wait* in its current location. The task is to find a sequence of move/wait actions for each agent a_i, moving it from $\alpha_0(a_i)$ to $\alpha_+(a_i)$ such that agents do not *conflict*, i.e., do not occupy the same location at the same time. Formally, an MAPF instance is a tuple $\Sigma = (G = (V, E), A, \alpha_0, \alpha_+)$. A *solution* for Σ is a sequence of arrangements $\mathcal{S}(\Sigma) = [\alpha_0, \alpha_1, ..., \alpha_\mu]$ such that $\alpha_\mu = \alpha_+$ where α_{t+1} results from valid movements from α_t for $t = 1, 2, ..., \mu - 1$. An example of MAPF and its solution are shown in Figure 1.

MAPF has practical applications in video games, traffic control, robotics etc. (see [17] for a survey). The scope of this paper is limited to the setting of *fully cooperative* agents that are centrally controlled. MAPF is usually solved aiming to minimize one of the two commonly-used global cumulative cost functions:

(1) sum-of-costs (denoted ξ) is the summation, over all agents, of the number of time steps required to reach the goal location [8, 23, 18, 17]. Formally, $\xi = \sum_{i=1}^{m} \xi(a_i)$, where $\xi(a_i)$ is an *individual path cost* of agent a_i.

---
[1] Charles University Prague, Malostranské náměstí 25, 11800, Praha, Czech Republic, email: pavel.surynek@mff.cuni.cz
[2] Ben Gurion University, Beer-Sheva, Israel 84105, email: felner@bgu.ac.il
[3] Ben Gurion University, Beer-Sheva, Israel 84105, email: stern-ron@bgu.ac.il
[4] Bar-Ilan University, Ramat-Gan, Israel, email: eli.boyarski@gmail.com

[5] Some variants of MAPF relax the empty location requirement by allowing a chain of neighboring agents to move, given that the head of the chain enters an empty locations. Most MAPF algorithms are robust (or at least easily modified) across these variants.


Figure 1. Example of MAPF for agents a_1, a_2, and a_3 over a 4-connected grid (left) and its optimal solution (right)

(2) makespan: (denoted μ) is the total time until the last agent reaches its destination (i.e., the maximum of the individual costs) [25, 27, 30].

It is important to note that in any solution $\mathcal{S}(\Sigma)$ it holds that $\mu \leq \xi \leq m \cdot \mu$ Thus the optimal *makespan* is usually smaller than the optimal *sum-of-costs*.

Finding optimal solutions for both variants is NP-Hard [34, 25]. Therefore, many suboptimal solvers were developed and are usually used when m is large [15, 6, 20, 14, 12, 32]

1.1 Optimal MAPF Solvers

The focus of this paper is on optimal solvers which are divided into two main classes:

(1) Reduction-based solvers. Many recent optimal solvers reduce MAPF to known problems such as COP [15], SAT [26], Inductive Logic Programming [33] and Answer Set Programming [9]. These papers mostly prove a polynomial-time reduction from MAPF to these problems. These reductions are usually designed for the *makespan* variant of MAPF; they are not applicable for the sum-of-costs variant.

(2) Search-based solvers. By contrast, many recent optimal MAPF solvers are search-based. Some are variants of the A* algorithm on a global *search space* – all different ways to place m agents into V vertices, one agent per vertex [23, 31]. Other employ novel search trees [18, 17, 5]. These search-based solvers are usually designed for the *sum-of-costs* MAPF variant.

A major weaknesses is that connection/comparison between different algorithms was usually done only within a given class of algorithms and cost variant but not across these two classes.

Figure 2. An example of time expansion graph.

1.2 Contributions

This paper aims to start and close the gap. Most of the search-based algorithms can be easily modified to the makespan variant by modifying the cost function and the way the state-space is represented. Some initial directions are given by [17]. By contrast, the reduction-based algorithms are not trivially modified to the sum-of-costs variant and sometimes a completely new reduction is needed.

In this paper we develop the first SAT-based solvers for the sum-of-costs variant which is based on adding *cardinality constraints* [3, 19] for bounding the sum-of-costs. We show how to use known lower bounds on the sum-of-costs to reduce the number of variables that encode these cardinality constraints so as to be practical for current SAT solvers. We then present an *enhanced SAT-solver* which adapts ideas from the ICTS algorithm [18] and uses *multi-value decision diagrams* (MDDs) [22] to further reduce the encoding. This demonstrates the potential of combining ideas from both classes of approaches (search-based and SAT solvers). Experimental results show that our enhanced SAT solver outperforms the best existing search-based solvers for the sum-of-costs variant on many scenarios.

2 SAT Encoding for Optimal Makespan

SAT solvers encompass boolean variables and answer binary questions. The challenge is to apply SAT for MAPF where there is a cumulative cost function. This challenge is stronger for the sum-of-costs variant where each agent has its own cost. We first describe existing SAT encodings for makespan. Then, we present our SAT encoding for sum-of-costs.

A *time expansion graph* (denoted TEG) is a basic concept used in SAT solvers for makespan [27]. We use it too in the sum-of-costs variant below. A TEG is a directed acyclic graph (DAG). First, the set of vertices of the underlaying graph G are duplicated for all time-steps from 0 up to the given bound μ. Then, possible actions (move along edges or wait) are represented as directed edges between successive time steps. Figure 2 shows a graph and its TEG for time steps 0, 1 and 2 (vertical layouts). It is important to note that in this example (1) horizontal edges in TEG correspond to *wait* actions. (2) diagonal moves in TEG correspond to real moves. Formally a TEG is defined as follows:

Definition 1 *Time expansion graph of depth μ is a digraph (V, E) where $V = \{u_j^t | t = 0, 1, ..., \mu \wedge u_j \in V\}$ and $E \subseteq \{(u_j^t, u_k^{t+1}) | t = 0, 1, ..., \mu - 1 \wedge (\{u_j, u_k\} \in E \vee j = k)\}$.*

The encoding for MAPF introduces propositional variables and constraints for a single time-step t in order to represent any possible arrangement of agents at time t. Given a desired makespan μ, the formula represents the question of whether there is a solution in the TEG of μ time steps. The search for optimal makespan is done by iteratively incrementing μ (=0, 1, 2...) until a satisfiable formula is obtained. This ensures optimality in case of a solvable MAPF instance. More information on SAT encoding for the makespan variant can be found, e.g. in [27, 28, 29]

3 Basic-SAT for Optimal Sum-of-costs

The general scheme described above for finding optimal makespan is to convert the optimization problem (finding minimal makespan) to a sequence of decision problems (is there a solution of a given makespan μ). We apply the same scheme for finding optimal sum-of-costs, converting it to a sequence of decision problems – is there a solution of a given sum-of-costs ξ.

