Lazy Incremental Search for Efficient Replanning with Bounded Suboptimality Guarantees

Inr. Journal of Robotics Research XX(X):1-29 @The Author(s) 2023 Reprints and permission: sagepub.co.uk/journalsPermissions.nav DOI: 10.1177/ToBeAssigned www.sagepub.com/ SAGE

Jaein Lim¹, Mahdi Ghanei¹, R. Connor Lawson¹, Siddhartha Srinivasa² and Panagiotis Tsiotras¹

Abstract

We present a lazy incremental search algorithm, Lifelong-GLS (L-GLS), along with its bounded suboptimal version, Bounded L-GLS (B-LGLS) that combine the search efficiency of incremental search algorithms with the evaluation efficiency of lazy search algorithms for fast replanning in problem domains where edge-evaluations are more expensive than vertex-expansions. The proposed algorithms generalize Lifelong Planning A* (LPA*) and its bounded suboptimal version, Truncated LPA* (TLPA*), within the Generalized Lazy Search (GLS) framework, so as to restrict expensive edge evaluations only to the current shortest subpath when the cost-to-come inconsistencies are propagated during repair. We also present dynamic versions of the L-GLS and B-LGLS algorithms, called Generalized D* (GD*) and Bounded Generalized D* (B-GD*), respectively, for efficient replanning with non-stationary queries, designed specifically for navigation of mobile robots. We prove that the proposed algorithms are complete and correct in finding a solution that is guaranteed not to exceed the optimal solution cost by a user-chosen factor. Our numerical and experimental results support the claim that the proposed integration of the incremental and lazy search frameworks can help find solutions faster compared to the regular incremental or regular lazy search algorithms when the underlying graph representation changes often.

Keywords

Replanning, Lazy Search, Incremental Search, Bounded Suboptimality

Introduction

Replanning is essential to every decision-making agent operating in a complex or dynamic environment. Agents in such environments often operate with limited computational resources and with partial information, making it prohibitive to construct an accurate model of a complex environment once and for all. Even if constructing an accurate representation of the environment were feasible, it may not be prudent to do so, as the model itself becomes out-of-date in a dynamic environment, and hence any initial plan quickly becomes obsolete or irrelevant. The ability to replan fast in order to adapt to environment changes is crucial for robust and responsive autonomy.

In this paper, we (re)consider path-planning problems on graph representations of a complex and dynamic environment suitable for efficient replanning. To better understand the issues involved, consider the case of a mobile robot navigating through a partially known environment with limited sensing. To produce a motion plan all the way to a distant location, the robot must make assumptions about the feasibility or cost of out-of-sight, or yet unexplored, edges on the underlying search graph, some of which will likely be incorrect. As the robot executes its initial plan and perceives new information about the environment, it may need to update the plan to avoid unknown or dynamic obstacles and leverage discovered shortcuts (Stentz, 1995; Koenig and Likhachev, 2005).

Replanning also appears in settings where the world may be fully known but is too complex to be modeled accurately. In this case, an initial solution on a coarse model may be improved via replanning on refined models as the on-board resources allow. Sampling-based motion planning algorithms offer a case in point. These algorithms generate an increasingly dense series of graphs by sampling the free configuration space in an anytime fashion. Sampling-based planning becomes especially useful in high-dimensional spaces where a global graph representation is neither sufficient nor tractable in order to find a good quality path (Lavalle, 1998; Karaman and Frazzoli, 2011). The rate of convergence toward the optimal solution is determined by the choice of the replanning strategy used to efficiently update the old plan upon refinement by sampling, and by employing an efficient replanning strategy. These have been shown to indeed improve the convergence rate in a variety of problems (Arslan and Tsiotras, 2013; Gammell et al., 2015; Strub and Gammell, 2020b,a). The key idea to efficient replanning is to restrict the replanning routine to only the relevant paths that could possibly improve the current solution upon subsequent graph densifications. As a result, the exploitation (replanning) phase can be shortened, leaving more resources available for the exploration (sampling) phase.

The necessity of replanning is not limited to improving solutions over time; replanning can also help find a solution

²University of Washington, Seattle, WA

Corresponding author:

¹Georgia Institute of Technology, Atlanta, Georgia

Jaein Lim, Georgia Institute of Technology, Atlanta, Georgia 30332. Email: jaeinlim126@gatech.edu

of a complex problem rather quickly by solving the problem sequentially in relaxed settings. For instance, the multiagent path finding problem (MAPF) (Stern et al., 2019) deals with finding collision-free paths of multiple agents on a graph. This problem has intractable complexity (Yu and LaValle, 2013), as the size of the joint planning space grows exponentially with the number of agents. There exist complete and optimal MAPF solvers (Sharon et al., 2015; Boyarski et al., 2015; Felner et al., 2018; Li et al., 2020) which relax the search complexity of the joint planning space by decomposing the problem into multiple single-agent planning problems with incremental inter-agent collision avoiding constraints. This decomposition reduces the exponential complexity of the search at the expense of many replannings of the individual agents, as they repair their old paths to satisfy additionally discovered interagent constraints online. Efficient replanning of individual planning problems, thus, helps find the overall solution faster (Boyarski et al., 2020). Fortunately, in these settings the change in the environment is often relatively small, such that a large part of the previous plan can be reused. Incremental search algorithms that reuse existing search results can facilitate the update of the current plan, resulting in reduced computations compared to searching from scratch (Ramalingam and Reps, 1996; Koenig et al., 2004; Koenig and Likhachev, 2005). These incremental methods minimally propagate the cost inconsistencies induced by changes in the graph to repair the search tree to become consistent again. The efficiency of this minimal inconsistency propagation has been widely manifested in many classical robotics applications (Arslan and Tsiotras, 2013; Gammell et al., 2015; Boyarski et al., 2020).

