



X*: Anytime Multi-Agent Path Finding for Sparse Domains using Window-Based Iterative Repairs



Kyle Vedder^{a,*}, Joydeep Biswas^b

^a University of Pennsylvania, Department of Computer and Information Science, 220 South 33rd Street, Philadelphia, PA 19104, United States of America

^b University of Texas Austin, Department of Computer Science, 2317 Speedway, Austin, TX 78712, United States of America

ARTICLE INFO

Article history:

Received 10 August 2019

Received in revised form 27 October 2020

Accepted 4 November 2020

Available online 17 November 2020

Keywords:

Multiagent systems

Motion and path planning

Multi-agent path finding

Anytime path finding

ABSTRACT

Real-world multi-agent systems such as warehouse robots operate under significant time constraints – in such settings, rather than spending significant amounts of time solving for *optimal* paths, it is instead preferable to find valid, collision-free paths quickly, even if suboptimal, and given additional time, to iteratively refine such paths to improve their cost. In such domains, we observe that agent-agent collisions are *sparse* – they involve small local subsets of agents, and are geographically contained within a small region of the overall space. Leveraging this insight, we can first plan paths for each agent individually, and in the cases of collisions between agents, perform small local repairs limited to local subspace *windows*. As time permits, these windows can be successively grown and the repairs within them refined, thereby improving the path quality, and eventually converging to the global joint optimal solution. Using these insights, we present two algorithmic contributions: 1) the Windowed Anytime Multiagent Planning Framework (WAMPF) for a class of anytime planners that quickly generate valid paths with suboptimality estimates and generate optimal paths given sufficient time, and 2) X*, an efficient WAMPF-based planner. X* is able to efficiently find successive valid solutions by employing re-use techniques during the repair growth step of WAMPF. Experimentally, we demonstrate that in sparse domains: 1) X* outperforms state-of-the-art anytime or optimal MAPF solvers in time to valid path, 2) X* is competitive with state-of-the-art anytime or optimal MAPF solvers in time to optimal path, 3) X* quickly converges to very tight suboptimality bounds, and 4) X* is competitive with state-of-the-art suboptimal MAPF solvers in time to valid path for small numbers of agents while providing much higher quality paths.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Multi-Agent Path Finding (MAPF) is the problem of finding a collision free, minimal cost global path π in the joint space of the set of agents α traveling from a set of start states s to a set of goal states g on a graph, often with one or more graph edges blocked at runtime [1]. The path cost, denoted $\|\pi\|$, is often defined as the makespan of π (i.e. the maximum cost for any agent) or the sum of costs for each agent; in this work we focus on optimizing for sum of costs, but this choice is not fundamental. Much of the prior art in MAPF focuses on finding optimal or bounded suboptimal global paths for large

* Corresponding author.

E-mail addresses: kvedder@seas.upenn.edu (K. Vedder), joydeepb@cs.utexas.edu (J. Biswas).

numbers of densely packed agents, often with a focus on how planners scale with an increasing number of agents [1–6]; however, there are many real-world multi-agent scenarios that have sparse agent distributions, are highly dynamic, and require valid paths in milliseconds such as warehouse robots [7], robot soccer [8–11], or drone swarms [12,13]. In such scenarios, finding a optimal global path is too time consuming; instead, it is desirable to employ an anytime solver that can quickly find a collision-free global path of reasonable quality and, if given additional time, improve the global path quality, ultimately converging to an optimal global path.

In this work we focus on the problem of producing an anytime planner which, in sparse domains, quickly finds a valid global path of reasonable quality and, if given sufficient time, will converge to an optimal global path. As part of this work, we leverage three key insights. 1) Unlike in domains like 8-puzzle [14] with each tile treated as an independent agent, in sparse domains, problem instances often have agent-agent collisions for individually planned global paths that involve only a small subset of the total agents and are isolated to a small area easily separable from other collisions. By exploiting sparsity, the MAPF problem can be decomposed into small subspaces, (i.e. small subsets of states and agents) and each subspace efficiently searched to produce a *repair* to the collision (i.e. a new, collision-free section of the global path for the colliding agents), thus producing a valid global path. 2) These subspaces can trade repair generation time for repair quality by varying their size; growing the area of a subspace will produce a repair of the global path of the same or better quality (i.e. lower contribution to the global path cost), but takes more time to produce a repair. 3) Iteratively growing the subspace and generating repairs monotonically improves the global path quality. When a repair search proceeds *unimpeded*, i.e. unrestricted by the constraints the subspace imposes on the full space, from the global start to the global goal of the agents involved, the global path is known to be optimal for those agents.

By combining these key insights, we present an anytime MAPF framework called Windowed Anytime Multiagent Planning Framework (WAMPF), along with an efficient WAMPF-based planner called Expanding A* (X^*) that performs search reuse for efficient iterative path repair. Experimentally, we demonstrate that in sparse domains:

1. X^* outperforms state-of-the-art anytime or optimal MAPF solvers in time to valid path.
2. X^* is competitive with state-of-the-art anytime or optimal MAPF solvers in time to optimal path.
3. X^* quickly converges to very tight suboptimality bounds.
4. X^* is competitive with state-of-the-art suboptimal MAPF solvers in time to valid path for small numbers of agents while providing much higher quality paths.

An earlier version of this work presented a similar version of WAMPF, the naïve WAMPF implementation, and X^* [15], but this work provides refined pseudocode, more detailed explanations, walked through examples, and a completely new experimental results section.

The rest of this paper proceeds as follows: We first introduce relevant background (Section 2) and provide an overview of related MAPF solvers (Section 3). We then present WAMPF, our MAPF solving framework, along with a naïve implementation and two worked out examples (Section 4). We then present X^* , an efficient WAMPF-based planner that performs search reuse for efficient successive path repair (Section 5). Finally, we present several experiments to characterize X^* and compare it to prior art in sparse domains (Section 6), and then discuss directions for future work (Section 7).

2. Background

To put our contributions in the context of the state-of-the-art, we begin by discussing the complexity of Single-Agent Path Finding along with the variety of solution approaches seen in the literature (Section 2.1). We then discuss the complexity of Multi-Agent Path Finding, comparing it to the single agent version, along with the variety of solution approaches seen in the literature (Section 2.2). We then discuss the breadth of both SAPF and MAPF prior art that employ three techniques which are relevant to our contributions, namely Bounded Search (Section 2.3), Search Reuse (Section 2.4), and Anytime Path Planning (Section 2.5). This presentation will prepare the reader for Section 3 where we analyze several MAPF solvers that utilize these techniques.

2.1. Single-agent path finding

Constructing a minimal cost, collision free path from a known start state to a known goal state for a single agent in the face of obstacles and under time constraints is a problem faced in many domains, from robotics to videogame agents. This problem, known as the Single-Agent Path Finding problem (SAPF), appears in domains with both discrete and continuous state spaces.

In discrete spaces, the problem can be modeled in a variety of ways, including integer linear programming [16,17], satisfiability [18], and answer set programming [19]; however, solutions most commonly model the problem as a graph with vertices that represent a state in the state space and with edges that represent the valid transitions between these states. Graph search algorithms are then used to find minimal cost paths between the start vertex and the goal vertex on the graph, and the resulting path can be mapped to a minimal cost set of transitions from the start state to the goal state. These graph search algorithms can be *uninformed*, meaning they know nothing about the problem beyond the given graph

(e.g. Uniform Cost Search [14]) or they can be *informed*, meaning they have additional information about the graph such as a heuristic, e.g. A* [20], or regular problem structure, e.g. Jump Point Search [21].

In continuous spaces, the most computationally challenging problems are intractable; for linked polyhedra moving through three-dimensional space with a fixed set of polyhedral obstacles, commonly known as the Moving Sofa problem or the Couch Mover's problem, finding an optimal, collision free path is PSPACE hard [22]. A common way to simplify continuous problems is to convert them to discrete problems [23,24]; this is often done by imposing a grid-structure, such as a four-connected grid or an eight-connected grid [7,25], or by randomly sampling the space [26]. Imposing a grid adds additional structure to the problem that can be exploited to speed search [21], but environments can be adversarially designed to admit no collision free path along a given grid, but admit many collision free paths in the continuous space version of the problem. To address this problem, the search space can be sampled online, ensuring probabilistic completeness [27]. Two common ways this is done is via initial random graph construction and then search [28] or via construction of the data structure during search [29]; the latter approach enables planning to be joined with perception, thereby dramatically lowering their overall computational cost [30].

2.2. Multi-agent path finding

The problem of finding collision-free paths for *multiple* agents that also avoid agent-agent collisions, known as the Multi-Agent Path Finding problem (MAPF), presents another layer of difficulty. Not only is the continuous, two dimensional case of path finding for multiple rectangles, a simplification of the Couch Mover's problem setup, PSPACE hard [31], the discrete MAPF problem is also significantly more challenging than the discrete SAPF problem. In general, planning jointly for all agents requires planning in a state space with the dimensionality that is at least linear in the number of agents, meaning the cardinality of the state space is at least exponential in the number of agents. Under common conditions, SAPF operates on a *polynomial* domain, i.e. the difficulty of the problem grows polynomially relative to the depth of the optimal solution due to duplicate detection; under these same conditions, MAPF operates on an *exponential* domain, i.e. the difficulty of the problem grows exponentially in the depth of the solution [32,33]. Similar to SAPF, discrete MAPF problems can be modeled via integer linear programming [34], satisfiability [35–37], and answer set programming [38], but many solutions operate directly on graphs [2,3,5].

2.3. Bounded search

Bounded Search is a technique where artificial limits are placed on the search space. While bounds usually produce a suboptimal solution, they prevent planning far into the future on a model of the world that is less likely to be accurate, thereby speeding solution generation. This bound can be enforced via time such as with a time-bounded lattice [39], via search depth such as Hierarchical Cooperative A* [2], or via restricted cost propagation such as Truncated D* Lite [40].

2.4. Search reuse

Search Reuse is a technique where information from previous searches is used to speed up future searches. One of the most widely used families of reuse algorithms, D* [41] / D* Lite [42] and their variants [40,43,44], operates by propagating changes in the environment up the search tree, only modifying states g -values as needed. Other examples of algorithms that employ reuse are from the predator-prey domain, where the predator prunes the search tree of a prior search to make it suitable for the current search, thus saving the cost of re-expanding the remaining states in the pruned tree [45–47].

2.5. Anytime path planners

Anytime Path Planners are planners that can quickly develop a solution to the given problem and, if given more computation time, iteratively improve the plan quality. Anytime algorithms are desirable for many domains as they allow for metareasoning to make online tradeoffs between solution quality and planning time [48–50]. A naïve way to construct an anytime planner is to run a standard planner with parameters which trade solution optimality for a runtime improvement (e.g. A* heuristic inflation [14]), and then iteratively re-run the planner with tighter bounds if computation time remains [51]. While this first plan generation is often fast, successive iterations grow increasingly slow due to lack of information reuse. Anytime planners that instead reuse information from prior searches are typically faster at generating successive plans [52–54].

There exist other, non A*-like anytime path planners that also leverage reuse techniques, such as RRT* [29], which finds a feasible solution and then, given more time, repeatedly improves it by further sampling the space and updating the tree with cheaper intermediate nodes when applicable, converging to the optimal solution in the limit. Reuse and bounded search techniques can also be combined to further speed anytime search [55,56].

3. MAPF related work

In this work we focus on MAPF solving for general graphs. In principle, any uninformed weighted graph search algorithm such as Uniform Cost Search (UCS) [14] is sufficient to find an optimal path for any MAPF problem by treating each joint

state as a position of a single high dimensional meta-agent; however, providing additional information such as a heuristic, framing the problem differently, or exploiting additional properties often present in relevant domains can produce more efficient algorithms, provide different runtime characteristics, or provide different guarantees, thus motivating the variety of MAPF solvers.

MAPF solvers fall into two major classes: global search and decoupled search. Like UCS, global search techniques solve a single large meta-agent search problem; however, these techniques attempt to leverage problem substructure to speed search [3,57–60]. Decoupled search approaches decompose the problem by planning for each agent serially, forcing later agents to account for sections or the entirety of earlier agents plans [2,4,6,61–66]. In order to discuss our approach in the context of prior art, we present a unified notation as follows: every state s is in the joint space of the agent set α which contains one or more possibly heterogeneous agents. In order to refer to the part of s associated with a subset of its agents, we introduce a state filter function $\Phi(s, \alpha')$, where $\alpha' \subseteq \alpha$. For example, if s 's agent set $\alpha = \{a, b, c\}$ and we want to refer to the part of s associated with agents b and c , this is denoted by $\Phi(s, \{b, c\})$. This notation allows us to reason about the subspaces that we introduce shortly. Importantly, in our notation states do not contain time bookkeeping; while time is relevant for collision checking, the bookkeeping for collision checking is well understood [1] and abstracted away by the state neighbor function $N(s)$, so we omit it for simplicity.

M^* [3] is a state-of-the-art global MAPF solver that exploits domain sparsity in order to speed its search. M^* operates by first computing an optimal individual space *policy* to $\Phi(g, \{a\})$ for all $a \in \alpha$. It then traces a path in the space of α from s to g using the policies of each agent. If a collision is encountered, M^* is able to use the policy information to compute the relevant $\alpha' \subseteq \alpha$ to involve in a joint search. In sparse domains, the number of agents involved in this joint search is small, allowing M^* to avoid the aforementioned combinatorial explosion, and collisions are typically separate from one another, avoiding the need to merge joint searches. Due to the expensive nature of the policy computation for each agent, even if lazily computed with approaches like Reverse Resumable A^* [2], M^* is ill-suited to the task of quickly generating a valid solution in sparse domains. Furthermore, while M^* can produce optimal and ϵ -suboptimal paths, it is not anytime nor does its ϵ -suboptimal version allow for efficient path refinement if given additional time.

Conflict-Based Search (CBS) [4] is a state-of-the-art decoupled MAPF solver that exploits domain sparsity to speed search. CBS first computes an optimal path from $\Phi(s, \{a\})$ to $\Phi(g, \{a\})$ for all $a \in \alpha$; if a collision occurs between agents i and j , CBS forms two models of the world, one where the path of i is constrained through the collision point and the path of j is replanned, and one where the path of j is constrained through the collision point and the path of i is replanned. This approach is then applied recursively to each model, forming a conflict tree. In sparse domains, the number of agents involved in a collision is often small, therefore producing a small conflict tree. A characteristic of CBS is it sometimes struggles with open areas; when there are many short paths that collide and a longer path needs to be employed, the conflict tree grows very large before the optimal solution is considered. Furthermore, while CBS can produce optimal paths and its extended counterpart ECBS can produce ϵ -suboptimal paths [66], neither are anytime nor does ECBS allow for efficient path refinement of ϵ -suboptimal paths if given additional time.

Anytime Focal Search (AFS) [5] is a state-of-the-art global MAPF solver that exploits the availability of “good enough” solutions in order to quickly find a valid solution and improves this path if given more time. AFS maintains open set O and closed set C structures similar to A^* and an additional structure *focal list* of states that have *f*-values of no more than ϵ times larger than the smallest value in O . Rather than constraining itself to only expand minimal cost states, AFS is willing to expand other states in the focal list, determined via a *priority function*, thereby allowing it to quickly find a path to g that is ϵ -suboptimal. Given more time, the bookkeeping done in the focal list allows AFS to tighten ϵ and improve its path without searching from scratch, ultimately producing an optimal solution. As AFS is anytime, it is able to provide intermediate results along with a confidence bound. AFS does not attempt to decompose the problem as it always plans in the full joint space of α from s to g , leading to higher valid solution runtimes compared to planners that exploit sparsity.

Push and Rotate (PR) [6] is a state-of-the-art decoupled MAPF solver. Unlike the other solvers presented, PR does not attempt to find an optimal or bounded suboptimal solution; instead, it uses graph transformations (Push and Rotate) to quickly find a *valid* solution, allowing it to scale to large numbers of agents with highly dense agent distributions. As PR is not an optimal or bounded suboptimal solver, it provides no guarantees of path quality; in our experimentation, PR commonly generated paths of cost 2x greater than optimal paths. Due to the high cost of the generated paths and an inability to refine them, PR is ill-suited for domains that require a high quality path.

Expanding A^* (X^*), which we introduce, combines many of the strengths of these algorithms. Like CBS, X^* first computes an optimal path from $\Phi(s, \{a\})$ to $\Phi(g, \{a\})$ for all $a \in \alpha$. Like M^* , when a collision is detected, it performs joint search only in a subspace, but without the need to compute individual policies and in a much smaller subspace. Like AFS, X^* is able to produce intermediate solutions while also exploiting domain sparsity. Like PR, X^* is able to quickly generate a valid solution in sparse domains but with tighter quality bounds.