However, encoding this decision problem is more challenging than the makespan case, because one needs to both bound the sum-of-costs, but also to predict how many time expansions are needed. We address this challenge by using two key techniques described next: **(1)** Cardinality constraint for bounding ξ and **(2)** Bounding the Makespan.

3.1 Cardinality Constraint for Bounding ξ

The SAT literature offers a technique for encoding a *cardinality constraint* [3, 19], which allows calculating and bounding a numeric cost within the formula. Formally, for a bound $\lambda \in \mathbb{N}$ and a set of propositional variables $X = \{x_1, x_2, ..., x_k\}$ the *cardinality constraint* $\leq_\lambda \{x_1, x_2, ..., x_k\}$ is satisfied iff the number of variables from the set X that are set to TRUE is $\leq \lambda$.

In our SAT encoding, we bound the sum-of-costs by mapping every agent's action to a propositional variable, and then encoding a cardinality constraint on these variables. Thus, one can use the general structure of the makespan SAT encoding (which iterates over possible makespans), and add such a cardinality constraint on top. Next we address the challenge of how to connect these two factors together.

3.2 Bounding the Makespan for the Sum of Costs

Next, we compute how many time expansions (μ) are needed to guarantee that if a solution with sum-of-costs ξ exists then it will be found. In other words, in our encoding, the values we give to ξ and μ must fulfill the following requirement:

R1: *all possible solutions with sum-of-costs ξ must be possible for a makespan of at most μ.*

To find a μ value that meets R1, we require the following definitions. Let $\xi_0(a_i)$ be the cost of the shortest individual path for agent a_i, and let $\xi_0 = \sum_{a_i \in A} \xi_0(a_i)$. ξ_0 was called the *sum of individual costs* (SIC) [18]. ξ_0 is an admissible heuristic for optimal sum-of-costs search algorithms, since ξ_0 is a lower bound on the minimal sum-of-costs. ξ_0 is calculated by relaxing the problem by omitting the other agents. Similarly, we define $\mu_0 = \max_{a_i \in A} \xi_0(a_i)$. μ_0 is length of the *longest* of the shortest individual paths and is thus a lower bound on the minimal makespan. Finally, let Δ be the extra cost over SIC (as done in [18]). That is, let $\Delta = \xi - \xi_0$.

Algorithm 1: SAT consult

1   **MAPF-SAT(MAPF** $\Sigma = (G = (V, E), A, \alpha_0, \alpha_+))$
2     $\mu_0 = \max_{a_i \in A} \xi_0(a_i)$; $\Delta \leftarrow 0$
3   **while** *Solution not found* **do**
4        $\mu \leftarrow \mu_0 + \Delta$;
5   **for** *each agent a_i* **do**
6   build $TEG_i(\mu)$;
7   **end**
8   Solution=Consult-SAT-SOLVER(Σ, μ,Δ);
9   **if** *Solution not found* **then**
10   Δ++;
11   **end**
12   **end**
13   **return** (Solution);
14 **end**

Figure 3. A TEG for an agent that needs to go from u_1 to u_3.

Proposition 1 *For makespan μ of any solution with sum-of-costs ξ, R1 holds for $\mu \leq \mu_0 + \Delta$.*

Proof outline: The worst-case scenario, in terms of makespan, is that all the Δ extra moves belong to a single agent. Given this scenario, in the worst case, Δ is assigned to the agent with the largest shortest-path. Thus, the resulting path of that agent would be $\mu_0 + \Delta$, as required. $\square$

Using Proposition 1, we can safely encode the decision problem of whether there is a solution with sum-of-costs ξ by using $\mu = \mu_0 + \Delta$ time expansions, knowing that if a solution of cost ξ exists then it will be found within $\mu = \mu_0 + \Delta$ time expansions. In other words, Proposition 1 shows relation of both parameters μ and ξ which will be both changed by changing Δ. Algorithm 1 summarizes our optimal sum-of-costs algorithm. In every iteration, μ is set to $\mu_0 + \Delta$ (Line 4) and the relevant TEGs (described below) for the various agents are built. Next a decision problem asking whether there is a solution with sum-of-costs ξ and makespan μ is queried (Line 8). The first iteration starts with $\Delta = 0$. If such a solution exists, it is returned. Otherwise ξ is incremented by one, Δ and consequently μ are modified accordingly and another iteration of SAT consulting is activated.

This algorithm clearly terminates for solvable MAPF instances as we start seeking a solution of $\xi = \xi_0$ ($\Delta = 0$) and increment Δ (which increments ξ and μ as well) to all possible values. The unsolvability of an MAPF instance can be checked separately by a polynomial-time complete sub-optimal algorithm such as PUSH-AND-ROTATE [7].

3.3   Efficient Use of the Cardinality Constraint

The complexity of encoding a cardinality constraint depends linearly in the number of constrained variables [19, 21]. Since each agent a_i must move at least $\xi_0(a_i)$, we can reduce the number of variables counted by the cardinality constraint by only counting the variables corresponding to extra movements over the first $\xi_0(a_i)$ movement a_i makes. We implement this by introducing a TEG for a given agent a_i (labeled TEG_i).

TEG_i differs from TEG (Definition 1) in that it distinguishes between two types of edges: E_i and F_i. E_i are (directed) edges whose destination is at time step $\leq \xi_0(a_i)$. These are called *standard edges*. F_i denoted as *extra edges* are directed edges whose destination is at time step $> \xi_0(a_i)$. Figure 3 shows an underlying graph for agent a_1 (left) and the corresponding TEG_1. Note that the optimal solution of cost 2 is denoted by the diagonal path of the TEG. Edges that belong to F_i are those that their destination is time step 3 (dotted lines). The
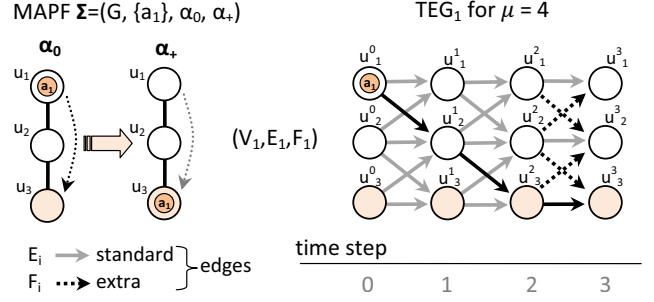
key in this definition is that the cardinality constraint would only be applied to the extra edges, that is, we will only bound the number of extra edges (they sum up to Δ) making it more efficient.

