

Candidate Path Selection Heuristics for Multi-Agent Path Finding: A Novel Compilation-Based Method

Pavel Surynek

Faculty of Information Technology, Czech Technical University in Prague
Thakurova 9, 160 00 Praha 6, Czech Republic

Keywords: Multi-Agent Path Finding, Problem Compilation, Boolean Satisfiability, Path Selection Heuristics.

Abstract: Multi-Agent path finding (MAPF) is a task of finding non-conflicting paths connecting agents' specified initial and goal positions in a shared environment. We focus on compilation-based solvers in which the MAPF problem is expressed in a different well established formalism such as mixed-integer linear programming (MILP), Boolean satisfiability (SAT), or constraint programming (CP). As the target solvers for these formalisms act as black-boxes it is challenging to integrate MAPF specific heuristics in the MAPF compilation-based solvers. We show in this work how to build a MAPF encoding for the target SAT solver in which domain specific heuristic knowledge is reflected.

1 INTRODUCTION AND BACKGROUND

Multi-agent path finding (MAPF) represents a fundamental problem in artificial intelligence (Silver, 2005; Ryan, 2007; Standley, 2010; Luna and Bekris, 2011; Yu and LaValle, 2013a). The task is to navigate each agent from the set of agents $A = \{a_1, a_2, \dots, a_k\}$ from its initial position to a specified goal position. In the standard discrete variant of MAPF, the environment is modeled as an undirected graph $G = (V, E)$ where vertices represent positions and edges represent the topology across which the agents move between vertices. There are two requirements that make the MAPF problem challenging: (1) the agents must not collide with each other, that is they never can share a vertex nor can traverse an edge in opposite directions and (2) an objective such as the total number of move actions must be optimized.

We address the MAPF problem from the perspective of compilation-based techniques. *Compilation* is one of the most important techniques used across computing. In the context of problem solving in artificial intelligence, the compilation approach reduces an input problem instance from its source formalism to a different, usually well established, target formalism for which an efficient solver exists.

Target formalisms are often combinatorial optimization frameworks like *constraint programming / optimization* (CP) (Dechter, 2003), *mixed integer*

linear programming (MILP) (Jünger et al., 2010; Rader, 2010), Boolean satisfiability (SAT) (Biere et al., 2009), satisfiability modulo theories (SMT) (Barrett and Tinelli, 2018), or answer set programming (ASP) (Lifschitz, 2019).

Currently **compilation-based** optimal solvers for MAPF represent a major alternative to **search-based** solvers, that model and solve the problem directly.

We focus in this paper on a recent SMT-based approach to solving MAPF (Surynek, 2021) that uses multi-level lazy compilation of MAPF to Boolean satisfiability. The previous approach builds on top the SMT-CBS, an optimal algorithm, (Surynek, 2019), that encodes an incomplete specification of the input MAPF instance into SAT. Concretely, collision avoidance constraints are not encoded at the beginning which may result in solutions that lead to a collisions. After the collisions are detected, collision avoidance constraints are added to the encoding and the process is repeated until the collision-free solution is obtained.

The further step from SMT-CBS is the next level of laziness called *sparsification* that adds laziness in encoding of the set of candidate paths for each agent. While the original SMT-CBS gives each agent a chance to choose any of the possible paths, the sparse version starts with a restricted set of candidate paths containing only the most promising path. If the search for non-colliding paths with sparse set of candidate paths is unsuccessful, the set of candidate paths is ex-

tended. The process may end up with all paths included but again the solution is often found before all paths are considered.

Since Boolean formulae, to which the input MAPF instance is reduced in SMT-CBS, are derived from the set of candidate paths, the effect of sparsification of the set is twofold: **(1)** it leads to smaller target Boolean formulae that can be constructed faster and **(2)** the satisfiability of formulae can be decided by the SAT solver faster, altogether improving the reduction-solving-interpretation loop in SMT-CBS.

The original sparsification technique used the selection of candidate paths so that at least one path for any subset of detected collisions is taken. Since the collisions are treated using the ‘or’ connective in this approach, we call it an *OR-path selection heuristic*.

In this work, we contribute by a different heuristically guided selection of candidate paths. In contrast to OR-path selection, we suggest to include one path that avoids all recently discovered collisions. As collisions are threatened via the ‘and’ connective, we call our new approach an *AND-path selection heuristic*.

An extended version of this paper is available (Surynek, 2022).

1.1 Multi-Agent Path Finding

We assume discrete time in MAPF. The configuration of agents at timestep t is denoted as s_t . Each agent a_i has a start position $s_0(a_i) \in V$ and a goal position $s_+(a_i) \in V$.