There exists a number of extensions to CBS and M^* that either utilize optimizations to underlying solvers that are orthogonal to the approach itself or exploit regular domain structure when available. Examples of orthogonal optimizations include Operator Decomposition (OD) [58], which operates by first considering neighbors that only change the path of one agent; these approaches are applicable to any A^* -based solver, including X^* . Examples of optimizations that exploit domain structure to speed search include Enhanced Partial Expansion A^* (EPEA *) [60], which exploits domain structure to only generate a subset of neighbors at a specific *f*-value and Prioritize Conflicts in Improved CBS (ICBS) [67], which relies upon avoiding alternate paths of the same cost for one or both agents involved in pair-wise collisions, as typically found in

structured domains, in order to reduce conflict tree size; however, the approaches of X*, AFS, CBS, and M* do not exploit domain structure in this way.

4. Windowed anytime multiagent planning framework

As discussed in Section 1, the size of the joint state space grows exponentially in the number of agents; this motivates subspace-based approaches such as M* that speed up search by decomposing the full MAPF problem into smaller subproblems involving fewer agents. A key insight is that while subspaces can be used to limit the search to a subset of *agents*, they can also be used to limit the search to a subset of *states*.

We present a construct called a *window* that encapsulates a subset of agents and a connected subset of states. A window is placed around a collision in the global path in order to produce a *repair* to the global path by performing a search within the window. The start of the repair search in w_k , denoted s_k , is the first state on the global path in the window and the goal of the repair search in w_k , denoted g_k , is the last state on the global path in the window. Every window w_k has a successor window w_{k+1} that shares the same agent set but has a superset of states. This allows for the concept of iteratively growing a window by replacing it with its successor that considers more of the domain in its repair. Two windows can be merged together to form a larger window that incorporates both smaller windows via the \cup operator. For example, w and w' can be joined together to form a larger window $w'' := w \cup w'$; w'' must have an agent set $\alpha'' = \alpha \cup \alpha'$ and all of the states in w and w' must be part of the joint states of w'' . Finally, two windows can be checked for overlap via the \cap operator. For example, $w \cap w'$ is true if and only if their agent sets α and α' overlap and they share one or more individual agent states. These window definitions and mechanics are demonstrated in Section 4.4.

While a window-based repair does not ensure the resulting repaired global path is optimal, a repair in a successor window w_{k+1} ensures that its repaired global path will be *at most* the same cost as the global path repaired by w_k and often cost less. Thus, repeatedly growing the subspace and generating repairs monotonically improves the global path quality. Furthermore, if a window w_k is sufficiently large that s_k and g_k are the global start $\Phi(s, \alpha)$ and goal $\Phi(g, \alpha)$ for its agents α and w_k does not *impede the search* from s_k to g_k , i.e. limit search exploration with w_k state restrictions, then the joint paths for the agents α in w_k are jointly optimal and w_k can be discarded. If no more windows exist, then the joint path is an optimal solution. Using this insight, we introduce an anytime MAPF framework called the Windowed Anytime Multiagent Planning Framework (WAMPF).

4.1. WAMPF overview

We present the pseudocode for WAMPF in Algorithm 1 featuring the eponymous top level procedure, the recursive procedure `RecWAMPF` which does the heavy lifting, and the overlapping window helper `PLANINOVERLAPWINDOWS`. The WAMPF pseudocode only manages the state of search windows; all searches are conducted by the implementation defined components `PLANIN` and `GROWANDREPLANIN`, discussed in Section 4.2, in order to make WAMPF domain-agnostic.

WAMPF operates by initially forming a potentially colliding global path by planning for each agent in individual space. `RecWAMPF` is then invoked, and this recursive procedure makes tail-recursive calls until the global path is provably optimal, each time improving the quality of the global path. `RecWAMPF` operates by first growing and replanning in all existing windows, merging them with existing windows if they overlap (Lines 6 – 12), then creating new windows to encapsulate any remaining collisions, merging them with existing windows if they overlap (Lines 13 – 15). At this point, no more collisions exist in the global path and thus the global path is valid. `RecWAMPF` then removes any window searches which have optimally repaired the global path (Lines 16 – 17); if no more windows exist, then the global path is proven optimal (Line 18) and `RecWAMPF` terminates. Otherwise, the current valid global path is reported as an intermediary solution along with its optimality bound estimate. This bound is computed via the current global path cost, an exact or over-estimate of the optimal global path cost, divided by the individual space planned global path cost, an exact estimate or an under-estimate of the optimal global path cost (Line 19). `RecWAMPF` then recursively invokes itself for another iteration.

One of the important features of WAMPF is it repairs collisions chronologically, thus ensuring that each window added and repaired is making progress towards a valid path. Newly added and repaired windows can potentially change the relative timing of agents *later* along the path, inadvertently fixing later collisions or introducing new ones; however, these repairs cannot cause changes *earlier* along the path, only later. By sweeping from the beginning to the end of the path, WAMPF ensures that once a window is added its repair work cannot be undone by other repairs and any collisions induced by a repair must be later along the path and thus handled by WAMPF.

Another important feature of WAMPF is it avoids invalidating repair windows during valid path improvement. It is possible that an earlier window can be grown and replanned in, producing a new repair of higher quality that changes the relative time that agents enter a later window; this change would invalidate the start state of the later window, forcing its repair efforts to be discarded. As such, repair searches and improvements are responsible for not invalidating any windows that exist later along the path; this can be implemented via padding as discussed in Section 4.3 and illustrated in Section 4.4.3.

Together, these two features ensure that WAMPF's running time for a valid path is a function of the number of agent-agent collisions and their separability from other collisions, i.e. domain sparsity, and running time for successive repairs is a function of the number of windows and the number of agents involved in each window.

Algorithm 1 Windowed Anytime Multiagent Planning Framework.

```

1: procedure WAMPF
2:    $\pi \leftarrow$  joint plan comprised of optimal paths planned in individual space
3:    $W \leftarrow \emptyset$ 
4:   return RecWAMPF( $\pi, W, \|\pi\|$ )
5: procedure RecWAMPF( $\pi, W, c$ )
6:   for all  $w_k \in W$  do
7:     if  $\exists w' \in W : w' \neq w_k \wedge w' \cap w_{k+1} \neq \emptyset$  then
8:        $W \leftarrow W \setminus \{w_k\}$ 
9:        $W, \pi \leftarrow$  PLANINOVERLAPWINDOWS( $w_{k+1}, W, \pi$ )
10:      continue
11:       $w_{k+1}, \pi \leftarrow$  GROWANDREPLANIN( $w_k, \pi$ )
12:       $W \leftarrow (W \setminus \{w_k\}) \cup \{w_{k+1}\}$ 
13:   while FIRSTCOLLISIONWINDOW( $\pi) \neq \emptyset$  do
14:      $w \leftarrow$  FIRSTCOLLISIONWINDOW( $\pi$ )
15:      $W, \pi \leftarrow$  PLANINOVERLAPWINDOWS( $w, W, \pi$ )
16:   for all  $w \in W$  do
17:     if SHOULDQUIT( $\pi, w$ ) then  $W \leftarrow W \setminus \{w\}$ 
18:   if  $W = \emptyset$  then return ( $\pi, 1$ )
19:   report  $(\pi, \frac{\|\pi\|}{c})$ 
20:   return RecWAMPF( $\pi, W, c$ )
21: function PLANINOVERLAPWINDOWS( $w, W, \pi$ )
22:   for all  $w' \in W : w' \cap w \neq \emptyset$  do
23:      $W \leftarrow W \setminus \{w, w'\}$ 
24:      $w \leftarrow w \cup w'$ 
25:    $\pi \leftarrow$  PLANIN( $w, \pi$ )
26:    $W \leftarrow W \cup \{w\}$ 
27:   return ( $W, \pi$ )

```

4.2. WAMPF components

As WAMPF is a domain agnostic framework for anytime MAPF planners, it has several definitions/subroutines which must be provided by any planner implementing it:

Window definition: a window definition is state space specific, but a window w_k must uphold the aforementioned properties, namely:

- Contain a connected subset of states for a subset of agents
- Possess a start s_k and a goal g_k on the global path
- Possess a successor window w_{k+1} which contains a superset of states and the same agent set
- The ability to merge with another window to form a new window encapsulating the agent sets and states contained in w_k and the other window via the \cup operator which returns the new window
- The ability to check for overlap with another window via the \cap operator which returns a boolean

FIRSTCOLLISIONWINDOW(π): given a path π , this subroutine finds the first agent-agent collision along the time dimension, beginning with π_0 . If one or more collisions exist, return a window encapsulating the first collision; otherwise, return \emptyset .

PLANIN(w_k, π): the given path π has an associated agent set α and the given window w_k has an associated agent set α' , where $\alpha' \subseteq \alpha$. This subroutine generates a collision free repair in w_k by planning an optimal path from s_k to g_k , respecting the entry times of agents to s_k . The repair is inserted as a replacement to the relevant subset of π , respecting the relative timings of agents involved in later windows, and π is returned.

GROWANDREPLANIN(w_k, π): the given path π has an associated agent set α and the given window w_k has an associated agent set α' , where $\alpha' \subseteq \alpha$. This subroutine grows w_k by replacing it with its successor, w_{k+1} , and generates a repair in w_{k+1} by planning an optimal path from s_{k+1} to g_{k+1} , and inserting it as a replacement to the relevant subset of π , respecting the relative timings of agents involved in later windows, then returning (w_{k+1}, π) . GROWANDREPLANIN(w_{k+1}, π) is guaranteed to only be invoked when PLANIN(w_{k+1}, π) or GROWANDREPLANIN(w_k, π) have previously been invoked and is guaranteed that w_{k+1} does not overlap with any other existing window.

SHOULDQUIT(π, w_k): this subroutine is a predicate that determines if the given window w_k should be discarded. In order to ensure that WAMPF produces globally optimal solutions, a window w_k with an associated agent set α cannot be discarded until $s_k = \Phi(s, \alpha)$, $g_k = \Phi(g, \alpha)$, and w_k does not impede the repair search.

Assuming a WAMPF-based planner meets these conditions:

Theorem 1. The planner will produce a valid global path after a single iteration of RecWAMPF. See Appendix A, Theorem 1 for proof.

Theorem 2. Given sufficient iterations of RecWAMPF, the planner will produce a optimal global path. See Appendix A, Theorem 2 for proof.

4.3. Naïve Windowing A*

To provide a concrete example of a WAMPF-based planner, we present Naïve Windowing A* (NWA*), a naïve implementation of WAMPF with a window definition specific to unit cost four-connected grids. NWA* employs A* as the underlying window solver and makes no attempt at search re-use when the window is grown. We present the requisite WAMPF definitions/subroutines:

Window definition: The window is formulated as a high dimensional rectangular prism, characterized by its bottom left and upper right corners in the joint space of its agent set. New windows are initialized around a collision state by selecting all states that have an L_∞ distance from the collision state of less than or equal to a hyperparameter. An example of such a window is shown in Fig. 1b, where the window, drawn as a dashed rectangle, is in the joint space of a and b and created via an L_∞ norm of 1. A window is grown by moving its corners further away from the center by a fixed number of steps. An example of window growth is shown in the transition from Fig. 1b to Fig. 1c, where window is grown by increasing the radius by a state. Windows w and w' overlap if $\alpha \cap \alpha' \neq \emptyset$ and their rectangles overlap. An example of non-overlapping windows is shown in Fig. 2c, and an example of overlapping windows is shown in Fig. 2d. Windows w and w' are merged to create w'' by unioning their agent sets and constructing a containing rectangle. An example of a window merge is shown in Fig. 2e, where w^{ab} and w^{ac} merge to form w^{abc} .

FIRSTCOLLISIONWINDOW(π): This subroutine looks for collisions along the global path π , starting with π_0 and ending with state $\pi_{|\pi|-1}$. If a collision is detected, a window is initialized around the colliding state with the colliding agents; otherwise, \emptyset is returned.

PLANIN(w_k, π): the given global path π has an associated agent set α and the given window w_k has an associated agent set α' , where $\alpha' \subseteq \alpha$. s_k and g_k are computed from $\Phi(\pi, \alpha')$; s_k is the first state on $\Phi(\pi, \alpha')$ in w and g_k is the last state on $\Phi(\pi, \alpha')$ in w . An A* search is run in the space of w_k from s_k to g_k , with any expanded state's neighboring states not in w discarded rather than placed in the open set O . The resulting repair π' replaces the section of path in $\Phi(\pi, \alpha')$ from s_k to g_k . Importantly, if π is not already a valid solution, then π 's cost may stay the same or it may increase after π' is inserted; if π is already a valid solution, then π' will be of the same or reduced cost compared to the region of $\Phi(\pi, \alpha')$ from s_k to g_k , as π will have already been repaired by a window w_{k-1} , and so the larger w_k may find a repair π' for the same region of $\Phi(\pi, \alpha')$ that costs less. In the case where π' costs less, it must be padded in order to ensure all agents leave g_k at the same time as they did in prior to the insertion of π' in π ; an example of this is shown in Section 4.4.3. Additionally, if the A* search returns NOPATH, w_k is grown to form w_{k+1} and the result of PLANIN(w_{k+1}, π) is returned.

GROWANDREPLANIN(w_k, π): This subroutine grows w_k by replacing it with its successor, w_{k+1} , and then returning the result of PLANIN(w_{k+1}, π).

SHOULDQUIT(π, w_k): the global path π has an associated agent set α and the window w_k has an associated agent set α' . This subroutine returns true iff s_k and g_k are $\Phi(\pi, \alpha')_0$ and $\Phi(\pi, \alpha')_{|\pi|-1}$, respectively, and w_k did not impede the search during the last invocation of PLANIN(w_k, π), i.e. neighbors were not culled during any of A*'s state expansion due to w_k 's state space constraints.

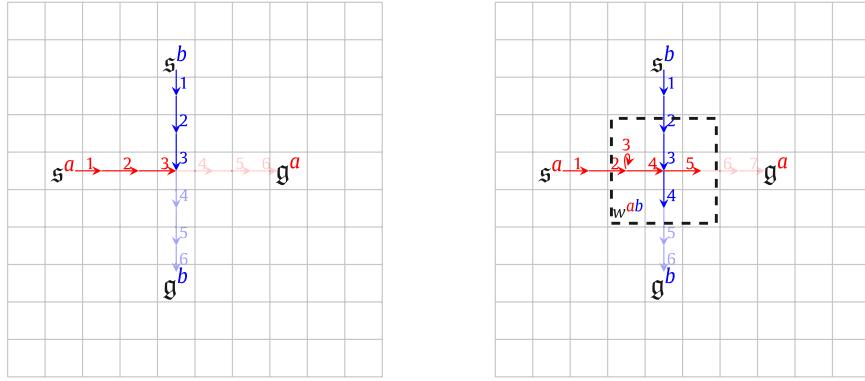
4.4. WAMPF examples

In order to illustrate the behavior of WAMPF (Algorithm 1), we present four worked out examples. The first example (Fig. 1) demonstrates how WAMPF operates for a single collision between two agents using NWA*'s window definition. The second example (Fig. 2) demonstrates how WAMPF operates for multiple collisions using NWA*'s window definition. The third example (Fig. 3) demonstrates how WAMPF can generate valid but suboptimal solutions, and how path insertion and padding operates using NWA*'s window definition. The fourth example (Fig. 4) demonstrates how WAMPF can operate on arbitrary graphs and how NWA*'s window definition can be generalized. All examples are applicable to NWA* as well as our efficient WAMPF-based planner, X* (Section 5), as both planners share the same window definition. The first three examples operate on a 10×10 unit cost four-connected grid and the fourth example operates on a random graph.

4.4.1. Single window example

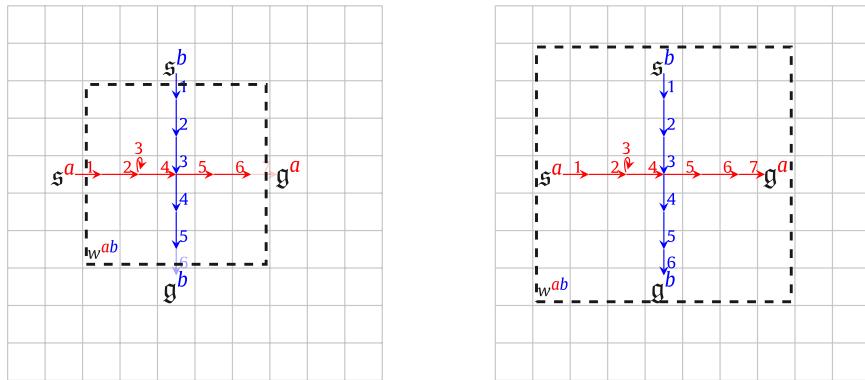
The single window example shown in Fig. 1 demonstrates the mechanics of window creation, window growth and replanning, and window termination using NWA*'s window definition. The example demonstrates a single collision between two agents resolved via a window search; this window is then repeatedly expanded and re-searched until it encompasses an unimpeded search from s to g . The associated figures depict how WAMPF planning for agents individually can induce a collision (Fig. 1a), how a window encapsulates a repair and what a repair looks like for joint plans (Fig. 1b), how a window can be grown to consider a larger search space, therefore potentially improving repair quality (Fig. 1c), and that a window can be terminated after it encapsulates a repair from the start to the goal and does not impede the repair search (Fig. 1d). A key takeaway from this example is that WAMPF windows do not need to encapsulate the entirety of the potentially infinite number of states in the space of their agents in order to terminate.