There are various possibilities to define what happens to an agent when it reaches the goal (disappears, waits etc.). In all cases, edges in TEGs corresponding to wait actions at the goal are not marked as extra. Importantly, our SAT approach is robust across all these variants.

3.4   Detailed Description of the SAT Encoding

Agent a_i must go from its initial position to its goal within TEG_i. This simulates its location in time in the underlying graph G. That is, the task is to find a path from $\alpha_0^0(a_i)$ to $\alpha_+^\mu(a_i)$ in TEG_i. The search for such a path will be encoded within the Boolean formula. Additional constraints will be added to capture all movement constraints such as *collision avoidance* etc. And, of course, we will encode the cardinality constraint that the number of extra edges must be exactly Δ.

We want to ask whether a sum-of-costs solution of ξ exists. For this we build TEG_i for each agent $a_i \in A$ of depth $\mu_0 + \Delta$. We use V_i to denote the set of vertices in TEG_i that agent a_i might occupy during the time steps. Next we introduce the Boolean encoding (denoted BASIC-SAT) which has the following Boolean variables:

1:) $\mathcal{X}_j^t(a_i)$ for every $t \in \{0, 1, ..., \mu\}$ and $u_j^t \in V_i$ – Boolean variable of whether agent a_i is in vertex v_j at time step t.

2:) $\mathcal{E}_{j,k}^t(a_i)$ for every $t \in \{0, 1, ..., \mu - 1\}$ and $(u_j^t, u_k^{t+1}) \in (E_i \cup F_i)$ — Boolean variables that model transition of agent a_i from vertex v_j to vertex v_k through any edge (standard or extra) between time steps t and $t + 1$ respectively.

3:) $\mathcal{C}^t(a_i)$ for every $t \in \{0, 1, ..., \mu - 1\}$ such that there exist $u_j^t \in V_i$ and $u_k^{t+1} \in V_i$ with $(u_j^t, u_k^{t+1}) \in F_i$ — Boolean variables that model cost of movements along **extra edges** (from F_i) between time steps t and $t + 1$.

We now introduce constraints on these variables to restrict illegal values as defined by our variant of MAPF. Other variants may use a slightly different encoding but the principle is the same. Let $T_\mu = \{0, 1, ..., \mu - 1\}$. Several groups of constraints are introduced for each agent $a_i \in A$ as follows:

C1: If an agent appears in a vertex at a given time step, then it must follow through exactly one adjacent edge into the next time step. This is encoded by the following two constraints, which are posted

for every $t \in T_\mu$ and $u_j^t \in V_i$

$$\mathcal{X}_j^t(a_i) \Rightarrow \bigvee_{(u_j^t, u_k^{t+1}) \in E_i \cup F_i} \mathcal{E}_{j,k}^t(a_i), \qquad (1)$$

$$\bigwedge_{(u_j^t, u_k^{t+1}),(u_j^t, u_l^{t+1}) \in E_i \cup F_i \wedge k < l} \neg \mathcal{E}_{j,k}^t(a_i) \vee \neg \mathcal{E}_{j,l}^t(a_i) \qquad (2)$$

C2: Whenever an agent occupies an edge it must also enter it before and leave it at the next time-step. This is ensured by the following constraint introduced for every $t \in T_\mu$ and $(u_j^t, u_k^{t+1}) \in E_i \cup F_i$:

$$\mathcal{E}_{j,k}^t(a_i) \Rightarrow \mathcal{X}_j^t(a_i) \wedge \mathcal{X}_k^{t+1}(a_i) \qquad (3)$$

C3: The target vertex of any movement except wait action must be empty. This is ensured by the following constraint introduced for every $t \in T_\mu$ and $(u_j^t, u_k^{t+1}) \in E_i \cup F_i$ such that $j \neq k$.

$$\mathcal{E}_{j,k}^t(a_i) \Rightarrow \bigwedge_{a_l \in A \wedge a_l \neq a_i \wedge u_k^{t+1} \in V_l} \neg \mathcal{X}_k^{t+1}(a_l) \qquad (4)$$

C4: No two agents can appear in the same vertex at the same time step (although the previous constraint ensures that an agent does not collide with an agent currently residing in a vertex it does not prevent simultaneous entering of the same vertex by multiple agents). That is the following constraint is added for every $t \in T_\mu$ and pair of of agents $a_i, a_l \in A$ such that $i \neq l$:

$$\bigwedge_{u_j^t \in V_i \cap V_l} \neg \mathcal{X}_j^t(a_i) \vee \neg \mathcal{X}_j^t(a_l) \qquad (5)$$

C5: Whenever an extra edge is traversed the cost needs to be accumulated. In fact, this is the only cost that we accumulate as discussed above. This is done by the following constraint for every $t \in T_\mu$ and extra edge $(u_j^t, u_k^{t+1}) \in F_i$.

$$\mathcal{E}_{j,k}^t(a_i) \Rightarrow \mathcal{C}^t(a_i) \qquad (6)$$

C6: Cardinality constraint. Finally the bound on the total cost needs to be introduced. Reaching the sum-of-costs of ξ corresponds to traversing exactly Δ extra edges from F_i. The following cardinality constrains ensures this:

$$\leq_\Delta \left\{ \begin{array}{c} \mathcal{C}^t(a_i) | i = 1, 2, ..., n \wedge t = 0, 1, ..., \mu - 1 \\ \wedge \{(u_j^t, u_k^{t+1}) \in F_i\} \neq \emptyset \end{array} \right\} \qquad (7)$$

Final formula. The resulting Boolean formula that is a conjunction of $C1 \ldots C6$ will be denoted as $\mathcal{F}_{BASIC}(\Sigma, \mu, \Delta)$ and is the one that is consulted by Algorithm 1 (line 4).

The following proposition summarizes the correctness of our encoding.