At each time step an agent can either *move* to an adjacent vertex or *wait* in its current vertex. The task is to find a sequence of move/wait actions for each agent a_i , moving it from $s_0(a_i)$ to $s_+(a_i)$ such that agents do not *collide*, i.e., do not occupy the same location at the same time and do not traverse the same edge in opposite directions.

Formally, a MAPF instance is a tuple $\Sigma = (G = (V, E), R, s_0, s_+)$ where $s_0 : R \rightarrow V$ is an initial configuration of agents and $s_+ : R \rightarrow V$ is a goal configuration of agents. A *solution* for Σ is a sequence of configurations $\mathcal{S}(\Sigma) = [s_0, s_1, \dots, s_\mu]$ such that s_{t+1} results from valid movements from s_t for $t = 1, 2, \dots, \mu - 1$, and $s_\mu = s_+$. Orthogonally to this, the solution can be represented as a set of paths for individual agents.

Various cumulative objectives are often optimized in MAPF. We will develop all concepts in this paper for the *sum-of-costs* objective, denoted *SoC*. *SoC* is the summation, over all agents, of the number of time steps required to reach the goal. Formally, $SoC = \sum_{i=1}^k cost(path(a_i))$, where $cost(path(a_i))$ is an *individual path cost* of agent a_i connecting $s_0(a_i)$ calculated as the number of edge traversals and wait

actions.¹

Finding an optimal solution with respect to the sum-of-costs objective is NP-hard (Yu and LaValle, 2013b; Surynek, 2010) and also determining the existence of a solution that differs from the optimum by a factor less than $4/3$ is NP-hard too (Ma et al., 2016). Therefore designing algorithms based on search and SAT for MAPF is justifiable.

1.2 Compilation-Based Approaches

The idea behind compilation that uses the SAT paradigm is to construct a Boolean formula whose satisfiability corresponds to existence of a solution of the given value of the objective, that is, the formula encodes bounded instance of MAPF. In our case we are encoding the question whether there is a solution of sum-of-costs *SoC* to a given MAPF Σ .

There are two ways how to connect satisfiability of the formula and solvability of Σ : using either **equivalence** or **implication**.

We say $\mathcal{F}(SoC)$ to be a *complete Boolean model* of MAPF.

Definition 1. (Complete Boolean Model) Boolean formula $\mathcal{F}(SoC)$ is a complete Boolean model of MAPF Σ if the following condition holds: $\mathcal{F}(SoC)$ is satisfiable $\Leftrightarrow \Sigma$ has a solution of sum-of-costs *SoC*.

Complete Boolean models were used in makespan optimal SAT-based solvers for MAPF (Surynek, 2017) and in MDD-SAT (Surynek et al., 2016), the first sum-of-costs optimal SAT-based solver. A natural relaxation from the complete Boolean model is an *incomplete Boolean model* where instead of the equivalence between solving MAPF and the formula we require an implication only. Incomplete models are inspired from the SMT paradigm and are used in the recent sum-of-costs optimal solver SMT-CBS (Surynek, 2019; Surynek, 2021).

Definition 2. (Incomplete Boolean Model). Boolean formula $\mathcal{H}(SoC)$ is an incomplete Boolean model of MAPF Σ if the following condition holds: $\mathcal{H}(SoC)$ is satisfiable $\Leftarrow \Sigma$ has a solution of sum-of-costs *SoC*.

Being able to construct formula \mathcal{F} one can obtain optimal MAPF solution by checking satisfiability of $\mathcal{F}(0), \mathcal{F}(1), \mathcal{F}(2), \dots$ until the first satisfiable $\mathcal{F}(SoC)$ is met. This is possible due to monotonicity of MAPF solvability with respect to increasing values

¹The notation $path(a_i)$ refers to path in the form of a sequence of vertices and edges connecting $s_0(a_i)$ and $s_+(a_i)$ while $cost$ assigns the cost to a given path.

of common cumulative objectives such as the sum-of-costs. In practice it is however impractical to start at 0; lower bound estimation is used instead - sum of lengths of shortest paths can be used in the case of sum-of-costs.

Construction of $\mathcal{F}(SoC)$ relies on the time expansion of underlying graph G . Having SoC , the basic variant of time expansion determines the maximum number of time steps μ (also referred to as a *makespan*) such that every possible solution of the given MAPF with the sum-of-costs less than or equal to SoC fits within μ timesteps (that is, no agent is outside its goal vertex after μ -th timestep if the sum-of-costs SoC is not to be exceeded).