A line-by-line analysis of Fig. 1 grounded in the WAMPF algorithm (Algorithm 1) is as follows:



(a) Individually planned paths for each agent from s to g are used to form a global path. An agent-agent collision occurs in the path between a and b at $t = 3$.

(b) Collision between a and b is repaired by jointly planning inside w^{ab} . The global path is now guaranteed to be valid, but not guaranteed to be optimal.



(c) w^{ab} is grown and a new repair is generated for a and b . The window does not yet encapsulate the search from s^{ab} and g^{ab} , so the repaired global path is not yet guaranteed to be optimal.

(d) w^{ab} is grown and a new repair is generated. The repair search is from s^{ab} to g^{ab} and unimpeded by w^{ab} , thus allowing w^{ab} to be removed and the global path returned as optimal.

Fig. 1. Single window WAMPF example using NWA*'s window definition.

- 1) Optimal paths are planned for each agent individually to form a global path; the paths for agents *a* and *b* collide at Step 1.
- 2) RECWAMPF invoked. There are no existing windows, so no window manipulations are done.
- 3) The collision between *a* and *b* is detected by FIRSTCOLLISIONWINDOW and w^{ab} is formed to encapsulate it.
- 4) PLANINOVERLAPWINDOWS is invoked to merge w^{ab} with existing windows if needed; however, there are no existing windows (W is empty) so no merging occurs.
- 5) PLANIN is invoked to generate a repair in w^{ab} . w^{ab} is added to the window set W .
- 6) No more collisions exist so FIRSTCOLLISIONWINDOW returns \emptyset and the collision detection loop exits.
- 7) w^{ab} does not allow for an unimpeded search from $\Phi(s, \{a, b\})$ to $\Phi(g, \{a, b\})$, so SHOULDQUIT returns false and W remains unchanged.
- 8) W is not empty so the global path π is not returned as optimal, but it is reported as an intermediary solution along with its optimality bound.
- 9) RECWAMPF is recursively invoked, with $W = \{w^{ab}\}$ and a valid but potentially suboptimal global path.
- 10) w^{ab} is grown and replanned in, producing a larger w^{ab} and a repair. The larger w^{ab} replaces its predecessor in W , and it does not overlap with any other windows so no merging is done.

- ▷ Lines 2 – 3.
Shown in Fig. 1a.
- ▷ Lines 6 – 12.
- ▷ Lines 13 – 14.
- ▷ Lines 22 – 23.
- ▷ Lines 25 – 26.
Shown in Fig. 1b.
- ▷ Line 22.
- ▷ Lines 16 – 17.
- ▷ Lines 18 – 19.
- ▷ Line 25.
- ▷ Lines 11 – 12.

11) No collisions exist and w^{ab} does not allow for an unimpeded search from $\Phi(s, \{a, b\})$ to $\Phi(g, \{a, b\})$, so the updated global path is reported as an intermediary solution and RecWAMPF is recursively invoked.

▷ Lines 13 – 18.
Shown in Fig. 1c.

12) RecWAMPF proceeds, growing w^{ab} and updating its repair and intermediary solutions, with no collisions introduced. The repair in w^{ab} allowed for an unimpeded search from $\Phi(s, \{a, b\})$ to $\Phi(g, \{a, b\})$, therefore allowing SHOULDQUIT to return true. This removes w^{ab} from W , making W empty and thus returns the global path as optimal.

▷ Lines 6 – 18.
Shown in Fig. 1d.

4.4.2. Multi-window example

The example shown in Fig. 2 expands on the mechanics demonstrated in Fig. 1 by demonstrating window merging and subspace planning capabilities using NWA*'s window definition. The example demonstrates a collision between two agents whose repair causes a cascading collision with another agent later along the path. The two repairs are then grown, eventually merging into the joint space of three agents, and eventually terminates after allowing an unimpeded search from s to g . The associated figures depict how WAMPF planning for agents individually can induce a collision, but often only for a subset of agents (Fig. 2a), how a window repair can cause collisions later in the path, creating the need for more windows (Fig. 2b), the creation of a second window, finally generating a collision free solution (Fig. 2c), that grown windows which overlap in the state and agent space need to be merged (Fig. 2d), the resulting merged window (Fig. 2e), and the repeatedly grown window which is finally terminated (Fig. 2f). A key takeaway from this example is WAMPF's window-based approach speeds search; while the given problem involves four agents, WAMPF never required a search in the joint space of more than three agents to produce an optimal path and only required two small searches in the joint space of two agents to produce a valid path.

A line-by-line analysis of Fig. 2 grounded in the WAMPF algorithm (Algorithm 1) is as follows:

1) Plans optimal paths for each agent individually and W is initialized. Note that agents a and b collide at step 1.

▷ Lines 2 – 3.
Shown in Fig. 2a.

2) RecWAMPF invoked. The collision between a and b is detected by FIRSTCOLLISIONWINDOW and w^{ab} is formed to encapsulate it, and there are no windows to collide with.

▷ Lines 6 – 23.

3) PLANIN is invoked to generate a repair in w^{ab} . w^{ab} is added to the window set W .

▷ Lines 25 – 26.
Shown in Fig. 2b.

4) The w^{ab} repair has created a new collision later in time between a and c . On the next iteration of the loop FIRSTCOLLISIONWINDOW detects the collision and w^{ac} is formed to encapsulate it.

▷ Lines 13 – 14.

5) PLANINOVERLAPWINDOWS is invoked to merge w^{ac} with existing windows as needed, but $W = \{w^{ab}\}$ and w^{ab} does not overlap with w^{ac} , so no window merges occur.

▷ Lines 22 – 23.

6) PLANIN is invoked to generate a repair in w^{ac} . w^{ac} is added to the window set W .

▷ Lines 25 – 26.
Shown in Fig. 2c.

7) No more collisions exist so FIRSTCOLLISIONWINDOW returns \emptyset and collision detection loop exits.

▷ Line 22.

8) w^{ab} does not allow for an unimpeded search from $\Phi(s, \{a, b\})$ to $\Phi(g, \{a, b\})$, and w^{ac} does not allow for an unimpeded search from $\Phi(s, \{a, c\})$ to $\Phi(g, \{a, c\})$, so SHOULDQUIT returns false for both windows and W remains unchanged.

▷ Lines 16 – 17.

9) W is not empty so the global path is not returned as optimal, but it is reported as an intermediary solution along with its optimality bound.

▷ Lines 18 – 19.

10) RecWAMPF is recursively invoked, with $W = \{w^{ab}, w^{ac}\}$ and the valid but potentially suboptimal plan.

▷ Line 25.

11) w^{ab} is grown and replanned in, producing a larger w^{ab} and a repair. The larger w^{ab} replaces its predecessor in W , and it does not overlap with w^{ac} so they do not merge.

▷ Lines 11 – 12.

12) w^{ac} is grown, and its successor overlaps with w^{ab} , so w^{ac} is removed from W such that $W = \{w^{ab}\}$, and w^{ac} 's successor is to be merged with w^{ab} .

▷ Lines 7 – 9.
Shown in Fig. 2d.

13) As w^{ab} and w^{ac} overlap, PLANINOVERLAPWINDOWS is invoked. These windows are merged together to form w^{abc} and a repair is generated inside it. w^{abc} is added to W , replacing w^{ab} and w^{ac} such that $W = \{w^{abc}\}$.

▷ Lines 22 – 27.
Shown in Fig. 2e.

14) No collisions exist and w^{abc} does not allow for an unimpeded search from $\Phi(s, \{a, b, c\})$ to $\Phi(g, \{a, b, c\})$, so the updated global path is reported as an intermediary solution and RecWAMPF is recursively invoked.

▷ Lines 13 – 18.

15) RecWAMPF proceeds, growing w^{abc} and updating its repair and intermediary solutions. No collisions are introduced and w^{abc} does not allow for an unimpeded search from $\Phi(s, \{a, b, c\})$ to $\Phi(g, \{a, b, c\})$, therefore allowing SHOULDQUIT to return true. This removes w^{abc} from W , making W empty and thus returns the global path as optimal.

▷ Lines 6 – 20.

16) RecWAMPF proceeds, growing w^{abc} and updating its repair and intermediary solutions, with no collisions introduced. The repair in w^{abc} allowed for an unimpeded search from $\Phi(s, \{a, b, c\})$ to $\Phi(g, \{a, b, c\})$, therefore allowing SHOULDQUIT to return true. This removes w^{abc} from W , making W empty and thus returns the global path as optimal.

▷ Lines 6 – 18.
Shown in Fig. 2f.

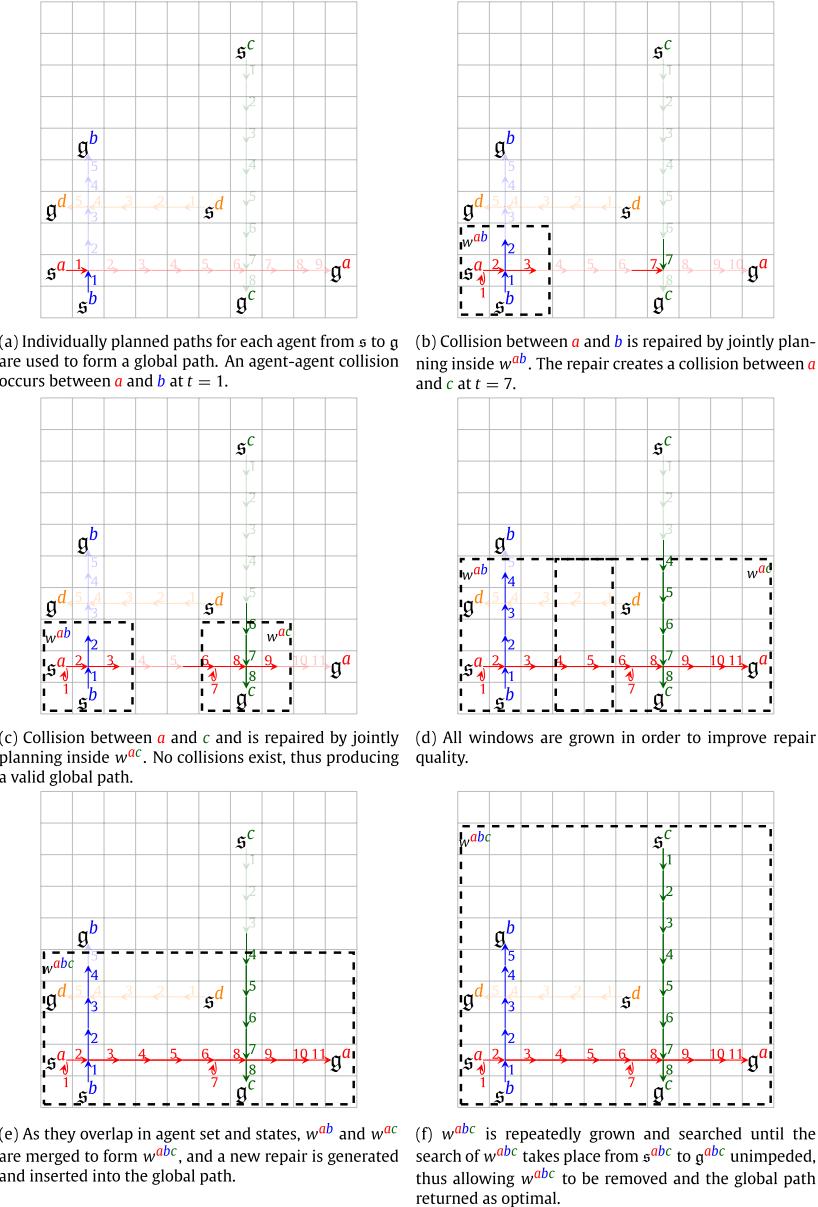


Fig. 2. Multi-window WAMPF example using NWA*’s window definition.

4.4.3. Globally suboptimal repairs and path padding example

The example shown in Fig. 3 demonstrates how WAMPF can produce a globally suboptimal path from an optimal repair within a window, and how higher quality repairs are padded to prevent breaking the entry state of windows further along the path using NWA*’s window definition. The associated figures first demonstrate an initial collision caused by WAMPF planning individually (Fig. 3a). WAMPF then creates a repair window w^{ab} that is searched to find an repair, forcing agent b to step inside the slot in the wall to let a pass, thus adding two more moves to the global path cost. Due to w^{ab} ’s constraints, the repair was unable to consider instead sending a above the upper wall, towards its goal which would produce no increase in global path cost; as such, the repair generated is optimal within w^{ab} but produces a suboptimal global path. This path also causes a collision later along the path (Fig. 3b) that is then repaired with a second window w^{bc} and WAMPF returns a valid solution (Fig. 3c).

Finally, the first window is grown, producing the repair of sending a above the upper wall and allowing b to travel without stepping into the slot; however, this improved repair would cause b to arrive at w^{bc} window too early as compared to its prior plan; in order to prevent this invalidation, the repair to b is padded with two waits to ensure that b leaves w^{ab} at the same time as it did previously in order to ensure it arrives at w^{bc} at the proper time. (Fig. 3d). By performing this padding, WAMPF ensures all agents leave the window at the times they did previously, thus guaranteeing leaving agents

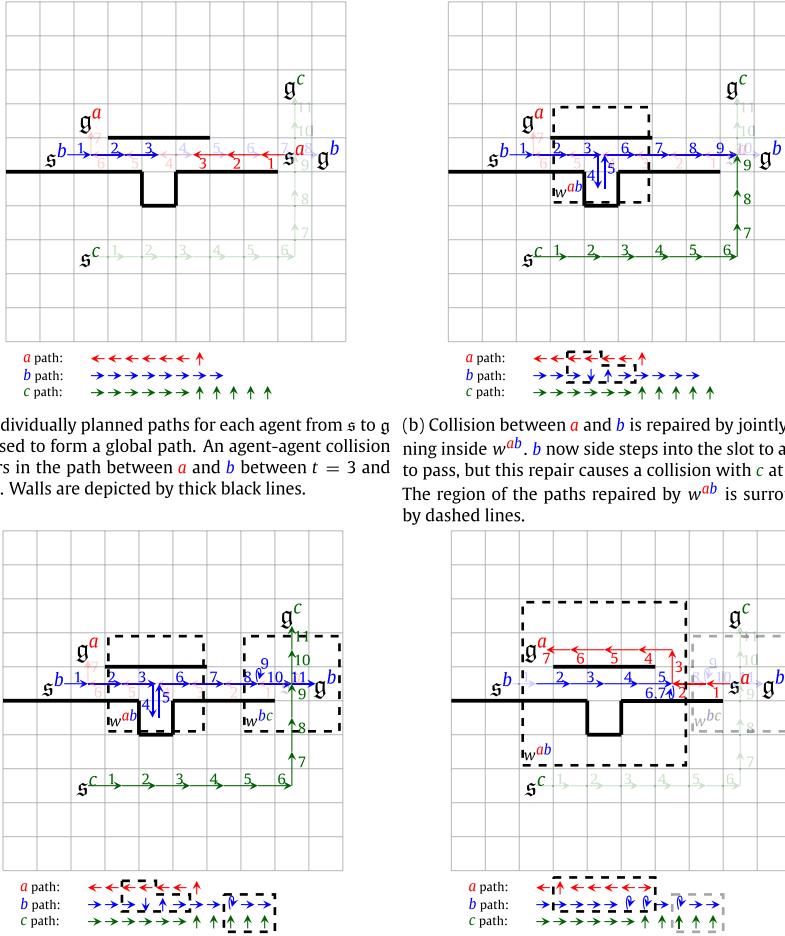
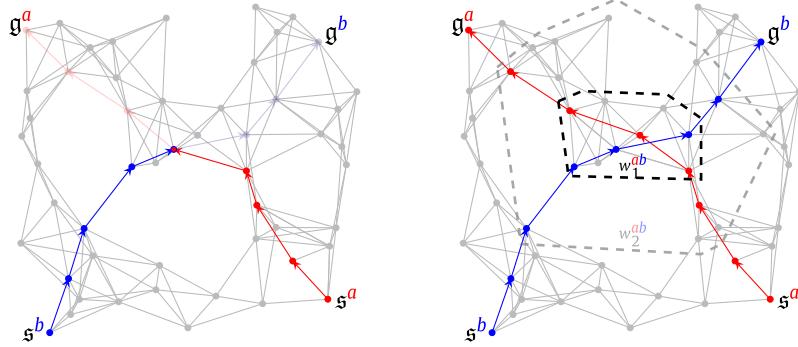


Fig. 3. Example of suboptimal, valid path generation and path padding using NWA*'s window definitions.

will travel the same paths and enter later windows as they did previously, thereby avoiding the introduction of any new collisions or invalidation of later windows and thus quickly generating successive solutions. Ultimately, w^{ab} and w^{bc} will merge, absorbing the padded section of b and allowing for the optimal global path to be generated.