Proposition 2 *MAPF* $\Sigma = (G = (V, E), A, \alpha_0, \alpha_+)$ *has a sum-of-costs solution of* ξ *if and only if* $\mathcal{F}_{BASIC}(\Sigma, \mu, \Delta)$ *is satisfiable. Moreover, a solution of MAPF* Σ *with the sum-of-costs of* ξ *can be extracted from the satisfying valuation of* $\mathcal{F}_{BASIC}(\Sigma, \mu, \Delta)$ *by reading its* $\mathcal{X}_j^t(a_i)$ *variables.*
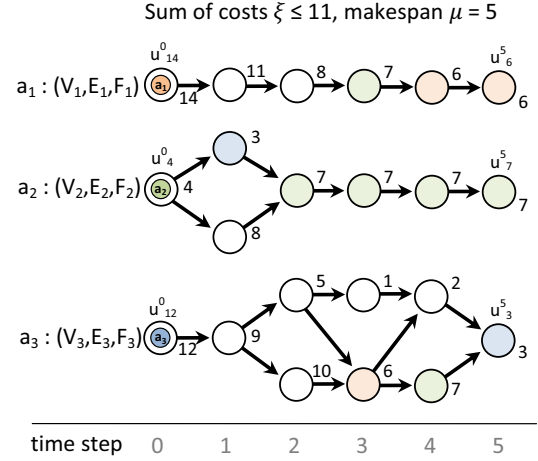

Figure 4. MDDs for agents a_1, a_2, and a_3 for the MAPF from Figure 1 for *sum of individual cost* $\xi \leq 11$.

Proof: The direct consequence of the above definitions is that a valid solution of a given MAPF Σ corresponds to non-conflicting paths in the TEGs of the individual agents. These non-conflicting paths further correspond to satisfying the variable assignment of $\mathcal{F}_{BASIC}(\Sigma, \mu, \Delta)$, i.e., that there are Δ extra edges in TEGs of depth $\mu = \mu_0 + \Delta$. $\square$

Proposition 3 *Let D be the maximal degree of any vertex in G and let m be the number of agents. If $m \cdot |E| \geq \Delta$ and $m \geq D$ then the number of clauses in $\mathcal{F}_{BASIC}(\Sigma, \mu, \Delta)$ is $O(\mu \cdot m^2 \cdot |E|)$, and the number of variables is $O(\mu \cdot |E| \cdot m)$.*

Proof: The components of $\mathcal{F}_{BASIC}(\Sigma, \mu, \Delta)$ is described in equations 1– 7. Equation 1 introduces at most $O(m \cdot \mu \cdot |E|)$ clauses. Equation 2 introduces at most $O(m \cdot \mu |E| \cdot D)$ clauses. Equation 3 introduces at most $O(m \cdot \mu \cdot |E|)$ clauses. Equation 4 introduces at most $O(m^2 \cdot \mu \cdot |E|)$ Equation 5 introduces at most $O(m^2 \cdot \mu \cdot |V|)$ clauses. Equation 6 introduces at most $O(m \cdot \mu \cdot |E|)$ clauses. Equation 7 introduces at most $O(m \cdot \mu \cdot (\xi - \xi_0))$ clauses, since a cardinality constraint checking that n variables has a cardinality constraint of m requires $O(n \cdot m)$ clauses [21]. Summing all the above results in a total of $O(\mu \cdot m \cdot (|E| \cdot (D + m) + (\xi - \xi_0)))$. If we assume that $m > D$ and that $m \cdot |E| > (\xi - \xi_0)$ then the number of clauses is $O(\mu \cdot m^2 \cdot |E|)$. The number of variables is easily computed in a similar way. $\square$

4 Improving Basic SAT by Adding MDDs

A major parameter that affects the speed of solving of Boolean formulae is their size [13]. The size of formulae in the BASIC-SAT encoding is affected mostly by the size of the TEGs (this is embodied in the $|E|$ factor in the encoding size). To obtain a significant speedup we reduce the size of TEG_i for agent a_i in terms of number of vertices while the soundness of encoding is preserved. To do this we borrow the ideas of *multi-value Decision Diagram* (MDD) from the search-based MAPF algorithm ICTS [18]. This shows the advantage of combining techniques from both classes of approaches (search-based and SAT).

Let TEG_i^μ denote TEG_i for μ time expansions. We set $\mu = \mu_0 + \Delta$ in our solution. The data structure we use for reducing
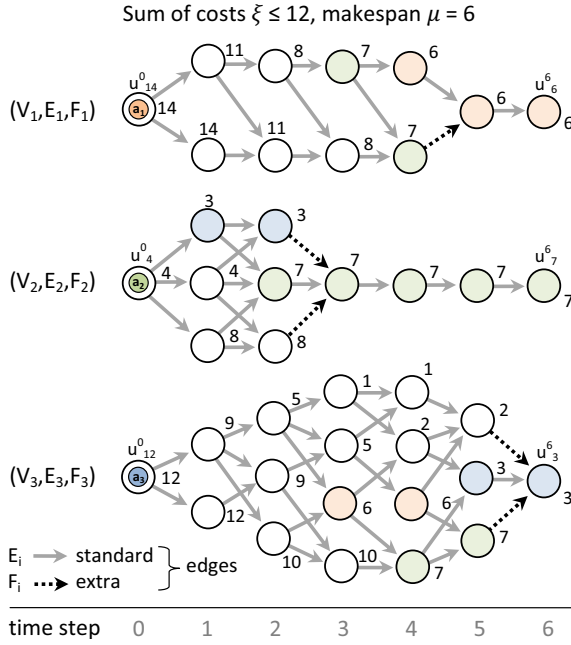
Figure 5. MDDs from Figure 4 for the incremented *sum of individual cost* from 11 to 12 ($\xi \leq 12$).

Figure 6. An illustration of MDD-SAT encoding using MDDs from Figure 5. Mutual exclusion constraints (C4) that prevent multiple agent occurrence in the same vertex are shown using dashed edges.

TEG_i^μ is a *multi-value Decision Diagram* (MDD). MDDs were already used in the search-based MAPF algorithm ICTS [18]. In our context, MDD_i^μ is a digraph that represents all possible valid paths from $\alpha_0(a_i)$ to $\alpha_+(a_i)$ of cost μ for agent a_i. MDD_i^μ has a single *source node* at level 0 and a single *sink node* at level μ. Every node at depth t of MDD_i^μ corresponds to a possible location of a_i at time t, that is on a path of cost μ from $\alpha_0(a_i)$ to $\alpha_+(a_i)$. It is easy to see that MDD_i^μ is subgraph of TEG_i. While TEG_i^μ includes all vertices of G at each time step, MDD_i^μ includes only those vertices and edges that represent possible valid paths, and thus vertices not in MDD_i^μ can be ignored.

Moreover, the maximum cost that can be consumed by single agent a_i under given sum-of-costs bound ξ is $\xi_0(a_i) + \Delta$ where, as defined above, $\xi_0(a_i)$ is the shortest path connecting $\alpha_0(a_i)$ with $\alpha_+(a_i)$ in G (assuming no other agent exist). Thus, it is sufficient to replace TEG_i^μ with $MDD_i^{\xi_0(a_i)+\Delta}$, which is useful since $\xi_0(a_i) + \Delta \leq \mu_0 + \Delta = \mu$.