The time expansion itself makes copies of vertices V for each timestep $t = 0, 1, 2, \dots, \mu$. That is, we have vertices v^t for each $v \in V$ time step t . Edges from G are converted to directed edges interconnecting timesteps in time expansion. Directed edges (u^t, v^{t+1}) are introduced for $t = 1, 2, \dots, \mu - 1$ whenever there is $\{u, v\} \in E$. Wait actions are modeled by introducing edges (u^t, u^{t+1}) . A directed path in time expansion corresponds to trajectory of an agent in time. Hence the modeling task now consists in construction of a formula in which satisfying assignments correspond to directed paths from $s_0^0(a_i)$ to $s_+^\mu(a_i)$ in the time expansion.

Assume that we have time expansion (V_i, E_i) for agent a_i . Propositional variable $X_v^t(a_i)$ is introduced for every vertex v^t in V_i . The semantics of $X_v^t(a_i)$ is that it is *TRUE* if and only if agent a_i resides in v at time step t . Similarly we introduce $\mathcal{E}_{u,v}^t(a_i)$ for every directed edge (u^t, v^{t+1}) in E_i . Analogically the meaning of $\mathcal{E}_{u,v}^t(a_i)$ is that it is *TRUE* if and only if agent a_i traverses edge $\{u, v\}$ between time steps t and $t + 1$.

Once we have the Boolean decision variables $X_v^t(a_i)$ and $\mathcal{E}_{u,v}^t(a_i)$ we can introduce constraints so that truth value assignments are restricted only to those that correspond to valid solutions of a given MAPF. The added constraints together ensure that $\mathcal{F}(SoC)$ is a *complete propositional model* for given MAPF.

We here illustrate the model by showing few representative constraints. For the detailed list of constraints we refer the reader to (Surynek et al., 2016).

Collisions between agents can be eliminated by the following constraint over $X_v^t(a_i)$ variables for every $v \in V$ and timestep t :

$$\sum_{a_i \in A \mid v^t \in V_i} X_v^t(a_i) \leq 1 \quad (1)$$

Next, there is a constraint stating that if agent a_i appears in vertex u at time step t then it has to leave

through exactly one edge (u^t, v^{t+1}) . This can be established by following constraints:

$$X_u^t(a_i) \Rightarrow \bigvee_{(u^t, v^{t+1}) \in E_i} \mathcal{E}_{u,v}^t(a_i), \quad (2)$$

$$\sum_{v^{t+1} \mid (u^t, v^{t+1}) \in E_i} \mathcal{E}_{u,v}^t(a_i) \leq 1 \quad (3)$$

Other constraints ensure that truth assignments to variables per individual agents form paths. That is if agent a_i enters an edge it must leave the edge at the next time step.

$$\mathcal{E}_{u,v}^t(a_i) \Rightarrow X_v^{t+1}(a_i) \wedge X_u^t(a_i) \quad (4)$$

A common measure how to reduce the number of decision variables derived from the time expansion is the use of *multi-valued decision diagrams* (MDDs) (Andersen et al., 2007; Sharon et al., 2013). The basic observation that holds for MAPF is that an agent can reach vertices in the distance d (distance of a vertex is measured as the length of the shortest path) from the current position of the agent no earlier than in the d -th time step. Analogical observation can be made with respect to the distance from the goal position.

The combination of SAT-based approach and MDD time expansion is the basis for the MDD-SAT algorithm. It is important to note that from our perspective the MDD can be regarded as a representation of the set of candidate paths for given agent. In MDD-SAT, MDDs represent all possible paths that meet the given makespan and sum-of-costs bounds.

1.3 CBS and Lazy Compilation

Conflict-based search (CBS) (Sharon et al., 2015) is a popular algorithm for solving MAPF optimally. From the compilation perspective, the CBS algorithm could be understood as a **lazy** method that tries to solve an underspecified problem and relies on to be lucky to find a correct solution even using this incomplete specification. There is another mechanism that ensures soundness of this lazy approach, the branching scheme. If the CBS algorithm is not lucky, that is, the candidate solution is incorrect in terms of the MAPF rules, then the search branches for each possible refinement of discovered MAPF rule violation and the refinement is added to the problem specification in each branch. Concretely, the MAPF rule violations are conflicts of pairs of agents such as collision of $a_i \in A$ and $a_j \in A$ in v at time step t and the refinements are conflict avoidance constraints for single agents in the form that $a_i \in A$ should avoid v at time step t (for a_j analogously).

The idea of laziness of CBS has been combined with compilation-based approaches in SMT-CBS algorithm. The major difference from the standard CBS is that there is no branching at the high level. The high level SMT-CBS roughly correspond to the main loop of MDD-SAT, that is, the algorithm checks the existence of MAPF solution for the sum of costs SoC_0 , $SoC_0 + 1$, $SoC_0 + 2$ until it gets a positive answer from the SAT solver. SMT-CBS differs from MDD-SAT in using the incomplete Boolean model in which it ignores specific between agents at beginning. Collisions are resolved lazily similarly as in CBS.