4.4.4. WAMPF in domains without regular structure

While the underlying planners for WAMPF may exploit additional domain structure, WAMPF itself exploits domain structure by using the window definition to carve the graph into small, self-contained collision repair problems. To do this effectively, the problem itself must be *sparse*, i.e. amenable to this carving approach, and WAMPF must be provided with a window definition that effectively performs the carving process. As defined in Section 4.3, NWA*'s window definition uses the regular structure of four-connected grids to compactly define such a window via a hyper-rectangle. In order to generalize to other regular grids, e.g. a hexagonal grid, this definition can be augmented to fit the grid's regular shape, e.g. a hyper-hexagon, and in order to generalize to an arbitrary graph with no known additional structure, this definition can be augmented to all states at most k degrees of separation away from one or more center states. This fully general definition is shown in Fig. 4; while the graph has L_2 cost edges for ease of presentation, WAMPF knows nothing beyond the graph's fundamental definition and is still able to operate. Alternative general window definitions include adding the state least expensive to reach from a center state and outside of the window, or the set of neighbors culled during the previous repair search by the existing window's constraints (this is provided for free by X*'s *out-of-window set*, presented in Section 5).



(a) For ease of presentation, this graph has L_2 cost edges, but despite making no such assumptions about the graph's structure, WAMPF is still able to operate. There is a collision in the individually planned paths of agents a and b .

(b) On arbitrary graphs, the windows can be defined as the set of states with at most k degrees of separation from the center state. This allows WAMPF to operate on any graph without additional information about its structure.

Fig. 4. Example of WAMPF run on an arbitrary graph with no graph structure assumptions.

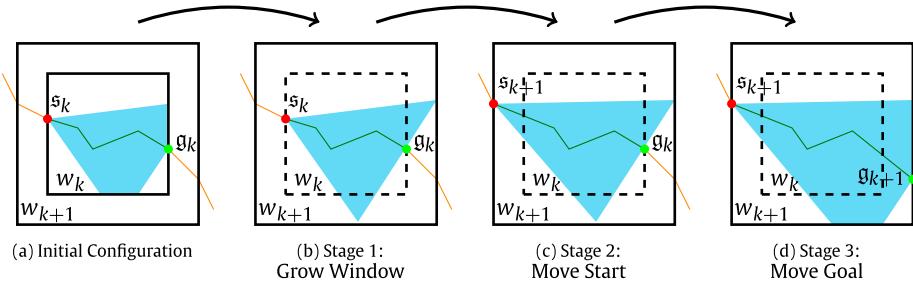


Fig. 5. Projected Illustrations of the three stage transformation employed by X^* to enable search tree re-use. w_k and w_{k+1} represent the k th and $k+1$ th windows, respectively. s_k and s_{k+1} represent the repair start for w_k and w_{k+1} , respectively. g_k and g_{k+1} represent the repair goal for w_k and w_{k+1} , respectively. Initial Configuration (Fig. 5a) show the initial search tree. Stage 1 (Fig. 5b) grows the window without moving the start or goal. Stage 2 (Fig. 5c) moves the start while keeping the same goal. Stage 3 (Fig. 5d) moves the goal.

5. Expanding A^*

Expanding A^* (X^*) is an efficient WAMPF-based planner. X^* is nearly identical to NWA* (Section 4.3), differing only in implementing additional bookkeeping to allow re-use of prior repair search information when solving for a successive repair. Due to this re-use, X^* is significantly more efficient than NWA* for successive plan generation. As we demonstrate empirically in Section 6, in sparse domains X^* outperforms the state-of-the-art in time to first solution while remaining competitive with the state-of-the-art in time to optimal solution.

5.1. X^* 's bookkeeping and search re-use for successive plan generation

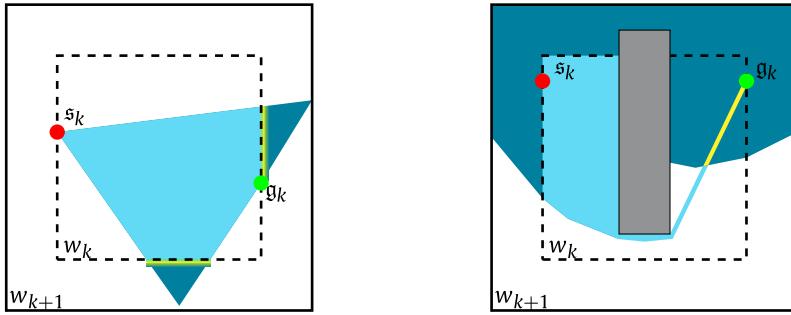
X^* 's bookkeeping during the search for a repair in the window w_k allows for the resulting search tree to be transformed into a search tree in w_{k+1} , saving computation during successive planning. The intuition behind X^* 's bookkeeping and transformations is depicted in Fig. 5 as a Projected Illustration, i.e. a two-dimensional illustration depicting the higher-dimensional joint space of w_k .

5.1.1. Search re-use: an A^* perspective

The three transformations depicted in Fig. 5 take an A^* -style Search Tree from a repair search in w_k (Fig. 5a) that produced an optimal repair in w_k and transform it into an A^* -style Search Tree for a repair search in w_{k+1} (Fig. 5d), producing an optimal repair in w_{k+1} .

Initial configuration The initial state, *Initial Configuration* (Fig. 5a), depicts a search tree from s_k to g_k restricted inside w_k ; for now, we can imagine that this search tree was produced by standard A^* .

Stage 1 The first stage, *Stage 1: Grow Window* (Fig. 5b), depicts this search tree transformed to be as if the search took place from s_k to g_k in the less restrictive w_{k+1} . In order to go from an A^* search tree in the smaller window w_k to a larger window w_{k+1} , we need to expand all the states that would have been expanded in a search of w_{k+1} but are blocked by



(a) First bookkeeping addition example. Yellow region shows states stored in out of window set during the search of w_k , where the light blue region was expanded. These yellow states form the frontier for the expansion of states when the window is grown to w_{k+1} and the search encompasses the dark blue area.

(b) Second bookkeeping addition example. Gray object is a joint space obstacle. Light blue region indicates area expanded during initial search of w_k to g_k . Dark blue region indicates area expanded during expansion of states in w_{k+1} based on the f -value of g 's expansion in w_k . Yellow region indicates states re-expanded with a lower g -value.

Fig. 6. Projected illustrations of motivating examples for the two bookkeeping additions. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

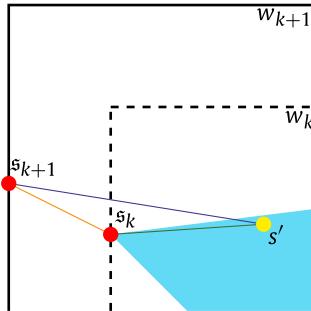


Fig. 7. Projected illustration of moving the search tree start from s_k to s_{k+1} along a section of the joint space between the two starts depicted in orange. The yellow point s' can be reached from s_{k+1} by traveling from s_{k+1} to s_k along the joint space path and then, using the information from the search tree shown in light blue, travel from s_k to s' as depicted in green. It may not always be the case that this is a minimal cost path from s_{k+1} to s' , as there may be a shorter path from s_{k+1} to s' without traveling through s_k such as the path depicted in blue, thus motivating a need to re-expansion of some states in the search tree.

w_k . These states, depicted in dark blue in Fig. 6a, must be reached via a state not in w_k whose direct predecessor is in w_k ; the set of these states is depicted in yellow in Fig. 6a. This motivates our **first bookkeeping addition: out of window set**. For each state $s \in w_k$ that was expanded, we keep track of the neighbors of s that were discarded due to the restrictions of w_k , i.e. $N(s) \setminus w_k$, placing them the *out of window set*. This bookkeeping allows us to add these states to A*'s open set O , thereby initializing the search frontier in w_{k+1} , depicted in yellow in Fig. 6a. Additionally, this bookkeeping provides a convenient way to track if the search was impeded when computing `SHOULDQuit`; if the out of window set is empty after a repair search in w_k , then the search in w_k was unimpeded.

When the window is grown, we also need to consider the possibility of new, shorter paths to already expanded states. An example of this is shown in Fig. 6b, where the gray obstacle forces a search constrained by w_k to travel below it to reach g_k , but a search in w_{k+1} allows for travel above the gray obstacle to not only reach g_k more quickly, but also more quickly reach the other states depicted in yellow. As such, we must allow for states which were expanded in the search of w_k to be re-expanded in the search of w_{k+1} if the search in w_{k+1} assigns these states a lower g -value. This motivates our **second bookkeeping addition: closed value**. In order to facilitate this re-expansion, during the initial A* search we also track the g -value at which a state is placed into the closed set C , called the state's *closed value*.

It is important to note that all states in C at the end of the search of w_k cannot be reached with a lower cost than their closed value via any path that stays entirely within w_k ; as the search in w_k is optimal, any lower cost path to any state in C must leave w_k , travel through a portion of w_{k+1} , and re-enter w_k , just as the path above the gray obstacle did in Fig. 6b. Thus, the addition of the states to O from our out of window set (first bookkeeping addition) ensures that all of such paths are able to be considered as long as states are able to be re-expanded if their closed g -value, recorded by our closed value (second bookkeeping addition), is higher than their g -value as they sit in O . With this modification, we can run A* until the minimal f -value in O is greater than the f -value of g_k . This will update all of the states in C and O to have the optimal g -value for a search in w_{k+1} and thus produce the search tree shown in Stage 1.

The second stage, *Stage 2: Move Start* (Fig. 5c), depicts the search tree transformed from a start s_k as seen in Stage 1 to a start s_{k+1} . In order to move the start backwards, we need to embed the search tree rooted at s_k into the search tree rooted at s_{k+1} . As illustrated in Fig. 7, to reach any state in the existing tree from s_{k+1} , e.g. s' (depicted in yellow), the cost of the minimal path (depicted in blue) is upperbounded by the cost to travel from s_{k+1} to s (depicted in orange) plus the cost to travel from s_k to s' (depicted in green). This holds because the orange path from s_{k+1} to s_k is extracted from the global path π , which is provably collision-free in this region (see Appendix B, Theorem 1 for proof), and thus serves a valid upperbound, and the green path from s_k to s' is provided by the g -values of the existing search tree and thus is the optimal cost from s_k to s' . Thus, if we increase every state's g -value and closed value (second bookkeeping addition) by the cost of the path from s_{k+1} to s_k , and we expand each state along the path from s_{k+1} to s_k , we can run A^* until the minimal f -value in O is greater than the f -value of g_k , leveraging the second bookkeeping addition to re-expand states in the s_k rooted tree as needed, as done in Stage 1.

The third stage, *Stage 3: Move Goal* (Fig. 5d), depicts the search tree rooted at s_{k+1} transformed from a goal g_k as seen in Stage 2 to a goal g_{k+1} . The states in O simply need to have their f -values updated with new h -values to g_{k+1} and then A^* can be run as normal until g_{k+1} is expanded.

5.1.2. Bookkeeping formalization

In order to be able to reason about a state's f -value, g -value, and h -value under different starts and goals, we augment the f , g , and h function with start and goal parameters. For example, given a state s , start s_k , and goal g_k , s 's f -value, g -value, and h -value are $f(s, s_k, g_k)$, $g(s, s_k)$, and $h(s, g_k)$, respectively. Like standard A^* , if any g -value entry has not been set, it returns ∞ .

Section 5.1.1 discusses two bookkeeping additions to standard A^* search trees that facilitate the search re-use depicted in Fig. 5. The first bookkeeping addition, called an *out of window* set X , maintains a set of all states that are neighbors of expanded states in w_k and themselves are not in w_k . In Stage 1, when w_k is grown to w_{k+1} , the states $\{s \mid s \in X \cap w_{k+1}\}$ are added to O and removed from X . The second bookkeeping addition, called a state s 's *closed cost*, is recorded in $\mathcal{C}(s, s) \leftarrow g(s, s)$ when s is placed in C ; this is similar to the bookkeeping done when running A^* with an inconsistent heuristic [14]. This table is checked during state expansions in Stage 1 and Stage 2's transformations in order to allow the re-expansion of states which have shorter paths. Like g -values, if an entry in \mathcal{C} has not been set, it returns ∞ .

5.2. WAMPF subroutine implementations

Three of X^* 's five key implementations are identical to NWA^* (Section 4.3); however, the other two make use of the guarantees provided by WAMPF regarding the ordering of `PLANIN` and `GROWANDREPLANIN` calls on successor windows to improve efficiency. Additionally, for these re-use techniques to work, we assume the heuristic is consistent, i.e. the triangle inequality holds.

`PLANIN(w, π)`: This subroutine is implemented almost identically to NWA^* 's `PLANIN` in Section 4.3, but with the implementation of the two bookkeeping additions from Section 5.1.2. `A*WITHBOOKKEEPING` in Algorithm 2 is A^* modified with these bookkeeping additions.

`GROWANDREPLANIN(w_k, π)`: As defined in Section 4, `GROWANDREPLANIN` will only be invoked on a window in which `GROWANDREPLANIN` of `PLANIN` were previously invoked. As such, this subroutine leverages the X^* Search Tree produced by the previous search of w_k to aid the current search of w_{k+1} via the transformation shown in Fig. 5 and discussed in Section 5.1. The algorithm and its supporting procedures are presented in Algorithm 2.

`CrowAndReplanIn` This algorithm performs setup and wraps three subroutines corresponding to the three stages shown in Fig. 5. Importantly, on Line 5, X^* replaces a section of the existing path π with its repair π' . Due to the fact that we are growing an existing repair, π is already a valid global path which we are improving. As such, we must ensure that if π' is shorter than the existing region in π , π' is padded so that all agents leave the state g_{k+1} at the same time they did in π ; this is critical to ensuring any window repairs further along π continue to have start states that are reachable from π .

A^* search until As discussed in Section 5.1.1, Stage 1 and Stage 2 need to expand all states with less than or equal to a given f -value in order to ensure that states have the minimal cost g -value for the given window. `A* SEARCH UNTIL` is a helper function which provides this functionality for a given f -value, f_{max} , by running a modified A^* search which only terminates when the minimal f -value of any state in O is greater than f_{max} . Note that the expansion skip condition for a state expansion (Line 17) also considers the closed value of the state, allowing for `A* SEARCH UNTIL` to re-expand a state if its g -value is lower than its closed value.

A^* with bookkeeping This procedure runs standard A^* from the given start s to the given goal g in the given window w using the given open set O , closed set C , and out of window set X . Note that, unlike `A* SEARCH UNTIL`, the expansion skip condition for `A*WITHBOOKKEEPING` is a standard A^* -style C membership check (Line 24).