MDDs for the agents of Figure 1 are shown in Figures 4 and 5. Indeed, the size of the MDDs is much smaller than the corresponding TEGs which include all states for all time steps. Though the increase in size caused by ability to reach more vertices under given sum-of-costs bounds is observable between Figures 4 and 5.

Grid 8x8	BASIC-SAT		MDD-SAT	
m	*Variables*	*Clauses*	*Variables*	*Clauses*
1	1 552.8	11 617.6	20.6	27.9
4	14 712.0	127 732.2	276.5	554.0
8	226 391.2	2 099 127.6	18 355.6	68 826.0
16	4 075 187.2	32 108 347.2	2 253 508.2	13 128 646.9

Table 1. The effect of using MDDs in the encoding in terms of the number of variables and clauses.

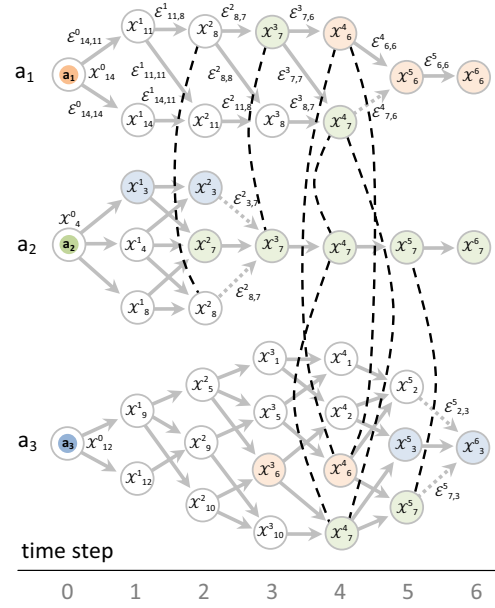The encoding that uses MDD-based time expansion will be called

MDD-SAT and the corresponding formulae will be denoted as $\mathcal{F}_{MDD}(\Sigma, \mu, \Delta)$. $\mathcal{F}_{MDD}(\Sigma, \mu, \Delta)$ are similar to BASIC-SAT. The only difference is that in BASIC-SAT there is a variable for all vertices and edges of the TEGs while in MDD-SAT, only variables for the vertices and edges of the MDDs are needed. This difference can be significant. Table 1 presents the number of propositional variables and clauses accumulated over all the constructed formulae for a given MAPF instance for BASIC-SAT and for MDD-SAT over 8×8 grid with 10% obstacles. The average values out of 10 random instances per number of agents is shown. Up to two orders of magnitude reduction is shown.

An illustration of the $\mathcal{F}_{MDD}(\Sigma, \mu, \Delta)$ formula is shown in Figure 6. It is particularly observable that MDDs reduce the number of mutual exclusion constraints (dashed edges) by omitting unreachable vertices (and all the constraints incident with them).

5 Experimental Evaluation

We experimented on 4-connected grids with randomly placed obstacles [20, 23] and on *Dragon Age* maps [17, 24]. Both settings are a standard MAPF benchmarks. The initial position of the agents was randomly selected. To ensure solvability the goal positions were selected by performing a long *random walk* from the initial arrangement.

We compared our SAT solvers to several state-of-the-art search-based algorithms: the *increasing cost tree search* - ICTS [18], *Enhanced Partial Expansion A** - EPEA* [10] and *improved conflict-based search* - ICBS [5]. For all the search algorithms we used the best known setup of their parameters and enhancements suitable for solving the given instances over 4-connected grids.

The SAT approaches were implemented in C++. The implementation consists of a top level algorithm for finding the optimal sum-of-costs ξ and *CNF* formula generator [4] that prepares input formula for a SAT solver into a file. The SAT solver is an external module
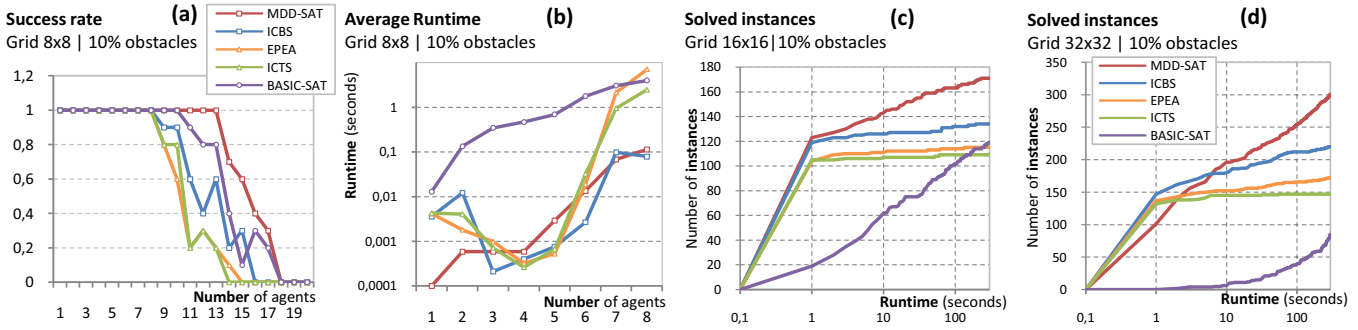
Figure 7. Results on 8×8 grid (left). Number of solved instances in the given runtime on 16×16 and 32×32 grids. (right)

of this architecture. We used `Glucose 3.0` [2, 1] which is a top performing SAT solver in the *SAT Competition* [11, 27].

The cardinality constraint was encoded using a simple standard circuit based encoding called *sequential counter* [21]. In our initial testing we considered various encodings of the cardinality constraint such as those discussed in [3, 19]. However, it turned out that changing the encoding has a minor effect.[6]

ICTS and ICBS were implemented in C#, based on their original implementation (here we used a slight modification in which the target vertex of a move must be empty). All experiments were performed on a Xeon 2Ghz, and on Phenom II 3.6Ghz, both with 12 Gb of memory.

5.1 Square Grid Experiments

We first experimented on 8×8, 16×16, and 32×32 grids with 10% obstacles while varying the number of agents from 1 up to the number where at least one solver was able to solve an instance (in case of the 8×8 grid this is 20 agents; and 32 and 58 in case of 16×16 and 32×32 grids respectively). For each number of agents 10 random instances were generated.