While the collision between a_i and a_j in vertex v at time step t in standard CBS is resolved as high-level branching is, in SMT-CBS it is resolved by refinement of $\mathcal{F}(SoC)$ with a new disjunction $\neg\mathcal{X}_v^t(a_i) \vee \neg\mathcal{X}_v^t(a_j)$.

2 HEURISTICS AND SPARSIFICATION

Although lazy compilation in SMT-CBS is an important factor that reduces the size of generated Boolean formulae in SMT-CBS compared to MDD-SAT, the reduction concerns only the conflict avoidance constraints for pairs of agents (binary disjunctions). The number of nodes as well as the number of directed edges in the MDDs is not reduced by laziness and hence the number of variables and constraints modeling the existence of directed paths in MDDs in the target Boolean formula is the same as in MDD-SAT.

It has been suggested in (Surynek, 2021) as part of the Sparse-SMT-CBS algorithm to reduce the number of directed paths being encoded. This is done by using sparse sets of candidate paths for each agent that satisfy given sum-of-costs and makespan bounds. That is, instead of considering all such paths as done in MDDs only a relevant subset of them is taken into account.

Sparse-SMT-CBS uses MDDs to represent the set of candidate paths, to distinguish it from MDD containing all paths it is referred to as Sparse MDD or SMDD in short. Given the maximum number of steps μ and any (sparse) set of paths $\Pi(a_i)$ in the time expansion of G of length at most μ , we are able to construct $SMDD_i$ that represents $\Pi(a_i)$. It is however important to note that SMDD tends to overestimate the set of paths being represented, that is, the super-set of $\Pi(a_i)$ is actually represented.

Sparse-SMT-CBS uses identical sum-of-costs and makespan bounds increasing scheme as SMT-CBS at the high-level. Each iteration at the high-level solves a question whether there exists a solution to the

input MAPF Σ such that it fits in the current sum-of-costs SoC and makespan μ . This question is compiled as a series of Boolean formulae and consulted with the SAT solver at the low-level.

The low-level in Sparse-SMT-CBS is different from SMT-CBS. In both algorithms it tries to find a non-conflicting set of paths satisfying SoC and μ , but the set of candidate paths from which SMT-CBS selects is fixed in advance in MDD, while Sparse-SMT-CBS starts with a minimal set of candidate paths and each time a new conflict is discovered the set of candidate paths is extended to reflect the new conflict. Sparse-SMT-CBS does so by the *OR-path* selection heuristic that finds a path for **every subset** of the set of the of conflicts and represents this path in the next SMDD.

The example of SMDDs is shown in Figure 1. Initially shortest paths are included in $\Pi(a_i)$ connecting agent's starting vertex and its goal, one shortest path per agent. However, the initial choice of paths is poor leading to a collision in v_5 at timestep 2 which is detected by Sparse-SMT-CBS and alternative paths reflecting the conflict are suggested for each agent by the new-Paths function. The new paths are included in SMDDs as shown in the right part of Figure 1. For these new SMDDs, non-conflicting path can be found.

Despite the extension of SMDD with paths suggested by the OR-path selection heuristic yields SMDD that is smaller than the full MDD. The SMDD can still be very large for the growing number of conflicts. Therefore we suggested a new heuristic called *AND-path* selection that adds to SMDD only one path that avoids all conflicts. Taking into account the overestimation of SMDD even one path can add enough freedom of navigation for the agent using paths being represented by the SMDD. If such path does not exist then the algorithm switches to the full MDD mode that represents all paths satisfying given makespan and sum-of-costs bounds.

We integrated the SMDD reasoning and the AND-path selection heuristic into the SMT-CBS framework, designing a new algorithm we called **Heuristic-SMT-CBS**. The pseudo-code of Heuristic-SMT-CBS is shown as Algorithm 1. The encoding used in Heuristic-SMT-CBS follows MDD-SAT with collision avoidance constraints omitted but instead of deriving decision Boolean variables and constraints from MDDs they are derived from SMDDs. That is, a Boolean variable $\mathcal{X}(a_i)_v^t$ is introduced for each node v^t in $SMDD_i$ and a variable $\mathcal{E}(a_i)_{u,v}^t$ is introduced for each edge in $SMDD_i$. Constraints are introduced analogously. Observe that the sparser the SMDDs are the smaller Boolean model is obtained.