Stage 1 This procedure converts the tree shown in Initial Configuration (Fig. 5a) into Stage 1 (Fig. 5b). It does this by initializing O with the frontier of the search for states in w_{k+1} but not in w_k and then leverages `A* SEARCH UNTIL` to expand

Algorithm 2 GrowAndReplanIn.

```

1: function GROWANDREPLANIN( $w_k, \pi$ )
2:   STAGE1
3:   STAGE2
4:    $\pi' \leftarrow \text{STAGE3}$ 
5:   Replace section of  $\Phi(\pi, \alpha)$  from  $w_{k+1}$ 's  $s$  to  $g$  with  $\pi'$ 
6:   return  $(w_{k+1}, \pi)$ 
7: procedure EXPANDSTATE( $s, s$ )
8:    $C \leftarrow C \cup \{s\}$ 
9:    $\mathcal{C}(s, s) \leftarrow g(s, s)$ 
10:   $O \leftarrow O \cup \{n \mid n \in N(s) : n \in w\}$ 
11:   $X \leftarrow X \cup \{n \mid n \in N(s) : n \notin w\}$ 
12:  for all  $n \in N(s)$  do  $g(n, s) \leftarrow \min(g(n, s), g(s, s) + c(s, n))$ 
13: procedure A*SEARCHUNTIL( $O, C, X, w, f_{\max}$ )
14:   while  $f(\text{top}(O, s, g), s, g) \leq f_{\max}$  do
15:      $s \leftarrow \text{top}(O, s, g)$ 
16:      $O \leftarrow O \setminus \{s\}$ 
17:     if  $s \in C \wedge \mathcal{C}(s, s) \leq g(s, s)$  then continue
18:     EXPANDSTATE( $s, s$ )
19: procedure A*WITHBOOKKEEPING( $O, C, X, w, s, g$ )
20:   while  $O \neq \emptyset$  do
21:      $s \leftarrow \text{top}(O, s, g)$ 
22:     if  $s = g$  then return UNWINDPATH( $C, g, s$ )
23:      $O \leftarrow O \setminus \{s\}$ 
24:     if  $s \in C$  then continue
25:     EXPANDSTATE( $s, s$ )
26:   return NOPATH
27: procedure STAGE1
28:    $O \leftarrow O \cup \{s \mid s \in X : s \in w_{k+1}\}$ 
29:    $X \leftarrow \{s \mid s \in X : s \notin w_{k+1}\}$ 
30:   A*SEARCHUNTIL( $O, C, X, w_{k+1}, f(g_k, s_k, g_k)$ )
31: procedure STAGE2
32:    $\pi' \leftarrow \text{path between } s_{k+1} \text{ and } s_k \text{ extracted from } \pi$ 
33:   for all  $s \in O \cup C$  do  $g(s, s_{k+1}) \leftarrow g(s, s_k) + \|\pi'\|$ 
34:   for all  $s \in C$  do  $\mathcal{C}(s, s_{k+1}) \leftarrow \mathcal{C}(s, s_k) + \|\pi'\|$ 
35:   for all  $s \in \pi'$  do EXPANDSTATE( $s, s_{k+1}$ )
36:   A*SEARCHUNTIL( $O, C, X, w_{k+1}, f(g_k, s_k, g_k) + \|\pi'\|$ )
37: procedure STAGE3
38:   Reorder  $O$  using  $f$ -value from  $s_{k+1}$  to  $g_{k+1}$ . ▷  $h$ -values changed
39:   if  $g_{k+1} \in C$  then return UNWINDPATH( $C, g_{k+1}, s_{k+1}$ )
40:   return A*WITHBOOKKEEPING( $O, C, X, w_{k+1}, s_{k+1}, g_{k+1}$ )

```

or re-expand states with f -values less than the f -value of g_k , as these states would have been expanded during a direct search of w_{k+1} .

Stage 2 This procedure converts the tree shown in Stage 1 (Fig. 5b) into Stage 2 (Fig. 5c). As discussed in Section 5.1.1, it does this by extracting the relevant section of π from s_{k+1} to s_k (Line 32). The path cost is used to increase the g -value of each state in O and C (Line 33), as well as the closed value of each state in C (Line 34). Then, all states along this path are expanded (Line 35). Note that as each state's g -value in O is increased by a fixed amount, no reordering of O is required even if backed by an ordered data structure (e.g. a heap).

Stage 3 This procedure converts the tree shown in Stage 2 (Fig. 5c) into Stage 3 (Fig. 5d). As the goal moves from g_k to g_{k+1} , the heuristic evaluation for each state in O will change by differing amounts for various states and thus, if O is backed by a structure such as a heap, it will require reordering (Line 39). The rest of STAGE3 is standard A* with bookkeeping additions (Line 40).

6. Empirical results

In this section, we evaluate X^* using randomly generated four-connected grids (e.g. Fig. 8a) and several standard benchmark domains¹ (e.g. Fig. 8b). All experiments treat the domains as uniform cost four-connected grids with randomly selected starts and goals. Unless stated otherwise, X^* is configured with an initial window $L_\infty = 2$ and window expansions grow the window by a single step. All boxplot whiskers are at most the length of the interquartile range, with the lower whisker fit

¹ den520d, brc202d, lak303d, ht_mansion_n, ost003d, and w_woundedcoast domains were used. Benchmarks available at <https://movingai.com/benchmarks/mapf/index.html>.

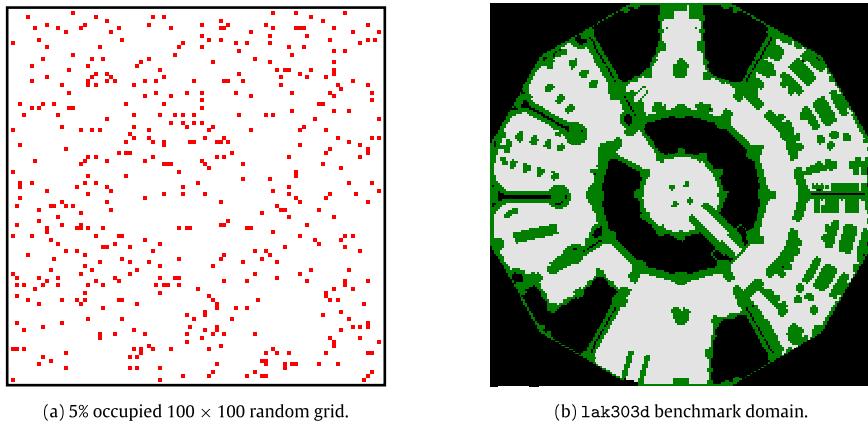


Fig. 8. Examples of the domains in which experiments all experiments were performed.

to the lowest datapoint above this value and the upper whisker fit to the highest datapoint below this value. We use these domains in multiple experiments to evaluate:

1. How X^* compares to state-of-the-art MAPF planners in time to generate a valid path in sparse domains (Section 6.1).
2. How X^* compares to state-of-the-art MAPF planners in time to generate an optimal path in sparse domains (Section 6.2).
3. How X^* compares to NWA* in valid path generation performance and optimal path generation performance (Section 6.3).
4. Which components of X^* dominate its runtime (Section 6.4).
5. The effect of domain characteristics on the performance of X^* (Section 6.5).
6. Suboptimality bounds of X^* 's first and intermediary paths (Section 6.6)
7. The effect of parameters on the performance of X^* (Section 6.7).

All planners were implemented in C++. X^* and NWA* were implemented by the authors of this paper,² AFS was implemented by its original authors, CBS was implemented by a third party,³ M* was implemented by its original authors,⁴ (Operator Decomposition version used), and PR was implemented by a third party and modified by the authors of this paper.⁵ All runtime measurements were performed on a dedicated computer with an Intel i7 CPU (TurboBoost disabled) and access to 60GB of DDR4 RAM. Any trial that exceeded the memory limit was recorded as a timeout.

6.1. Comparison for time to valid path

In order to evaluate the performance of X^* compared to state-of-the-art anytime or optimal MAPF solvers for time to valid path generation in sparse domains, we run X^* , AFS, CBS, and M* with varied numbers of agents on randomly generated 100×100 grids with 1%, 5%, and 10% of the states blocked (Fig. 9) and on standard benchmark domains for fixed number of agents (Table 1, first rows).

Fig. 9 demonstrates that in random domains, X^* outperforms the state-of-the-art MAPF planners in time to a valid path. X^* 's improved performance is most distinct in domains with 1% of states blocked, as these domains are especially sparse and thus amenable to X^* 's approach; as the density of obstacles increases and thus domain sparsity decreases, the gap between X^* 's performance and the state-of-the-art MAPF planners shrinks but is still pronounced. Compared to AFS and M*, X^* on average produces a path at least an order of magnitude faster; while some of this performance difference may be the result of differing implementation quality, much of it can be attributed to the overhead of requiring a full joint search for AFS or individual space policy computations for M*. Compared to CBS, X^* performs just as well for small numbers of agents, producing paths for 10 agents in under 10 milliseconds, but as the number of agents increases, performance diverges in favor of X^* .

Table 1 reaffirms that X^* is significantly faster than AFS and M* for time to valid path generation with over a two order of magnitude faster time, while CBS and X^* are highly competitive; this result is a reflection of the high degree of sparsity in these domains and the initial overhead of AFS and M*.

² Source code available at <https://github.com/kylevedder/libMultiRobotPlanning>.

³ Source code available at <https://github.com/whoenig/libMultiRobotPlanning/>.

⁴ Source code available at https://github.com/gswagner/mstar_public/.

⁵ Source code available at <https://github.com/kylevedder/Push-and-Rotate>.

Table 1

X^* , CBS, AFS, and M^* run on various standard benchmarks for 50 agents on all 25 provided random instances with a timeout of 300 seconds. Median time to valid path is reported in the first row and time to optimal path is reported in the second row in seconds.

| Scenario | X^* | CBS | M^* | AFS |
|----------------|--------|--------|--------|---------|
| den520d | 0.0026 | – | – | 12.0885 |
| | 0.0027 | 0.0024 | 4.2496 | 12.0885 |
| brc202d | 0.0038 | – | – | 8.1243 |
| | 0.0050 | 0.0037 | 4.6910 | 8.1310 |
| lak303d | 0.0023 | – | – | 2.6335 |
| | 0.0052 | 0.0023 | 1.5907 | 2.6335 |
| ht_mansion_n | 0.0021 | – | – | 0.7354 |
| | 0.0035 | 0.0017 | 0.7301 | 0.7357 |
| ost003d | 0.0022 | – | – | 2.1470 |
| | 0.0037 | 0.0018 | 1.4160 | 2.1470 |
| w_woundedcoast | 0.0104 | – | – | 1.9448 |
| | 0.0180 | 0.0173 | 3.1426 | 1.9466 |

Like all other planners in Fig. 9 and Table 1, X^* fails to generate a path in a reasonable amount of time for particularly challenging problems. As discussed in Section 6.4, this is caused by high dimensional searches resulting from repairs in windows with a large number of agents.

In order to evaluate the performance of X^* compared to suboptimal MAPF solvers for time to valid path generation in sparse domains, we run X^* and PR with varied numbers of agents on randomly generated 100×100 grids with 1%, 5%, and 10% of the states blocked (Fig. 10). These results demonstrate that in random domains for small numbers of agents, X^* outperforms PR in time to first path; while some of the performance difference can be attributed to implementation quality, much of it can be attributed to the fact that X^* exploits the sparsity present in these test domains while PR does not. PR provides much more consistent runtimes in its valid path generation, solving all scenarios for all agent counts in under $\frac{1}{5}$ th of a second and scaling quasi-linearly with increasing agent counts; however, PR provides significantly lower path quality than X^* . PR's median path suboptimality factor, computed against an optimal path generated post-hoc, was (2.0020, 2.0673, 2.1372) across all runs for 1%, 5%, and 10% obstacle density, respectively. X^* 's online suboptimality factor, an exact or overestimate of the true suboptimality factor, was (1.0029, 1.0029, 1.0029) across all runs for 1%, 5%, and 10% obstacle density, respectively (a full analysis of X^* 's path quality bounds is presented in Section 6.6). This experiment demonstrates X^* 's advantage for time to valid path generation for small numbers of agents or when path quality is important.

6.2. Comparison for time to optimal path

In order to evaluate the performance of X^* compared to state-of-the-art MAPF solvers for time to optimal path generation in sparse domains, we run X^* , AFS, CBS, and M^* with varied numbers of agents on randomly generated 100×100 grids with 1%, 5%, and 10% of the states blocked (Fig. 11) and on standard benchmark domains for fixed number of agents (Table 1, second rows).

Fig. 11 demonstrates that in random domains, X^* is competitive with state-of-the-art MAPF planners in time to an optimal path. Like with time to valid path, X^* is most competitive when the domains are sparser, i.e. lower numbers of agents or fewer blocked states. Against AFS and M^* , for small numbers of agents X^* exhibits a significantly lower mean and lower quartile runtime; for larger numbers of agents, X^* exhibits similar or higher means and but significantly faster lower quartile times. Against CBS, X^* has either higher or similar means with heavily overlapping interquartile ranges and lower quartiles. As discussed in Section 6.4 the variance in X^* 's optimal path generation time can be attributed to the variance in the number of agents involved in any window search; as X^* repeatedly grows windows the likelihood window merges thus requiring higher dimensional searches increases, contributing to the large spread in runtimes.

Table 1 show that X^* 's is significantly faster than AFS and M^* for time to optimal path generation with over a two order of magnitude faster time, while CBS and X^* are highly competitive; this is a reflection of the high sparsity of the benchmark domains and the initial overhead of AFS and M^* .

6.3. X^* versus baselines

X^* operates by restricting the initial repair search space, quickly finding a repair in this restricted search space to produce a valid global path, then relaxing the restriction and repeating the process until an optimal global path is found. While this approach provides WAMPF's anytime property, it also incurs computational overhead, even in planners like X^* which perform reuse between repair searches.

To demonstrate this overhead, we run X^* , NWA* and A* on a 20×20 four-connected grid scenario with an agent starting on the center of each edge and with a goal on the center of the opposite edge, thereby inducing a four agent collision in

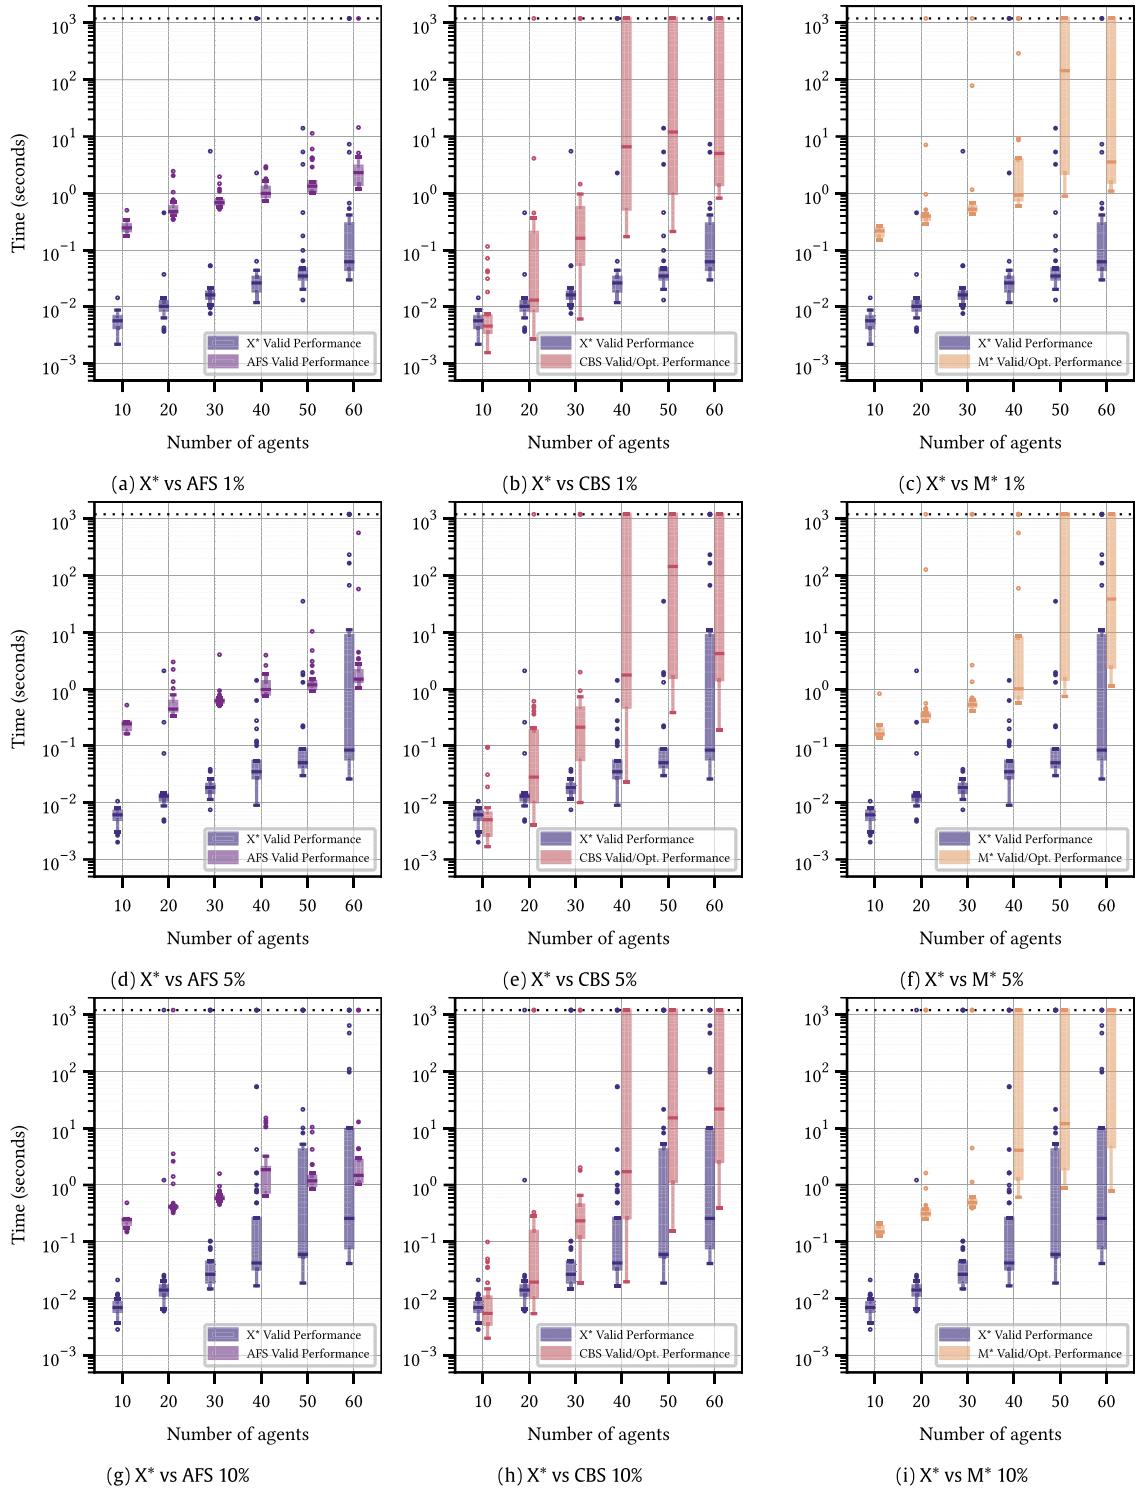


Fig. 9. Box plots of time to valid path for X^* , AFS, CBS, and M^* on a log scale. For each agent count, 30 trials are run, each with a 20 minute timeout, with each trial run on a randomly generated 100×100 four-connected grid. The percentage of states blocked is listed in each caption.

the center of the scenario. While A^* will directly solve for an optimal path, X^* and NWA^* will quickly produce a valid global path, multiple intermediary global paths, and terminate with a provably optimal global path.