Figure 7 presents results where each algorithm was given a time limit of 300 seconds (as was done by [18, 5, 16]). The leftmost plot (Plot (a)) shows the *success rate* (=percentage out of given 10 random instances solved within the time limit) as a function of the number of agents for the 8×8 grid (higher curves are better). The next plot (Plot (b)) reports the average runtime for instances that were solved by all algorithms (lower curves are better). Here, we required 100% success rate for all the tested algorithms to be able to calculate average runtime; this is also the reason why the number of agents is smaller. The two right plots visualize the results on 16×16 grid (Plot (c)) and 32×32 grid (Plot (d)) but in a different way. Here, we present the number of instances (out of all instances for all number of agents) that each method solved (y-axis) as a function of the elapsed time (x-axis). Thus, for example Plot (c) says that MDD-SAT was able to solve 145 instances in time less than 10 seconds (higher curves are better).

The first clear trend is that MDD-SAT significantly outperforms BASIC-SAT in all aspects. This shows the importance of developing enhanced SAT encodings for the MAPF problem. The performance of the BASIC-SAT encoding compared to the search-based algorithm degrades as the size of the grids grow larger: in the 8x8 grids it is

Figure 8. Success rate and runtime on the 8×8 grid with increasing number of obstacles (out of 64 cells).

second only to MDD-SAT, in the 16x16 grid it is comparable to most search-based algorithms, and in the 32x32 grid it is even substantially worse. For the rest of the experiments we did not activate BASIC-SAT.

In addition, a prominent trend observed in all the plots is that MDD-SAT has higher success rate and solves more instances than all other algorithms. In particular, in highly constrained instances (containing many agents) the MDD-SAT solver is the best option.

However, on the 32×32 grid (rightmost figure) for easy instances when the available runtime was less than 10 seconds, MDD-SAT was weaker than the search-based algorithms. This is mostly due to the architecture of the MDD-SAT solver which has an overhead of running the external SAT solver and passing input in the textual form to it. This effect is also seen in the 8x8 plot (Plot (b)) as these were rather easy instances (solved by all algorithms) and the extra overhead of activating the external SAT solver did not pay off.

Next, we varied the number of obstacles for the 8×8 grid with 10 agents to see the impact of shrinking free space and increasing the frequency of interactions among agents. Results are shown in Figure 8. Again, MDD-SAT clearly solves more instances over all settings. MDD-SAT was always faster except for some easy instances (that needed up to 1 second) where ICBS was slightly faster which is again due to the overhead in setup of the SAT solving by an external solver. Interestingly, increasing the number of obstacles reduces the number of open cells. This is an advantage for the SAT formula generator in MDD-SAT as the formula has less variables and constraints. By contrast, the combinatorial difficulty of the instances increases with adding obstacles for all the solvers as it means that the graphs gets denser and harder to solve.

---

[6] Due to the knowledge of lower bounds on the sum-of-costs, the number of variables involved in the cardinality constraint is relatively small and hence the different encoding style has not enough room to show its benefit.

brc202d   **den520d**   **ost003d**

Figure 9. Three structurally diverse `Dragon-Age` maps used in the experimental evaluation. This selection includes: narrow corridors in `brc202d`, large open space in `den520d`, and open space with almost isolated rooms in `ost003d`.

5.2 Results on the Dragon Age Maps

Next, we experimented on three structurally different Dragon-Age maps - `ost003d`, `den520d`, and `brc202d`, that are commonly used as testbeds [18, 10, 5] - see Figure 9. On these maps we only evaluated the most efficient algorithms, namely, MDD-SAT, ICTS, and ICBS. Generally, in these maps there is a large number of open cells but the graph is sparse with agents but there are topological differences. `brc202d` has many narrow corridors. `ost003d` consists of few open areas interconnected by narrow doors. Finally, `den520d` has wider open areas.

To obtain instances of various difficulties we varied the distance between start and goal locations. Ten random instances were generated for each distance in the range: $\{8, 16, 24, \ldots, 320\}$ in order to have instances of different difficulties (total of 400 instances). With larger distances, the problems are more difficult as the probability for interactions (avoidance) among agents increases as they need to travel through a larger part of the graph.

The results for the three Dragon-Age maps are shown in Figure 10 (`brc202d`), Figure 11 (`den520d`), and Figure 12 (`ost003d`). Two setups were used for each map - one with 16 agents, the other with 32 agents. The left plot of each figure shows the number of solved instances (y-axis) as a function of the elapsed time (x-axis). Again, higher curves correspond to better performance). The right plot is interpreted as follows. For each solver the 400 instances are ordered in increasing order of their solution time (this has strong correlation with the distance between the start and goal configurations). Thus, the numbers in the x-axis give the relative location (out of the 400) in this sorted order. The y-axis gives the actual running time for each instance. Here, lower curves correspond to better performance.

All these figures show a similar clear trend with the exception of `ost003d` with 32 agents (discussed below). On the easy instances where little time is required (left of the figures), MDD-SAT is not the best. But, for the harder instances that need more time (right of the figures), MDD-SAT clearly outperform all the other solvers.

Intuitively, one might think that the search-based solvers will have an advantage in these domains since they contain many open spaces (low combinatorial difficulty) while the MDD-SAT approach will suffer here as it will need to generate a large number of formulae (as the domains are large). This might be true for the easy instances. Nevertheless, the effectiveness of MDD-SAT was clearly seen on the harder instances where generating the formulae and the external time to activate the architecture of the SAT solver seemed to pay off. This trend was also seen in the case of small densely occupied grids discussed above.
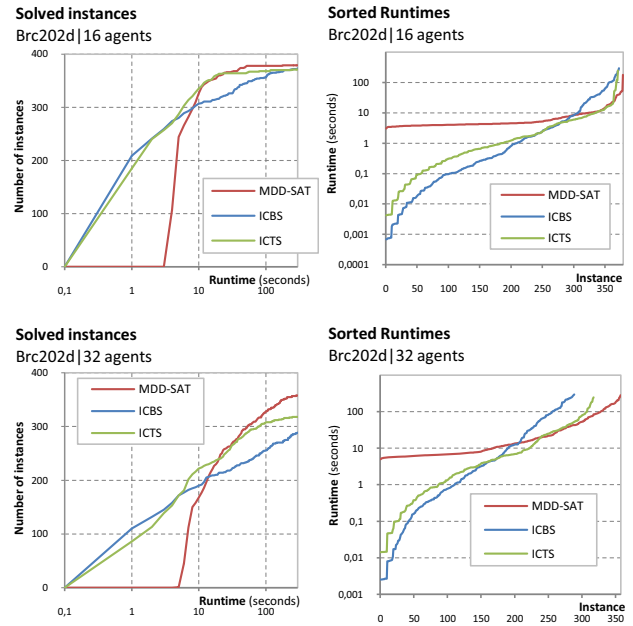