The sparse set of candidate paths is used to build

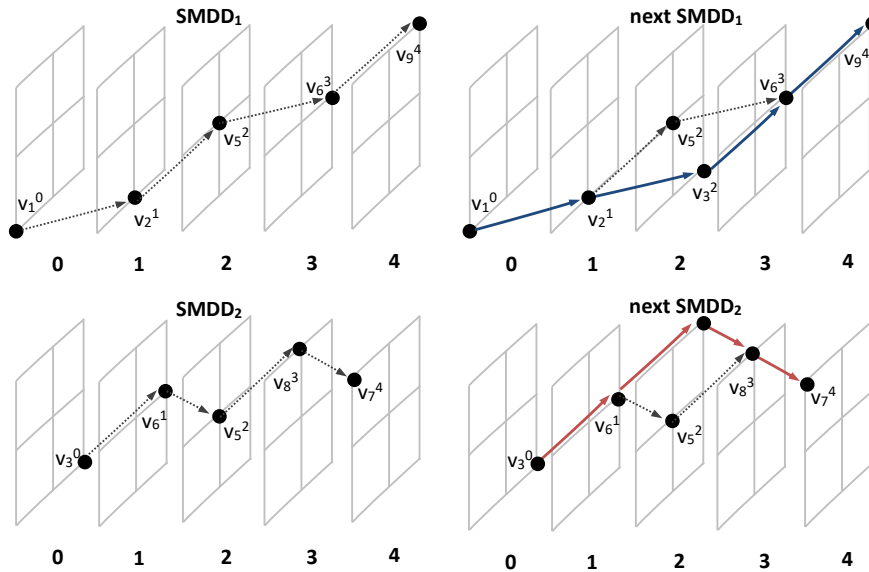


Figure 1: An example of sparse MDDs (SMDDs) for agents a_1 and a_2 . The first iteration yields a conflict in v_5 at timestep 2 between the agents which can be avoided via newly represented paths in the next iteration.

an incomplete Boolean MAPF model $\mathcal{F}(SoC)$ (line 14) which is subsequently solved by the SAT solver. If a satisfying truth-value assignment of $\mathcal{F}(SoC)$ is found (line 16), then it has to be interpreted and checked against MAPF rules, that is, it is checked for collisions between agents (line 18). If there are no collisions, the algorithm can return a valid MAPF solution (line 20). If a collision is detected (lines 21-28), a proper treatment must be applied.

This includes (1) extending the model $\mathcal{F}(SoC)$ with a collision elimination constraint, a disjunction forbidding a pair of conflicting actions to take place at once (lines 21-23) and (2) and extending the sparse set of candidate paths Π with a new path reflecting newly discovered conflicts (line 24-29). It may however happen that no new path could be found after which the algorithm switches to SMT-CBS mode using the full MDD as the set of candidate paths.

Without proof let us state the Heuristic-SMT-CBS algorithm is sound and optimal.

Path included into Π suggested by the AND-path selection can be found by various shortest path algorithms like A* that take into account the set of conflicts.

3 EXPERIMENTAL EVALUATION

We performed an experimental evaluation of Heuristic-SMT-CBS and compared it against SMT-CBS, and Sparse-SMT-CBS on standard benchmarks from *movingai.com* (Boyarski et al., 2015; Sharon et al., 2013; Sturtevant, 2012). To provide broader

perspective outside compilation-based techniques we also include comparison with a vanilla implementation of the CBS algorithm, a search-based algorithm. Representative part of the results is presented in this section.

We implemented Heuristic-SMT-CBS in C++ using the existing implementations of Sparse-SMT-CBS and SMT-CBS. The existing implementations of Sparse-SMT-CBS and SMT-CBS were used in the evaluation as well. All variants of SMT-CBS algorithms are built on top of the Glucose 3.0 SAT solver (Audemard and Simon, 2018)

As for CBS, we used implementation from (Barer et al., 2014) written in C#. Although this implementation of CBS does not use recent MAPF heuristics like rectangle reasoning (Li et al., 2019) we consider it suitable for our experiments as implementations of both Sparse-SMT-CBS and SMT-CBS also do not rely on heuristics derived from the grid.

The SAT solver we used supports incremental mode which means that when the formula is extended with new variables and clauses the solver does not need to start from scratch but it can use learned clauses from its previous runs. This incremental mode is highly utilized across the low-level phase of SMT-CBS, Sparse-SMT-CBS, as well as Heuristic-SMT-CBS whenever the target Boolean formula is extended.

All experiments were run on system consisting of Xeon 2.8 GHz cores, 32 GB RAM, running Ubuntu Linux 18.