The runtime results are presented in Table 2, with 95% confidence intervals over 30 trials. Due to the nearly identical structure of their initial path generation, NWA^* and X^* have nearly identical performance for time to first path, outper-

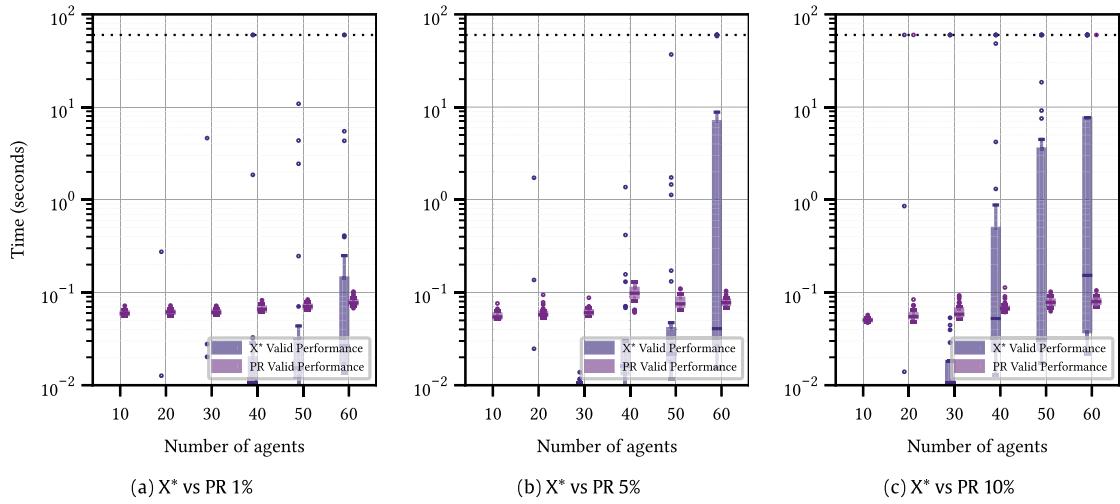


Fig. 10. Box plots of time to valid path for X^* vs PR. For each agent count, 30 trials are run, each with a 1 minute timeout, with each trial run on a randomly generated 100×100 four-connected grid. The percentage of states blocked is listed in each caption.

Table 2

X^* , NWA^* , and A^* run on a 20×20 grid with an agent starting on the center of each edge and with a goal on the center of the opposite edge to demonstrate the overhead of WAMPF-style window growth compared to A^* . Each result is reported as a percentage of the total A^* runtime. Column *current iteration* represents the RecWAMPF iteration the given planner was on when A^* terminated with the optimal solution and column *total iterations* represents the total RecWAMPF iterations needed to generate an optimal solution.

| Planner | Valid Path Runtime median A^* runtime | Optimal Path Runtime median A^* runtime | current iteration | total iterations |
|---------|--|--|-------------------|------------------|
| X^* | 6.32% | 175.18% | 6 | 9 |
| NWA^* | 6.28% | 547.20% | 4 | 9 |
| A^* | 100.00% | 100.00% | – | – |

forming A^* 's time to its first path by over an order of magnitude. Due to the window overhead, X^* takes approximately 1.5 times longer than A^* to produce an optimal global path, having finished 5 of the needed 9 window expansions when A^* terminates, and NWA^* takes approximately 6x longer than A^* to produce an optimal global path due to a lack of search re-use, having finished 3 of the needed 9 window expansions when A^* terminates. This result demonstrates the efficacy of X^* 's search reuse techniques in improving its optimal path generation performance and demonstrates to practitioners that, while X^* and NWA^* have the same first path runtime, X^* strictly dominates NWA^* in time to optimal path.

6.4. X^* components that dominate runtime

In order to optimize X^* , be it from an implementation perspective or a theoretical one, is important to understand which components dominate its runtime. X^* 's runtime is dominated by `PLANIN` and `GROWANDREPLANIN`, where the window searches with the highest number of agents dominate both time to valid path (Fig. 12a) and time to optimal path (Fig. 12b). Fortunately, for random domains with various agent counts, as the magnitude of the Largest Number of Agents In Any Window (LNAIAW) grows linearly, the number of occurrences of such a window decreases exponentially for valid path generation (Fig. 12a) and linearly for optimal path generation (Fig. 12b). This finding also provides an opportunity for practitioners to build an X^* -based composite WAMPF solver that falls back on another MAPF solver when a high dimensional window is detected, preventing X^* from performing a potentially expensive search.

6.5. X^* runtime versus sparsity of domain

X^* is designed to exploit sparsity of agent-agent collisions in order to quickly develop a suboptimal but valid path as well as produce an optimal global path. First, to demonstrate that X^* does exploit available sparsity in practice, we look at X^* 's success at keeping the number of agents involved in each window low, measured by the magnitude of the Largest Number of Agents In Any Window (LNAIAW). Fig. 12 demonstrates that as the obstacle density of the domain increases, and thus sparsity decreases, the magnitude of LNAIAW increases; this is especially clear in time to valid path (Fig. 12a), where there is a clear increasing trend in the distribution of LNAIAW from 1% occupied to 10% occupied grids, but a similar trend exists in time to optimal path (Fig. 12b). These trends are the result of the fact that in domains with relatively high sparsity, e.g. the 1% occupied grids, X^* is able to cleanly separate collisions from one another, while in less sparse domains, e.g. the 10% occupied grids, X^* cannot separate collisions as well and must form windows with more agents.

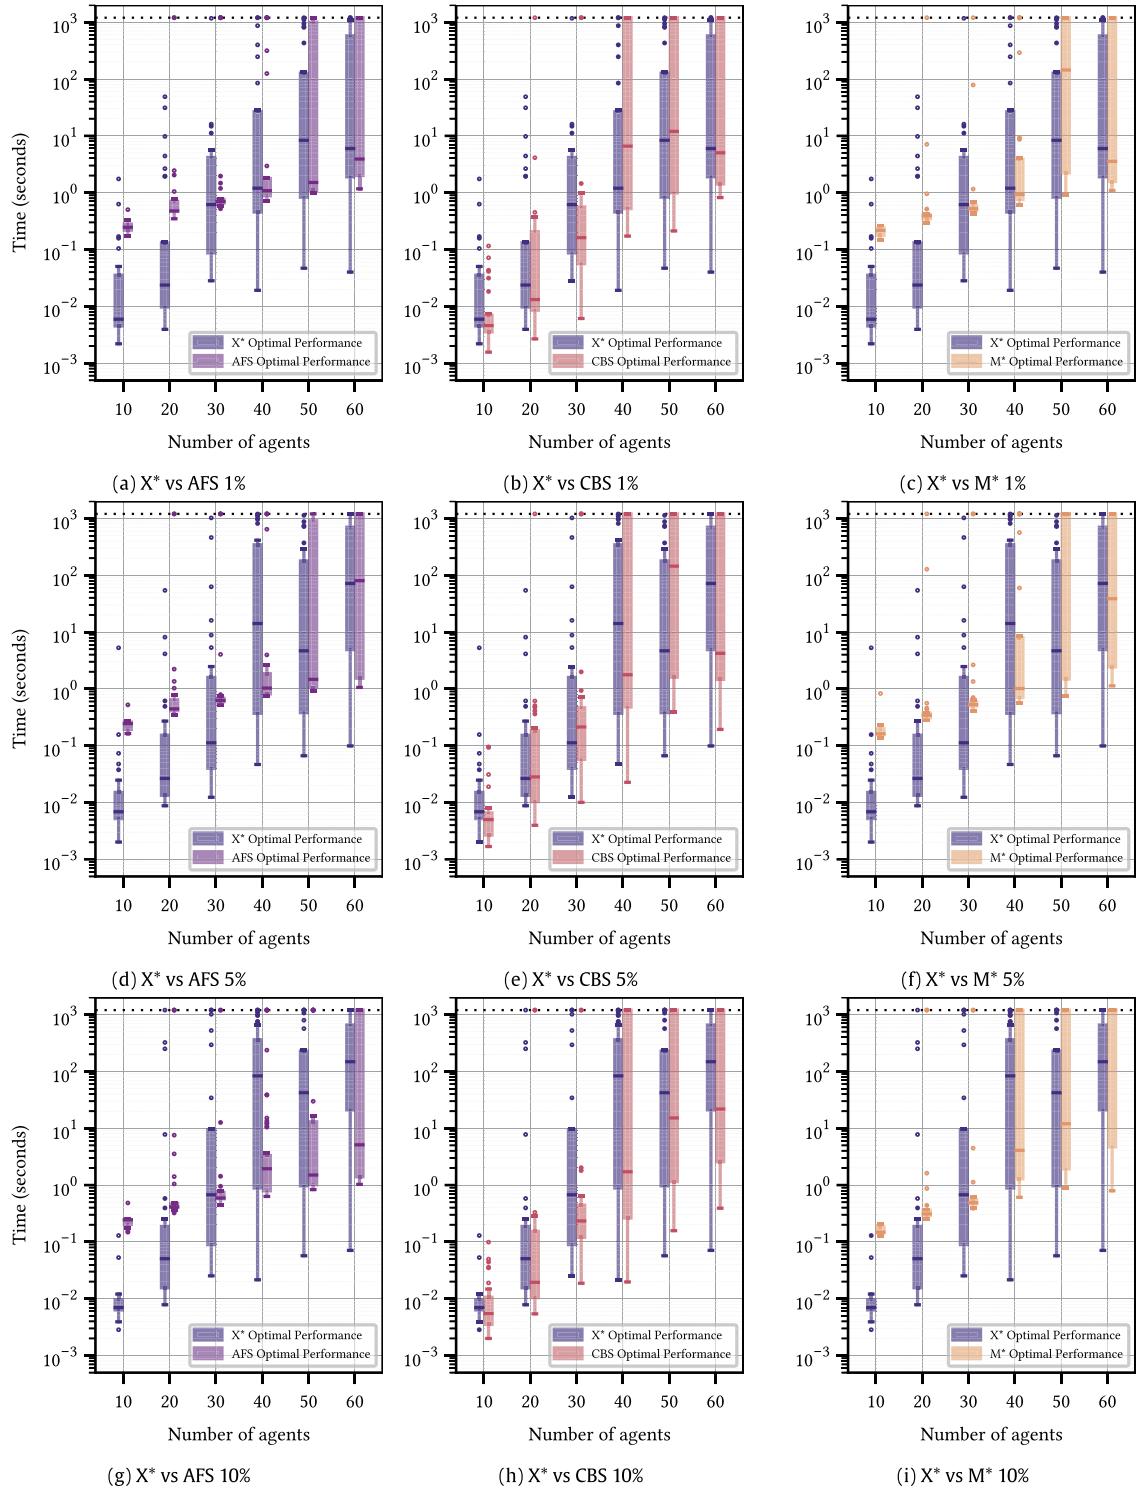


Fig. 11. Box plots of time to optimal path for X^* , AFS, CBS, and M^* on a log scale. For each agent count, 30 trials are run, each with a 20 minute timeout, with each trial run on a randomly generated 100×100 four-connected grid. The percentage of states blocked is listed in each caption.

Second, as a result of the fact that X^* exploits sparsity, it is expected that X^* will scale well when the number of agents in a domain increases but the level of sparsity stays the same. To validate this expectation, we run X^* on varying sized four-connected grids with a 10% obstacle density and constant $\frac{\text{grid area}}{\text{agent count}}$ ratio of 500 in an attempt to maintain similar

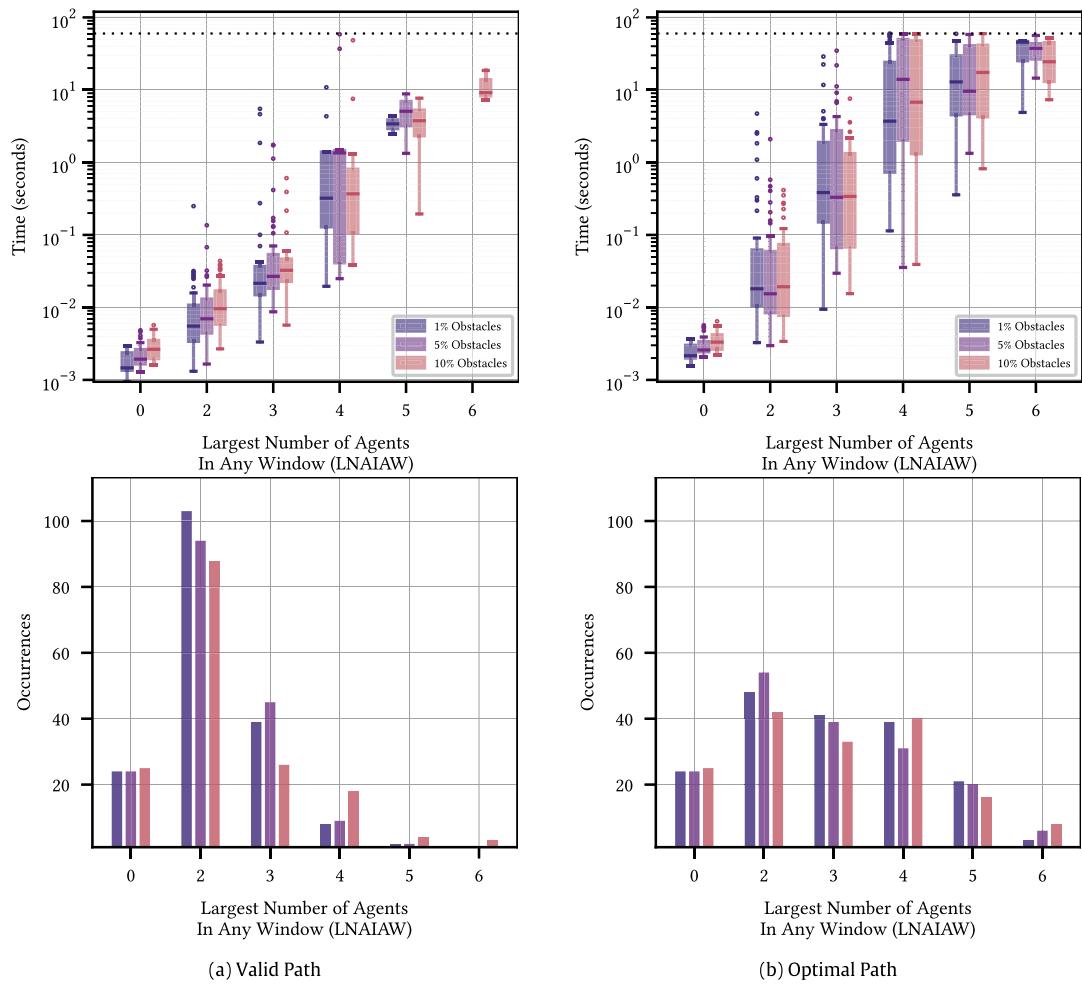


Fig. 12. Time to valid and optimal paths for X^* on a log scale vs Largest Number of Agents In Any Window (LNIAIW), along with the frequency of each LNIAIW. Run across 20 to 60 agents in steps of 10, each of 30 trials, on 100×100 four-connected grids with 1%, 5%, and 10% obstacle density with a timeout of 60 seconds.