Figure 10. Results for dragon age map `brc202d` with 16 and 32 agents. The left part shows the number of instances (y-axis) a solver manages to solve in the given timeout (x-axis). The right part shows all the runtimes for a given solver sorted in the ascending order.
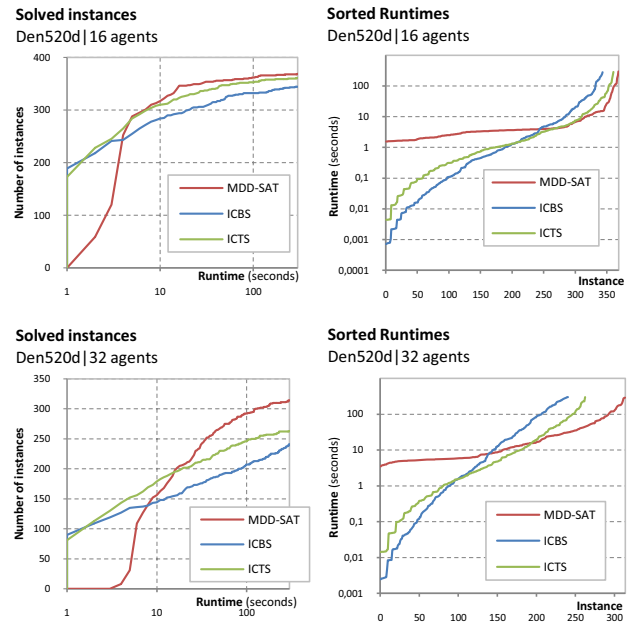

Figure 11. Results for dragon age map `den520d` with 16 and 32 agents. MDD-SAT is the best option on hard instances with more agents.

The `ost003d` map with 32 agents is the only case where MDD-SAT was outperformed by ICTS. This is probably due to the specific structure of `ost003d` which has a number of isolated open spaces. This gives an advantage to ICTS with relatively many agents (32) as conflicts mostly occur at the exits/doors of the open areas. ICTS handles this on a per-agent cost basis while the other solvers are less effective here.
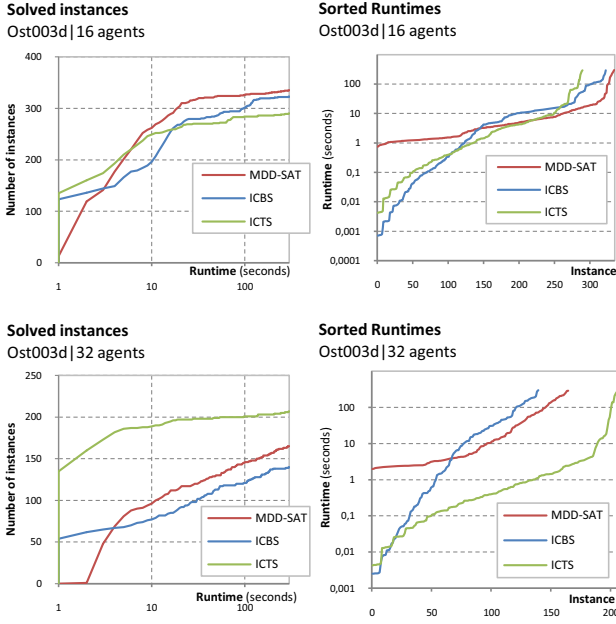
Figure 12. Results for dragon age map `ost003d` with 16 and 32 agents. Although MDD-SAT performs as best with 16 agents, it gets outperformed in the case with 32 agents by ICTS. This case shows that there is no universal winner among the tested algorithms.

The entire set of experiments show a clear trend. For the easy instances when a small amount of time is given the search-based algorithm may be faster. But, given enough time MDD-SAT is the correct choice, even in the large maps where it has an initial disadvantage. One of the reasons for this is modern SAT solvers have the ability to learn and improve their speed during the process of answering a SAT question. But, this learning needs sufficient time and large search trees to be effective. By contrast, search algorithms do not have this advantage.

5.3 Size of the Formulae

Concrete runtimes for 10 instances of *ost003d* are given in Table 2. MDD-SAT solves the hardest instance (#1) while other solvers ran out of time. The right part of the table illustrates the cumulative size of the formulae generated during the solving process. Although the map is much larger than the square grids, the size of formulae is comparable to the densely occupied grid (see Figure 1). This is because ξ_0 is a good lower bound of the optimal cost in the sparse maps.

The observation from this experiments is that the large underlaying graph does not necessarily imply generating large Boolean formulae in the MDD-SAT solving process. Though in harder scenarios (where start and goals are far apart) large formulae are eventually generated but still do not represent any significant disadvantage for the MDD-SAT solver according to presented measurements. We observed that generating large formulae takes considerable portion of the total runtime (up to 10%-30%) within the MDD-SAT solver. Hence efficient implementation of this part of the solver has significant impact on the overall performance.

6 Summary and Conclusions

We introduced the first state-of-the-art SAT-based solver for the sum-of-costs variant of MAPF. The resulting enhanced encoding, called MDD-SAT migrates ideas from the search-based methods to use with SAT solvers, was experimentally compared to the state-of-the-art search-based solvers over a variety of domains - we tested 4-connected grids with random obstacles and large maps from computer games. We have seen that MDD-SAT is a better option in hard scenarios while the search-based solvers may perform better in easier cases.

Nevertheless, as previous authors mentioned [17, 5] there is no universal winner and each of the approaches has pros and cons and thus might work best in different circumstances. For example, ICTS was best on `ost003d` with 32 agents. This calls for a deeper study of various classes of MAPF instances and their characteristics and how the different algorithms behave across them. Not too much is known at present to the MAPF community on these aspects.

There are several factors behind the performance of the SAT-based approach: clause learning, constraint propagation, good implementation of the SAT solver. On the other hand, the SAT solver does not understand the structure of the encoded problem which may downgrade the performance. Hence, we consider that implementing techniques such as learning directly into the dedicated MAPF solver may be a future direction. Finally, migrating of other ideas from both classes of approaches might further improve the performance.