The experimental evaluation has been done on diverse instances consisting of 4-connected *grid* maps

Algorithm 1: SMT-based optimal MAPF solver with heuristically guided sparsification of the set of candidate paths.

```

1 Heuristic-SMT-CBS ( $\Sigma = (G = (V, E), R, s_0, s_+)$ )
2    $conflicts \leftarrow \emptyset$ 
3    $\Pi \leftarrow \{\Pi^*(a_i) \mid \text{a shortest path from } s_0(a_i) \text{ to } s_+(a_i) \mid i = 1, 2, \dots, k\}$ 
4    $SoC \leftarrow \sum_{i=1}^k cost(\Pi(a_i))$ 
5    $\mu \leftarrow \max_{i=1}^k cost(\Pi(a_i))$ 
6   while TRUE do
7      $(path, conflicts) \leftarrow \text{Heuristic-SMT-CBS-Fixed}(conflicts, \Pi, \mu, SoC, \Sigma)$ 
8     if  $paths \neq UNSAT$  then
9       return  $paths$ 
10     $SoC \leftarrow SoC + 1$ 
11     $\mu \leftarrow \mu + 1$ 
12 Heuristic-SMT-CBS-Fixed( $\Pi, conflicts, \mu, SoC, \Sigma$ )
13 while TRUE do
14    $\mathcal{F}(SoC) \leftarrow \text{encode-Incomplete}(\Pi, conflicts, \mu, SoC, \Sigma)$ 
15    $assignment \leftarrow \text{consult-SAT-Solver}(\mathcal{F}(SoC))$ 
16   if  $assignment \neq UNSAT$  then
17      $paths \leftarrow \text{extract-Solution}(assignment)$ 
18      $collisions \leftarrow \text{validate}(paths)$ 
19     if  $collisions = \emptyset$  then
20       return  $(paths, conflicts)$ 
21     for each  $(a_i, a_j, v, t) \in collisions$  do
22        $\mathcal{F}(SoC) \leftarrow \mathcal{F}(SoC) \cup \{\neg \mathcal{X}_v^t(a_i) \vee \neg \mathcal{X}_v^t(a_j)\}$ 
23        $conflicts \leftarrow conflicts \cup \{(a_i, v, t), (a_j, v, t)\}$ 
24     for each  $a_i \in R$  do
25        $\pi(a_i) \leftarrow \text{new-AND-Path}(\Pi(a_i), conflicts(a_i), \mu)$ 
26       if  $\pi(a_i) \neq \square$  then
27          $\Pi(a_i) \leftarrow \Pi(a_i) \cup \{\pi(a_i)\}$ 
28       else
29          $\Pi(a_i) \leftarrow \{\text{all paths as in MDD}\}$ 
30   else
31     return  $(UNSAT, conflicts)$ 

```

ranging in sizes from small to relatively large.

We varied the number of agents to obtain instances of various difficulties while initial and goal configurations of agents were generated according to scenarios provided on movingai.com. Depending on the map size, we generated instances with up to 64 (small maps) or 128 (large maps) agents and 25 different instances per number of agents. The time limit in all test was set to 128 seconds. Presented results were obtained from instances solved within this timeout.

Part of the results are presented in Figures 2. We present *success rate* and *sorted runtimes*. Success

rate shows the ratio of instances solved under the time limit of 128 seconds out of 25 instances per number of agents.

The cactus plot for runtime has been generated by taking runtimes of all instances solved under the time limit by the given algorithm and sorting them along x-axis; so the x-th data-point represents the runtime of x-th fastest solved instance by the given algorithm. The faster algorithm hence typically yields to a lower curve in the cactus plot but differences for various regions of the plot can be observed too.

The general trend observable across almost all maps from success rate and sorted runtimes is that Heuristic-SMT-CBS represents an improvement over both SMT-CBS and Sparse-SMT-CBS. The one exception is the maze-128-128-10 map where Heuristic-SMT-CBS lost to Sparse-SMT-CBS but the gap is relatively small and SMT-CBS performs significantly worse than other two variants of SMT-CBS. The behavior of Heuristic-SMT-CBS on the maze-128-128-10 map can be explained by frequent non-existence of path generated by the AND-path selection heuristic.

Sorted runtimes comparison also shows that in Heuristic-SMT-CBS dominates in instances of easy and medium difficulty. However as the difficulty of instances grows often Heuristic-SMT-CBS become worse than Sparse-SMT-CBS. This behavior can be explained by significant reduction of the size of SMDD when the number of collisions is not high. In such a case Heuristic-SMT-CBS constructs only small SMDD which decreases overall solving runtime. On the other hand when there are many collisions between agents, the algorithm switches to full MDD, that is to the identical mode as SMT-CBS and all previous growing of SMDD becomes an overhead.