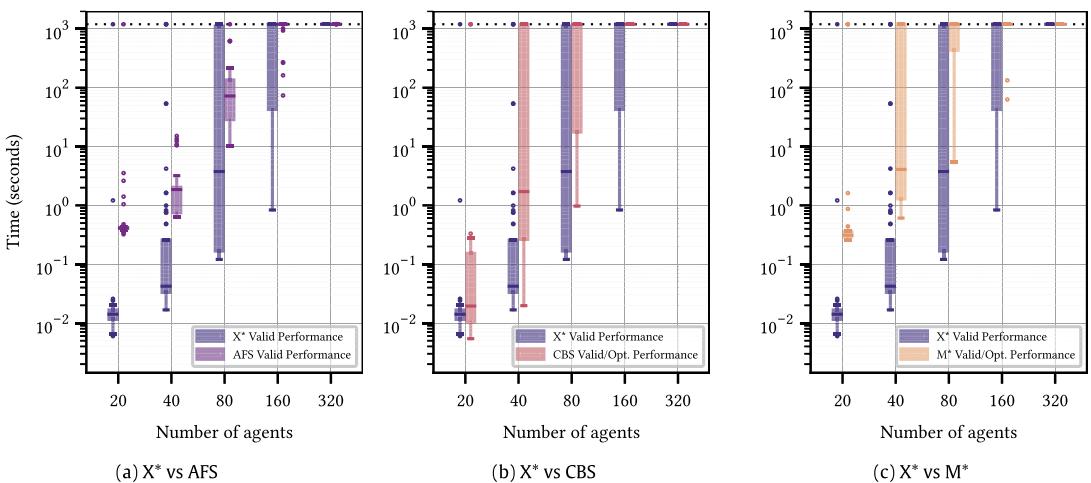


Fig. 13. Time to valid path for X^* vs AFS, CBS, and M^* on a log scale. For each agent count, 30 trials are run, each with a 20 minute timeout, with each trial run on with a constant $\frac{\text{grid area}}{\text{agent count}}$ ratio of 500 with a 10% obstacle density.

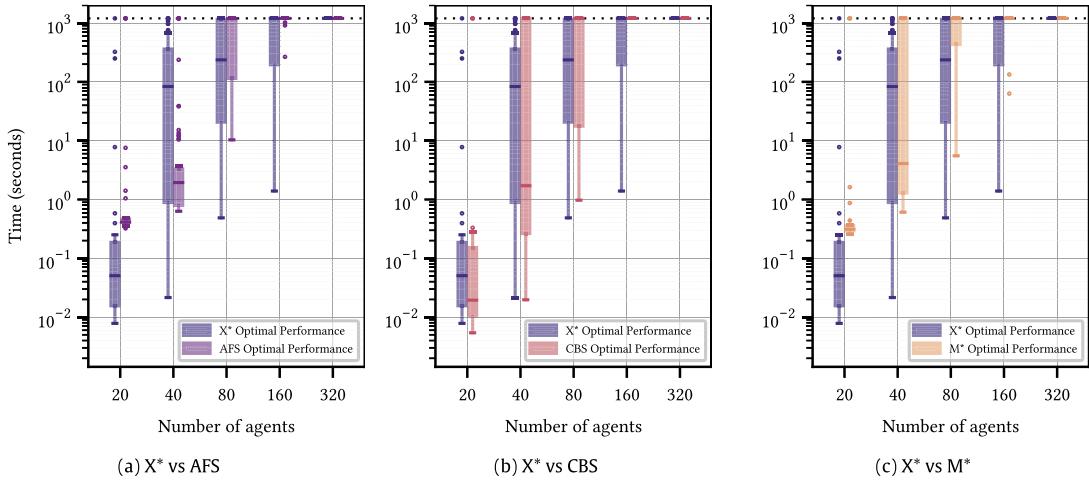


Fig. 14. Time to optimal path for X^* vs AFS, CBS, and M^* on a log scale. For each agent count, 30 trials are run, each with a 20 minute timeout, with each trial run on with a constant $\frac{\text{grid area}}{\text{agent count}}$ ratio of 500 with a 10% obstacle density.

levels of domain sparsity. We also run CBS, AFS, and M^* on the same domains to provide a frame of reference. Time to valid global path is presented in Fig. 13 and time to optimal global path is presented in Fig. 14.

For time to optimal path, X^* 's median time is consistently faster than any other planner, its lower quartile is consistently two orders of magnitude faster than AFS or M^* and it scales better than any other planner; with the exception of a few instances solved by AFS and M^* , X^* was the only planner able to produce paths for the 160 agent case, in some cases producing paths in under a second; X^* 's superior performance against these other planners is due to its ability to exploit domain sparsity to greater effect.

For time to optimal path, X^* has a higher median runtime than the other planners for lower numbers of agents; however, for 80 agents, X^* 's median runtime is below the timeout threshold while all other planners medians are at the timeout threshold and, with the exception of a few instances solved by AFS and M^* , X^* is the only planner able to generate optimal paths for 160 agents.

Together, these findings suggest that, compared to state-of-the-art algorithms, X^* 's approach scales well to large numbers of agents across domains with similar levels of sparsity.

6.6. Suboptimality bounds on intermediary paths

For the ϵ -suboptimal intermediary solutions of an anytime planner to be useful in practice, the ϵ bound must be reasonably tight. In order to characterize X^* 's ϵ bound in practice, we ran X^* for 30 trials on a 100×100 random grid with 30 agents and varied obstacle density. The results for the first 20 X^* iterations (recursive invocations of RecWAMPF), shown in Fig. 15, were collected from the same experiments shown in Fig. 9 and Fig. 11.

These results demonstrate that, in practice, X^* 's first valid path cost is almost always within 0.5% of the optimal path and outliers are quickly improved upon within a few additional iterations of X^* . For practitioners, these results indicate that X^* 's first path is often of sufficient quality and, if not, a few additional iterations of RecWAMPF should be sufficient to bring the path quality within a tight quality bound.

6.7. X^* window selection impact on runtime

As shown in Section 6.4, window dimensionality dominates runtime. As such, selecting the proper initial window size to repair a search in order to minimize window merges is an important factor in X^* 's valid path generation performance. Fig. 16a shows the impact of the initial window radius parameter on X^* 's time to valid path; unsurprisingly, smaller window radii more quickly produce a valid global path due to a decreased likelihood of requiring window merges.

However, smaller window radii can *increase* time to optimal path in some cases. Shown in Fig. 16b, an initial window radius of 1 or 2 result interval bounds that are roughly 5x higher than the bounds produced by initial window radii of 3, 4, and 5, with similar performance differences even in outliers. The root cause of this performance degradation is the expansion of states during a small window search which would not be expanded by a fresh search in a larger window, such as depicted in the large dark blue area of Fig. 6b. As such, these unnecessary expansions earlier in X^* 's search will add states to O to be expanded which would never be considered by a search that initially had a larger window. The exact radius values for which performance degrades changes across scenarios as a consequence of the structure of the domain, making this analysis important for practitioners who care about time to optimal path.

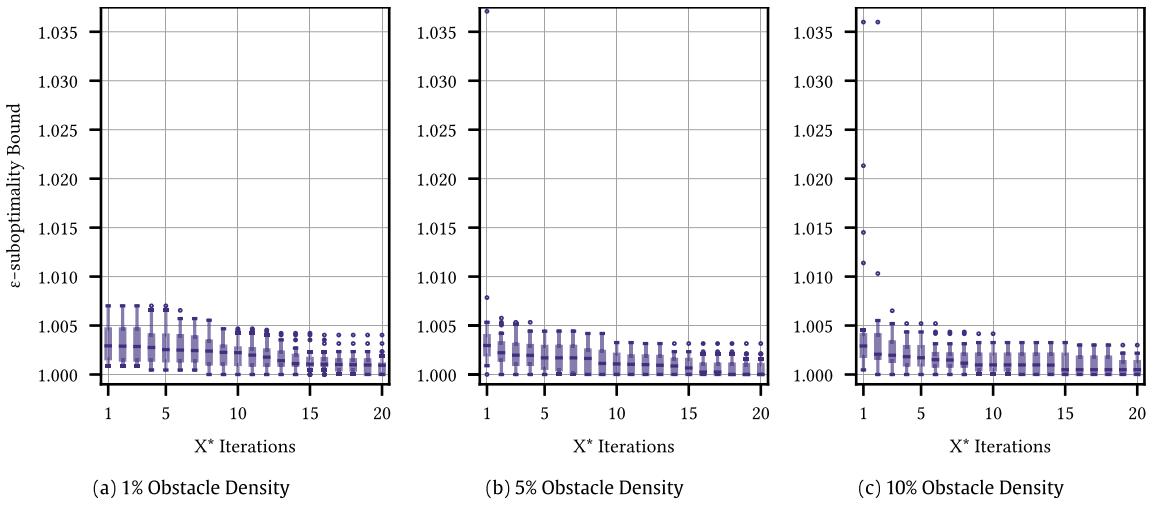


Fig. 15. ϵ -suboptimality bounds for the first 20 iterations of 30 trials of X^* on 100×100 random grids for 30 agents. The 3 trials that failed to produce any path in 10% Obstacle Density were not recorded. The trials that terminated in fewer than 20 iterations had their last bound duplicated for the remaining iterations.

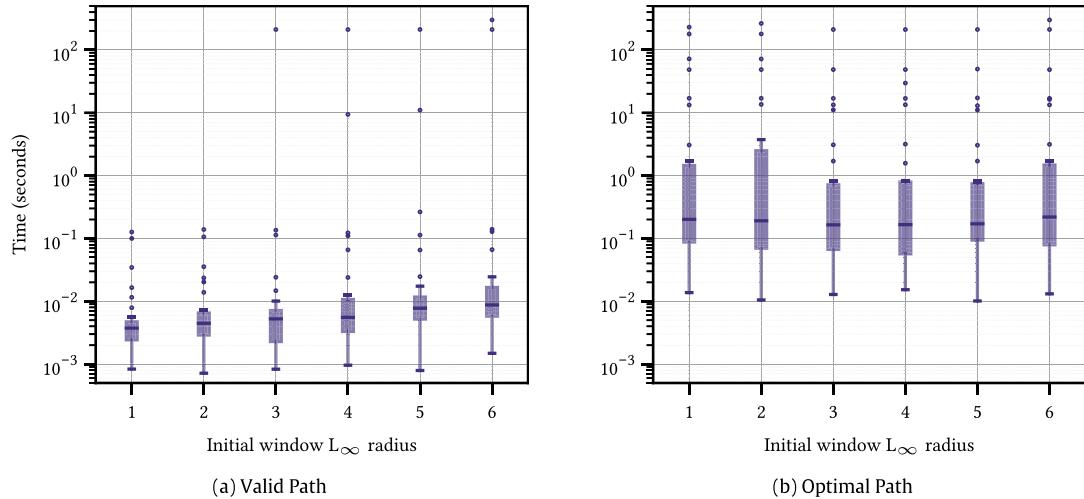


Fig. 16. Time to valid path and time to optimal path vs initial window radius for X^* on a log scale. Run across 30 trials of 30 agents on 100×100 four-connected grids with 5% obstacle density.

7. Future work

X^* uses standard A* to perform optimal window searches; if a fast optimal MAPF solver such as CBS or an anytime MAPF solver such as AFS were used to admit suboptimal repairs inside a window, and this search tree could be grown using X^* -style reuse, this approach may produce a WAMPF planner faster than X^* . This investigation would also lend itself to exploring ϵ -suboptimal WAMPF.

In addition, there is room for further exploration of window size and shape; in this work we used rectangular windows for NWA* and X^* because they are simple to reason about and performed better in our initial experimentation than rasterized spheres, but there may be other shapes that are better suited to WAMPF.

Finally, we believe that further investigation into quantifying sparsity of MAPF domains would provide great insight into the fundamental nature of MAPF and potentially allow for the development of an ensemble MAPF solver that switches techniques based on individual problem structure.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work is supported in part by AFRL and DARPA under agreement #FA8750-16-2-0042 and NSF grants IIS-1724101 and IIS-1954778. We would also like to thank Liron Cohen for his implementation of AFS, Wolfgang Hoenig for his implementation of CBS, Glen Wagner for his implementation of M*, and Ilja Ivanashev for his implementation of PR. Finally, we would like to thank the anonymous reviewers for their helpful feedback that improved this work.

Appendix A. WAMPF proofs

Theorem 1. *If we assume PLANIN and GROWANDREPLANIN produce optimal paths in w , a valid path exists, then WAMPF will produce a valid path after a single invocation of RECWAMPF.*

Proof 1. This is a special case of Case 1 or Case 2 in Proof 2; as shown, either π generated on Line 2 is optimal, in which case WAMPF terminates with π as its path, or π will be repaired to generate a valid path. \square

Theorem 2. *If we assume:*

1. *A valid path exists.*
2. *PLANIN and GROWANDREPLANIN produce optimal repairs in their given windows.*
3. *SHOULDQUIT(w_k) does not discard a window w_k with an associated agent set α until $s_k = \Phi(s, \alpha)$, $g_k = \Phi(g, \alpha)$, and the repair is unimpeded by w_k 's restrictions on the state space.*

Then, given sufficient time WAMPF will produce a minimal cost path.

Proof 2.

Lemma 2.1 (*Optimal merged paths are optimal*). *Given two paths, π for agent set α and π' for agent set α' , where π and π' are optimal, $\alpha \cap \alpha' = \emptyset$, and π and π' do not collide with each other, then if π and π' are joined to produce π'' , it follows that $\|\pi''\| = \|\pi\| + \|\pi'\|$, and thus π'' is optimal.*

Proof by contradiction. Consider a case where π'' constructed via the method above is not optimal. That would imply that there exists another, optimal path with the same s and g , π''' , such that:

$$\begin{aligned} \|\pi'''\| &= \|\Phi(\pi''', \alpha)\| + \|\Phi(\pi''', \alpha')\| \\ &< \|\pi''\| \\ &= \|\Phi(\pi'', \alpha)\| + \|\Phi(\pi'', \alpha')\| \\ &= \|\pi\| + \|\pi'\| \\ &\implies \end{aligned}$$

$$\|\Phi(\pi''', \alpha)\| < \|\pi\| \vee \|\Phi(\pi''', \alpha')\| < \|\pi'\|$$

which implies that π or π' are suboptimal, which violates the assumption that π and π' are optimal. \square

Lemma 2.2 (*Unrestricted window searches produce optimal paths*). *Given a joint path π and window w with an associated agent set α is used to repair $\Phi(\pi, \alpha)$, if w contains s and g , the global start and goal for agents α , and a repair within w has just been performed on $\Phi(\pi, \alpha)$ in which w did not constrain the search between s and g , then $\Phi(\pi, \alpha)$ is globally optimal for agent set α .*

We know from the definition of a window that if s and g associated with $\Phi(\pi, \alpha)$ are in w , then they are the s and g used by w . Thus, we know that as the given repair strategy produces an optimal solution in w between w 's s and g , w 's s and g are s and g of $\Phi(\pi, \alpha)$, and the search was not restricted by w , then this path would be optimal even for an arbitrarily large w , and thus $\Phi(\pi, \alpha)$ is globally optimal for agent set α . \square

Lemma 2.3 (*RECWAMPF always grows all windows*). *Given a set of windows W , all $w \in W$ will be enlarged by RECWAMPF to encompass more states.*

At the start of each iteration of RECWAMPF, GROWANDREPLANIN will be invoked $\forall w \in W$ (Lines 6 – 11), by definition causing all windows to be grown, thereby upholding the claim. Some of these windows may be merged with existing windows (Line 12), resulting in a larger, merged windows (Line 24), thereby upholding the claim. Some of these windows may be merged with newly created windows, resulting in larger, merged windows (Line 15), thereby upholding the claim. \square

When RECWAMPF is called, we know a given path is either:

1. Valid and globally optimal, with $W = \emptyset$
2. Invalid and at or below cost of globally optimal solution, with $W = \emptyset$
3. Valid and potentially globally suboptimal, with windows surrounding locally optimal repairs, i.e. $W \neq \emptyset$

We do an analysis of RECWAMPF in these three cases:

1. We invoke RECWAMPF with a valid and globally optimal solution and $W = \emptyset$. Lines 6 – 12 are skipped, as $W = \emptyset$. Lines 13 – 15 are skipped, as no collisions exist. Lines 16 – 17 are skipped, as $W = \emptyset$. Finally, $W = \emptyset$, so $(\pi, 1)$ is returned with π unmodified (Line 18) and thus RECWAMPF returns π , having proved it's a globally optimal solution.
2. We invoke RECWAMPF with an invalid solution at or below joint optimal cost and $W = \emptyset$. Lines 6 – 12 are skipped, as $W = \emptyset$. Lines 13 – 15 create windows and locally repair each collision as they occur along π , merging windows if they overlap. When Lines 13 – 15 are complete, π is a valid but potentially globally suboptimal solution. If Lines 16 – 17 can prove that all windows produces local repairs that are globally optimal, then RECWAMPF returns π , having proven it's a globally optimal solution. Otherwise, RECWAMPF has produced a valid and potentially globally suboptimal solution with windows surrounding locally optimal repairs, the scenario handled by Case 3.
3. We invoke RECWAMPF with a valid and potentially globally suboptimal path π with windows surrounding locally optimal repairs. We know from Lemma 2.3 that these windows will continue to grow with each recursive invocation of RECWAMPF, any overlapping windows will be merged together (Lines 7 – 12), and any new collisions induced by repairs will be encapsulated by a new window and merged with any overlapping existing windows (Lines 13 – 15). Thus, we know in a finite number of recursive invocations of RECWAMPF, every window w , associated with an agent set α , will eventually contain s and g associated with $\Phi(\pi, \alpha)$ such that the window based repair between s and g is not constrained by w . Thus, we know from Lemma 2.2 that the globally optimal path for α from s to g has been proven to be found, and thus w can be removed from W by SHOULDQUIT (Lines 16 – 17). Thus, after a finite number of iterations, RECWAMPF will terminate and from Lemma 2.1 we know that the globally optimal solution has been found.