MAPF	Ost003d (seconds) 16 agents, distance=168			m	MDD-SAT, 16 agents	
	MDD-SAT	*ICBS*	*ICTS*	**Distance**	*Variables*	*Clauses*
1	**101.4**	N/A	N/A	8	758.0	1 169.7
2	12.8	9.7	**2.4**	64	34 648.7	120 961.1
3	13.2	4.4	**2.4**	128	932 440.9	9 128 568.8
4	3.8	**0.6**	1.2			
5	13.5	9.6	**3.2**	m	MDD-SAT, 32 agents	
6	22.7	**10.7**	N/A	**Distance**	*Variables*	*Clauses*
7	N/A	N/A	N/A			
8	36.9	49.6	**2.5**	8	2 377.6	3 751.3
9	12.0	2.6	**1.4**	64	571 915.1	3 672 249.3
10	N/A	N/A	N/A	128	5 163 157.0	49 201 960.0

Table 2. Runtime for 10 instances (left) and the average size of the MDD-SAT formulae for `ost003d` (right)

7 Acknowledgements

REFERENCES

[1] G. Audemard, J. Lagniez, and L. Simon, 'Improving glucose for incremental SAT solving with assumptions: Application to MUS extraction', in *Theory and Applications of Satisfiability Testing - SAT 2013*, pp. 309–317, (2013).

[2] G. Audemard and L. Simon, 'Predicting learnt clauses quality in modern SAT solvers', in *IJCAI*, pp. 399–404, (2009).

[3] O. Bailleux and Y. Boufkhad, 'Efficient CNF encoding of boolean cardinality constraints', in *CP*, pp. 108–122, (2003).

[4] A. Biere, A. Biere, M. Heule, H. van Maaren, and T. Walsh, *Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications*, IOS Press, 2009.

[5]   E. Boyarski, A. Felner, R. Stern, G. Sharon, D. Tolpin, O. Betzalel, and S. Shimony, 'ICBS: improved conflict-based search algorithm for multi-agent pathfinding', in *IJCAI*, pp. 740–746, (2015).

[6]   L. Cohen, T. Uras, and S. Koenig, 'Feasibility study: Using highways for bounded-suboptimal mapf', in *SOCS*, pp. 2–8, (2015).

[7]   B. de Wilde, A. ter Mors, and C. Witteveen, 'Push and rotate: a complete multi-agent pathfinding algorithm', *JAIR*, **51**, 443–492, (2014).

[8]   K. Dresner and P. Stone, 'A multiagent approach to autonomous intersection management', *JAIR*, **31**, 591–656, (2008).

[9]   E. Erdem, D. G. Kisa, U. Oztok, and P. Schueller, 'A general formal framework for pathfinding problems with multiple agents', in *AAAI*, (2013).

[10] M. Goldenberg, A. Felner, R. Stern, G. Sharon, N. Sturtevant, R. Holte, and J. Schaeffer, 'Enhanced partial expansion A*', *JAIR*, **50**, 141–187, (2014).

[11] M. Järvisalo, D. Le Berre, O. Roussel, and L. Simon, 'The international SAT solver competitions', *AI Magazine*, **33**(1), (2012).

[12] M. M. Khorshid, R. C. Holte, and N. R. Sturtevant, 'A polynomial-time algorithm for non-optimal multi-agent pathfinding', in *Symposium on Combinatorial Search (SOCS)*, (2011).

[13] Justyna Petke, *Bridging Constraint Satisfaction and Boolean Satisfiability*, Artificial Intelligence: Foundations, Theory, and Algorithms, Springer, 2015.

[14] G. Röger and M. Helmert, 'Non-optimal multi-agent pathfinding is solved (since 1984).', in *SOCS*, (2012).

[15] M. Ryan, 'Constraint-based multi-robot path planning', in *ICRA*, pp. 922–928, (2010).

[16] G. Sharon, R. Stern, A. Felner, and N. Sturtevant, 'Conflict-based search for optimal multi-agent pathfinding', *Artif. Intell.*, **219**, 40–66, (2015).

[17] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, 'Conflict-based search for optimal multi-agent pathfinding', *Artif. Intell.*, **219**, 40–66, (2015).

[18] G. Sharon, R. Stern, M. Goldenberg, and A. Felner, 'The increasing cost tree search for optimal multi-agent pathfinding', *Artificial Intelligence*, **195**, 470–495, (2013).

[19] J. Silva and I. Lynce, 'Towards robust CNF encodings of cardinality constraints', in *CP*, pp. 483–497, (2007).

[20] D. Silver, 'Cooperative pathfinding', in *AIIDE*, pp. 117–122, (2005).

[21] C. Sinz, 'Towards an optimal CNF encoding of boolean cardinality constraints', in *CP*, pp. 827–831, (2005).

[22] A. Srinivasan, T. Ham, S. Malik, and R. Brayton, 'Algorithms for discrete function manipulation', in *(ICCAD*, pp. 92–95, (1990).

[23] T. Standley, 'Finding optimal solutions to cooperative pathfinding problems.', in *AAAI*, pp. 173–178, (2010).

[24] Nathan R. Sturtevant, 'Benchmarks for grid-based pathfinding', *Computational Intelligence and AI in Games*, **4**(2), 144–148, (2012).

[25] P. Surynek, 'An optimization variant of multi-robot path planning is intractable', in *AAAI*, (2010).

[26] P. Surynek, 'Towards optimal cooperative path planning in hard setups through satisfiability solving', in *PRICAI*, 564–576, (2012).

[27] P. Surynek, 'Compact representations of cooperative path-finding as SAT based on matchings in bipartite graphs', in *ICTAI*, pp. 875–882, (2014).

[28] P. Surynek, 'A simple approach to solving cooperative path-finding as propositional satisfiability works well', in *PRICAI*, pp. 827–833, (2014).

[29] P. Surynek, 'Simple direct propositional encoding of cooperative path finding simplified yet more', in *MICAI*, pp. 410–425, (2014).

[30] P. Surynek, 'Reduced time-expansion graphs and goal decomposition for solving cooperative path finding sub-optimally', in *IJCAI*, pp. 1916–1922, (2015).

[31] G. Wagner and H. Choset, 'Subdimensional expansion for multirobot path planning', *Artif. Intell.*, **219**, 1–24, (2015).

[32] K. Wang and A. Botea, 'MAPP: a scalable multi-agent path planning algorithm with tractability and completeness guarantees', *JAIR)*, **42**, 55–90, (2011).

[33] J. Yu and S. LaValle, 'Planning optimal paths for multiple robots on graphs', in *ICRA*, pp. 3612–3617, (2013).

[34] J. Yu and S. M. LaValle, 'Structure and intractability of optimal multi-robot path planning on graphs', in *AAAI*, (2013).