Comparison with CBS shows an expected trend that due to its negligible overhead, the CBS algorithm dominates in easy instances where SMT-CBS, Sparse-SMT-CBS, and Heuristic-SMT-CBS need to deal with complex instance compilation process that eventually does not pay off. As difficulty of instances grow, the advanced learning mechanisms and intelligent space search pruning implemented in SAT solvers has greater impact resulting in better performance of compilation-based solvers.

The known disadvantage of SMT-CBS is that in large maps like warehouse-10-20-10-2-1 or lak303d it can generate a formulae that are too large which causes degradation in performance. According to our experiments it seems that using sparse set of candidate paths in SMT-CBS helps the solver especially on large maps - there is a relatively big gap between SMT-CBS and its variants with sparsification.

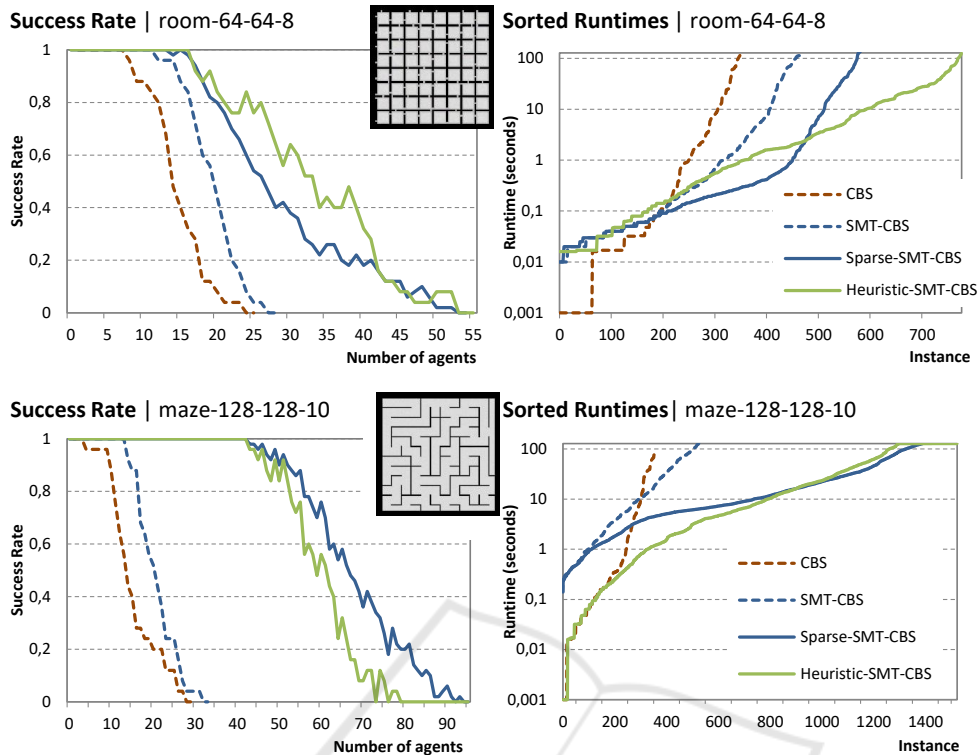


Figure 2: Success rate and runtime comparison on medium-sized maps.

The reason is that the size of MDDs for large maps is often prohibitive and causes significant difficulty for the SAT solver as it needs to deal with a large formula derived from MDD. Sparsification enables choosing few promising paths that are represented in SMDDs which improves both the Boolean formula construction process as well as its solving. Yet fewer paths represented in SMDDs generated by the AND-path selection heuristic leads to an improvement over Sparse-SMT-CBS.

4 CONCLUSION

We suggested a new heuristically guided selection of candidate paths for agents in sparsification schemes for SAT-based approach to multi-agent path planning (MAPF). The technique aims on currently the most significant drawback of compilation-based approaches to MAPF which is the performance on large graphs.

The original sparsification technique generates large sets of candidate paths. We mitigated this drawback in this work by suggesting a heuristic called *AND-path selection*. AND-path selection searches for the most constrained path with respect to discov-

ered collisions between agents. This path when represented multi-value decision diagram (MDD) yields small set of candidate path that still provides enough movement freedom for the agent.

According to our experiments on a number of benchmarks, the new algorithm called Heuristic-SMT-CBS that represents an integration of the AND-path selection into SMT-CBS, significantly outperform Sparse-SMT-CBS, a variant of SMT-CBS with previous sparsification technique.

ACKNOWLEDGEMENTS

This research has been supported by GAČR - the Czech Science Foundation, grant registration number 22-31346S.