We know that RECWAMPF will only be invoked in the three cases:

1. π is composed of individually planned, optimal paths (Line 2), and it is collision free. $W = \emptyset$ (Line 3), and so it qualifies for Case 1. Case 1 always terminates after a single invocation of RECWAMPF, and π has been proved to be optimal.
2. π is composed of individually planned, optimal paths (Line 2), and it is *not* collision free. $W = \emptyset$ (Line 3), and so it qualifies for Case 2. Case 2 either terminates after a single invocation of RECWAMPF, and π has been proved to be optimal, or it invokes RECWAMPF in Case 3.
3. π is in the process of being repaired, making it potentially globally suboptimal, and it has an associated window set $W \neq \emptyset$. Case 3 either terminates and π has been proved to be optimal, or it again invokes Case 3. \square

Appendix B. X^* proofs

Theorem 1. During Stage 2, the path between s_{k+1} and s_k extracted from the full joint path π is collision-free.

Proof 1. We know that after a single iteration of RECWAMPF, the full joint path π will be collision free (Appendix A, Theorem 1). Additionally, we know by construction that GROWANDREPLANIN will not be invoked until after the first iteration of RECWAMPF, and it will be invoked on the windows created during the initial iteration of RECWAMPF. All subsequent changes to π via window growth or merging are improvements; X^* 's PLANIN and GROWANDREPLANIN ensure that these improvements do not create collisions in regions of π not encompassed by existing windows by ensuring that the timing of agent movements in regions not encompassed by existing windows remains unchanged. To do this, GROWANDREPLANIN and PLANIN ensure that, after the first iteration of RECWAMPF, repairs within the given window are padded as necessary (Section 5.2). Additionally, WAMPF ensures that if a w_{k+1} overlaps with any another windows, the overlapping windows are merged and repaired via PLANIN rather than invoking GROWANDREPLANIN on w_k (Algorithm 1, Lines 6 – 12), thus ensuring that s_{k+1} cannot be inside another window. As a result, the section of π from s_{k+1} to g_{k+1} must be from a region of π unencompassed by another window, and regions of π unencompassed by another window must be collision free, and thus the section of π from s_{k+1} to g_{k+1} is collision-free. \square

References

- [1] R. Stern, N.R. Sturtevant, D. Atzmon, T. Walker, J. Li, L. Cohen, H. Ma, T.K.S. Kumar, A. Felner, S. Koenig, Multi-agent pathfinding: definitions, variants, and benchmarks, in: Symposium on Combinatorial Search (SoCS), 2019, pp. 151–158.
- [2] D. Silver, Cooperative pathfinding, in: Proceedings of the First AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE'05, AAAI Press, 2005, pp. 117–122.
- [3] G. Wagner, Subdimensional Expansion: A Framework for Computationally Tractable Multirobot Path Planning, Ph.D. thesis, The Robotics Institute Carnegie Mellon University, 2015.

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

[5] L. Cohen, M. Greco, H. Ma, C. Hernández, A. Felner, T.K.S. Kumar, S. Koenig, Anytime focal search with applications, in: *IJCAI*, 2018.

[6] B. DeWilde, A. Mors, C. Witteveen, Push and rotate: a complete multi-agent pathfinding algorithm, *J. Artif. Intell. Res.* 51 (2014) 443–492.

[7] P. Wurman, R. D'Andrea, M. Mountz, Coordinating hundreds of cooperative, autonomous vehicles in warehouses, *AI Mag.* 29 (2008) 9–20.

[8] K. Vedder, E. Schneeweiss, S. Rabiee, S. Nashed, S. Lane, J. Holtz, J. Biswas, D. Balaban, UMass MinuteBots 2017 Team Description Paper, 2017.

[9] K. Vedder, E. Schneeweiss, S. Rabiee, S. Nashed, S. Lane, J. Holtz, J. Biswas, D. Balaban, UMass MinuteBots 2018 Team Description Paper, 2018.

[10] B. Karasfi, H. RasamFard, B. Mostafavi, A. Abbasian, A. Saboohi, A.H. Najafdari, MRL Middle Size Team: Robocup2019 Team Description Paper, 2019.

[11] H. Lu, J. Xiao, Z. Zeng, Q. Yu, K. Huang, W. Dai, Z. Zhou, X. Li, B. Han, B. Chen, P. Zhu, Z. Guo, Z. Zhong, Y. Zhao, Z. Zheng, NuBot Team Description Paper 2019, 2019.

[12] A. Tahir, J. Böling, M.-H. Haghbayan, H.T. Toivonen, J. Plosila, Swarms of unmanned aerial vehicles – a survey, *J. Indust. Inf. Integr.* 16 (2019) 100–106.

[13] B. Araki, J. Strang, S. Pohorecky, C. Qiu, T. Naegeli, D. Rus, Multi-robot path planning for a swarm of robots that can both fly and drive, in: 2017 IEEE International Conference on Robotics and Automation (ICRA), 2017, pp. 5575–5582.

[14] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd edition, Prentice Hall Press, USA, 2009.

[15] K. Vedder, J. Biswas, X*: anytime multiagent path planning with bounded search, in: E. Elkind, M. Veloso (Eds.), *Autonomous Agents and Multiagent Systems*, 2019.

[16] C.S. Ma, R.H. Miller, MILP optimal path planning for real-time applications, in: 2006 American Control Conference, 2006.

[17] J. Berger, A. Boukhtouta, A. Benmoussa, O. Kettani, A new mixed-integer linear programming model for rescue path planning in uncertain adversarial environment, *Comput. Oper. Res.* 39 (2012) 3420–3430.

[18] W.N.N. Hung, X. Song, J. Tan, X. Li, J. Zhang, R. Wang, P. Gao, Motion planning with satisfiability modulo theories, in: 2014 IEEE International Conference on Robotics and Automation (ICRA), 2014, pp. 113–118.

[19] V. Nguyen, P. Obermeier, T.C. Son, T. Schaub, W. Yeoh, Generalized target assignment and path finding using answer set programming, in: *SOCS*, 2017.

[20] M. Hard, N. Nilsson, B. Raphael, A formal basis for the heuristic determination of minimum cost paths, in: *IEEE Transactions on Systems Science and Cybernetics SSC4*, 1968.

[21] D. Harabor, A. Grastien, Online graph pruning for pathfinding on grid maps, in: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI'11*, AAAI Press, 2011, pp. 1114–1119.

[22] J. Reif, Complexity of the generalized Mover's problem, in: J. Schwartz, J. Hopcroft, M. Sharir (Eds.), *Planning, Geometry, and Complexity of Robot Motion*, Ablex Publishing Corp, 1987, pp. 267–281, Ch. 11.

[23] W. van Toll, A.F. Cook, R. Geraerts, A navigation mesh for dynamic environments, *Comput. Animat. Virtual Worlds* 23 (2012) 535–546.

[24] S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics*, MIT Press, Cambridge, Mass, 2005.

[25] J. Lee, W. Yu, A coarse-to-fine approach for fast path finding for mobile robots, in: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009, pp. 5414–5419.

[26] S. Murray, W. Floyd-Jones, Y. Qi, D.J. Sorin, G. Konidaris, Robot motion planning on a chip, in: *Robotics: Science and Systems*, 2016.

[27] S. Karaman, E. Frazzoli, Sampling-based algorithms for optimal motion planning, *Int. J. Robot. Res.* 30 (7) (2011) 846–894.

[28] L.E. Kavraki, P. Svestka, J.C. Latombe, M.H. Overmars, Probabilistic roadmaps for path planning in high-dimensional configuration spaces, in: *IEEE Transactions on Robotics and Automation*, vol. 12, IEEE, 1996, pp. 566–580.

[29] S. Karaman, E. Frazzoli, Sampling-based algorithms for optimal motion planning, in: *International Journal of Robotics Research*, vol. 30, 2011, pp. 846–894.

[30] S. Ghosh, J. Biswas, Joint perception and planning for efficient obstacle avoidance using stereo vision, in: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2017, pp. 1026–1031.

[31] J. Hopcroft, J. Schwartz, M. Sharir, On the complexity of motion planning for multiple independent objects: PSPACE – hardness of the “warehouseman's problem”, in: *The International Journal of Robotics Research*, 1984, pp. 76–88.

[32] D. Harbor, S. Koenig, N. Sturtevant, *AAMAS 2019 Tutorial on Heuristic Search*, 2019.

[33] A. Felner, M. Barer, N.R. Sturtevant, J. Schaeffer, Abstraction-based heuristics with true distance computations, in: *SARA*, 2009.

[34] J. Yu, S.M. LaValle, Planning optimal paths for multiple robots on graphs, in: 2013 IEEE International Conference on Robotics and Automation, 2013, pp. 3612–3617.

[35] P. Surynek, Towards optimal cooperative path planning in hard setups through satisfiability solving, in: *PRICAI*, 2012.

[36] P. Surynek, A. Felner, R. Stern, E. Boyarski, Efficient Sat Approach to Multi-Agent Path Finding Under the Sum of Costs Objective, 2016.

[37] R. Barták, J. Svacara, On SAT-based approaches for multi-agent path finding with the sum-of-costs objective, in: *Symposium on Combinatorial Search (SoCS)*, 2019.

[38] E. Erdem, D.G. Kisa, U. Öztok, P. Schüller, A general formal framework for pathfinding problems with multiple agents, in: *AAAI*, 2013.

[39] A. Kushleyev, M. Likhachev, Time-bounded lattice for efficient planning in dynamic environments, in: 2009 IEEE International Conference on Robotics and Automation, 2009, pp. 1662–1668.

[40] S. Aine, M. Likhachev, Truncated incremental search, *Artif. Intell.* 234 (C) (2016) 49–77.

[41] A. Stentz, I.C. Mellon, Optimal and efficient path planning for unknown and dynamic environments, *Int. J. Robot. Autom.* 10 (1993) 89–100.

[42] S. Koenig, M. Likhachev, D* lite, in: *Proceedings of the AAAI Conference of Artificial Intelligence, AAAI*, 2002, pp. 476–483.

[43] M. Likhachev, D. Ferguson, G. Gordon, A. Stentz, S. Thrun, Anytime Dynamic A*: An Anytime, Replanning Algorithm, 2005, pp. 262–271.

[44] D. Ferguson, A.T. Stentz, Using interpolation to improve path planning: the field D* algorithm, *J. Field Robot.* 23 (2) (2006) 79–101.

[45] X. Sun, W. Yeoh, S. Koenig, Efficient incremental search for moving target search, in: *IJCAI*, 2009.

[46] X. Sun, W. Yeoh, S. Koenig, Generalized fringe-retrieving A*: faster moving target search on state lattices, in: *AAMAS*, 2010.

[47] X. Sun, W. Yeoh, S. Koenig, Moving target D* lite, in: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 2010, pp. 67–74.

[48] J. Svegliato, S. Zilberstein, Adaptive metareasoning for bounded rational agents, in: *CAI-ECAI Workshop on Architectures and Evaluation for Generality, Autonomy and Progress in AI (AEGAP)*, Stockholm, Sweden, 2018.

[49] J. Svegliato, K.H. Wray, S. Zilberstein, Meta-level control of anytime algorithms with online performance prediction, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 2018.

[50] J. Svegliato, P. Sharma, S. Zilberstein, A model-free approach to meta-level control of anytime algorithms, in: *Proceedings of the International Conference on Robotics and Automation*, 2020.

[51] R. Zhou, E.A. Hansen, Multiple sequence alignment using A*, in: *Proceedings of the AAAI Conference of Artificial Intelligence*, 2002.

[52] M. Likhachev, G. Gordon, S. Thrun, ARA*: anytime A* with provable bounds on sub-optimality, in: *Advances in Neural Information Processing Systems 16: Proceedings of the 2003 Conference*, 2003.

[53] S. Aine, P.P. Chakrabarti, R. Kumar, AWA*-a window constrained anytime heuristic search algorithm, in: *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, 2007, pp. 2250–2255.

[54] R. Natarajan, M.S. Saleem, S. Aine, M. Likhachev, H. Choset, A-MHA*: anytime multi-heuristic A*, in: *Twelfth Annual Symposium on Combinatorial Search, SoCS'19*, 2019.

- [55] S. Aine, M. Likhachev, Anytime truncated D*: anytime replanning with truncation, in: Proceedings of the Sixth Annual Symposium on Combinatorial Search, SOCS 2013, Leavenworth, Washington, USA, July 11–13, 2013, 2013.
- [56] V. Narayanan, M. Phillips, M. Likhachev, Anytime safe interval path planning for dynamic environments, in: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2012, pp. 4708–4715.
- [57] M.R.K. Ryan, Exploiting subgraph structure in multi-robot path planning, *J. Artif. Intell. Res.* 31 (2008) 497–542.
- [58] T. Scott Standley, Finding Optimal Solutions to Cooperative Pathfinding Problems, Vol. 1, 2010.
- [59] A. Felner, M. Goldenberg, G. Sharon, R. Stern, T. Beja, N.R. Sturtevant, J. Schaeffer, R.C. Holte, Partial-expansion A* with selective node generation, in: AAAI, 2012.
- [60] M. Goldenberg, A. Felner, R. Stern, G. Sharon, N. Sturtevant, R.C. Holte, J. Schaeffer, Enhanced partial expansion A*, *J. Artif. Intell. Res.* 50 (1) (2014) 141–187.
- [61] M. Erdmann, T. Lozano-Pérez, On multiple moving objects, *Algorithmica* 2 (1) (1987) 477.
- [62] K. Kant, S.W. Zucker, Toward efficient trajectory planning: the path-velocity decomposition, *Int. J. Robot. Res.* 5 (3) (1986) 72–89.
- [63] S. Leroy, J.P. Laumond, T. Simeon, Multiple path coordination for mobile robots: a geometric algorithm, in: Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'99, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999, pp. 1118–1123.
- [64] M. Saha, P. Isto, Multi-robot motion planning by incremental coordination, in: 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2006, pp. 5960–5963.
- [65] M. Crosby, A. Jonsson, M. Rovatsos, A single-agent approach to multiagent planning, in: Proceedings of the Twenty-First European Conference on Artificial Intelligence, ECAI'14, IOS Press, Amsterdam, The Netherlands, 2014, pp. 237–242.
- [66] M. Barer, G. Sharon, R. Stern, A. Felner, Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem, in: Proceedings of the Sixth International Symposium on Combinatorial Search, 2014.
- [67] E. Boyarski, A. Felner, R. Stern, G. Sharon, O. Betzalel, D. Tolpin, S.E. Shimony, ICBS: the improved conflict-based search algorithm for multi-agent pathfinding, in: SOCS, 2015.