REFERENCES

- Andersen, H. R., Hadzic, T., Hooker, J. N., and Tiedemann, P. (2007). A constraint store based on multivalued decision diagrams. In *Principles and Practice of Constraint Programming - CP 2007, Proceedings*, volume 4741 of *Lecture Notes in Computer Science*, pages 118–132. Springer.

- Audemard, G. and Simon, L. (2018). On the glucose SAT solver. *Int. J. Artif. Intell. Tools*, 27(1):1840001:1–1840001:25.
- Barer, M., Sharon, G., Stern, R., and Felner, A. (2014). Sub-optimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Proceedings of the Seventh Annual Symposium on Combinatorial Search, SOCS 2014*. AAAI Press.
- Barrett, C. W. and Tinelli, C. (2018). Satisfiability modulo theories. In Clarke, E. M., Henzinger, T. A., Veith, H., and Bloem, R., editors, *Handbook of Model Checking*, pages 305–343. Springer.
- Biere, A., Heule, M., van Maaren, H., and Walsh, T., editors (2009). *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- Boyarski, E., Felner, A., Stern, R., Sharon, G., Tolpin, D., Betzalel, O., and Shimony, S. E. (2015). ICBS: improved conflict-based search algorithm for multi-agent pathfinding. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, pages 740–746. AAAI Press.
- Dechter, R. (2003). *Constraint processing*. Elsevier Morgan Kaufmann.
- Jünger, M., Liebling, T. M., Naddef, D., Nemhauser, G. L., Pulleyblank, W. R., Reinelt, G., Rinaldi, G., and Wolsey, L. A., editors (2010). *50 Years of Integer Programming 1958-2008 - From the Early Years to the State-of-the-Art*. Springer.
- Li, J., Felner, A., Boyarski, E., Ma, H., and Koenig, S. (2019). Improved heuristics for multi-agent path finding with conflict-based search. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pages 442–449. ijcai.org.
- Lifschitz, V. (2019). *Answer Set Programming*. Springer.
- Luna, R. and Bekris, K. E. (2011). Push and swap: Fast cooperative path-finding with completeness guarantees. In Walsh, T., editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 294–300. IJCAI/AAAI.
- Ma, H., Tovey, C. A., Sharon, G., Kumar, T. K. S., and Koenig, S. (2016). Multi-agent path finding with payload transfers and the package-exchange robot-routing problem. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, 2016*, pages 3166–3173. AAAI Press.
- Rader, D. (2010). *Deterministic Operations Research: Models and Methods in Linear Optimization*. Wiley.
- Ryan, M. R. K. (2007). Graph decomposition for efficient multi-robot path planning. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2003–2008.
- Sharon, G., Stern, R., Felner, A., and Sturtevant, N. R. (2015). Conflict-based search for optimal multi-agent pathfinding. *Artif. Intell.*, 219:40–66.
- Sharon, G., Stern, R., Goldenberg, M., and Felner, A. (2013). The increasing cost tree search for optimal multi-agent pathfinding. *Artif. Intell.*, 195:470–495.
- Silver, D. (2005). Cooperative pathfinding. In *Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference*, pages 117–122. AAAI Press.
- Standley, T. S. (2010). Finding optimal solutions to cooperative pathfinding problems. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010*. AAAI Press.
- Sturtevant, N. R. (2012). Benchmarks for grid-based pathfinding. *IEEE Trans. Comput. Intell. AI Games*, 4(2):144–148.
- Surynek, P. (2010). An optimization variant of multi-robot path planning is intractable. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010*. AAAI Press.
- Surynek, P. (2017). Time-expanded graph-based propositional encodings for makespan-optimal solving of cooperative path finding problems. *Ann. Math. Artif. Intell.*, 81(3-4):329–375.
- Surynek, P. (2019). Unifying search-based and compilation-based approaches to multi-agent path finding through satisfiability modulo theories. In Kraus, S., editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pages 1177–1183. ijcai.org.
- Surynek, P. (2021). Sparsification for fast optimal multi-robot path planning in lazy compilation schemes. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, Prague, Czech Republic*, pages 7931–7938. IEEE.
- Surynek, P. (2022). Heuristically guided compilation for multi-agent path finding.
- Surynek, P., Felner, A., Stern, R., and Boyarski, E. (2016). Efficient SAT approach to multi-agent path finding under the sum of costs objective. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 810–818. IOS Press.
- Yu, J. and LaValle, S. M. (2013a). Planning optimal paths for multiple robots on graphs. In *2013 IEEE International Conference on Robotics and Automation, 2013*, pages 3612–3617. IEEE.
- Yu, J. and LaValle, S. M. (2013b). Structure and intractability of optimal multi-robot path planning on graphs. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, 2013*. AAAI Press.