Unfortunately, existing incremental search algorithms (Koenig et al., 2004; Koenig and Likhachev, 2005; Aine and Likhachev, 2016) are designed specifically to reduce the number of vertex expansions, and they are agnostic to the number of edge evaluations. These methods often evaluate edges excessively to find the new optimal (or bounded suboptimal) solution, incurring a significant computational overhead in problem domains where edge evaluations are expensive. Indeed, in many replanning problems (e.g., navigating in a partially known environment, nonholomic dynamics, sampling-based planning, multiagent path finding problem), evaluating an edge can be more expensive than expanding a vertex. An edge evaluation typically involves multiple collision checks in the configuration space (Kavraki et al., 1996; Lavalle, 1998), solving a two-point boundary value problem (Karaman and Frazzoli, 2010; Webb and van den Berg, 2013), propagating the system dynamics with a closed-loop controller (Kuwata et al., 2009), or checking all-to-all inter-agent collisions (Shome et al., 2020). In this work, we seek to remedy such excessive edge evaluations of prior incremental search algorithms by borrowing ideas from the lazy search framework of (Bohlin and Kavraki, 2000; Cohen et al., 2014; Hauser, 2015; Dellin and Srinivasa, 2016; Mandalika et al., 2018, 2019) in problem domains where edge evaluations are more expensive than vertex expansions.

Consider, for example, a scenario where a distributed team of nonholonomic vehicles needs to replan their trajectories to avoid collisions with one another (see Figure 1). A popular and practical approach to solve this problem is



Figure 1. Multi-agent motion planning scenario for distributed nonholonomic vehicles, where each agent needs to replan asynchronously to find a collision-free trajectory. Edge evaluations entail finding dynamically feasible trajectories, collision checks against static obstacles, and collision checks against all other agents trajectories.

to have each of the agents asynchronously communicate and replan a collision-free trajectory by considering other agents' pre-planned trajectories as dynamic obstacles until a consensus is reached (Liu et al., 2018; Tordesillas and How, 2022). Although such a deconfliction strategy does not guarantee completeness, it has been proven robust in realworld scenarios with imperfect communication links and stochastic environments (Zhou et al., 2022). The approach proposed in this paper can expedite each re-planning instance by deferring expensive edge evaluations, which may involve finding dynamically feasible trajectories, collision checks against static obstacles, and collision checks against all other trajectories.

In this paper we present a class of incremental search algorithms that improve upon existing incremental search algorithms such as Lifelong Planning A* (LPA*) (Koenig et al., 2004) and D*-Lite (Koenig and Likhachev, 2005) and their bounded suboptimal variants, namely, Truncated LPA* (TLPA*) and Truncated D*-Lite (TD*) (Aine and Likhachev, 2016). We address their common drawback, often overlooked in classical planning, namely, excessive edge evaluations, and we propose to generalize the current incremental search algorithms using a lazy search framework to mitigate unnecessary edge evaluations.

This paper extends our previous work in Lim et al. (2022), where we proposed L-GLS, an algorithm that combines the vertex efficiency of LPA* (Koenig et al., 2004) with the edge efficiency of GLS (Bohlin and Kavraki, 2000; Mandalika et al., 2018, 2019) for more efficient replanning in problem domains where edge-evaluations are expensive, by adding several implementation details, and by relaxing the optimality constraint so as to find a bounded suboptimal solution more quickly based on the two relaxation techniques, namely, truncation of inconsistency propagation during repair (Aine and Likhachev, 2016) and by adding an inflation heuristic edge estimate during search (Pohl, 1970). In addition to L-GLS we introduce three additional algorithms (B-LGLS, GD*, B-GD*) and show that by relaxing the optimality constraint of previous lazy incremental search algorithms we can further save

computational resources, especially when the changes are frequent, yet rarely significant. Extensive experimental results are also included to demonstrate the efficacy of the proposed methods for replanning applications in dynamic environments.

The paper is organized as follows: in the next section we provide a taxonomy and a comprehensive literature review of the relevant algorithms, namely, those that utilize either lazy search, bounded suboptimal search, or incremental search, highlighting the connections with the proposed algorithms. Then, we formally introduce the problem formulation and the notation used throughout the paper and we present an optimal lazy incremental search algorithm, namely, Lifelong-GLS (L-GLS). Afterwards we present its bounded suboptimal version, namely, Bounded L-GLS (B-LGLS), which ensures that a replanned solution does not exceed the optimal solution by a given factor. The next two sections present dynamic versions of the lazy incremental search algorithms L-GLS and B-LGLS, called Generalized D* (GD*) and Bounded Generalized D* (B-GD*), respectively, for non-stationary planning queries in a partially known environment. We provide extensive experimental results comparing the proposed algorithms for solving replanning problems in dynamic environments using sampling-based algorithms and navigation problems in dynamic graphs. Finally, we conclude the paper with a summary of our results and their impact on real-world replanning applications.

Contributions

We introduce a class of incremental search algorithms that improve upon existing incremental algorithms by minimizing their often excessive edge evaluations that hinder fast re-planning for nonholonomic systems. Using lazy replanning ideas, our proposed extension saves a significant amount of computation that is wasted when evaluating irrelevant edges. As a result, the proposed algorithms find the solution much faster compared to classical incremental search algorithms, especially in problem domains where edge evaluations are expensive.

We provide theoretical results that guarantee the completeness and correctness of the proposed algorithms. Specifically, we show that the returned solution is bounded above by a user-chosen multiplicative factor of the optimal solution, given the graph. Our numerical experiments support our claim that generalizing incremental search algorithms within the lazy search framework indeed result in much faster replanning.

Related Work

Many prior works have studied incremental graph search and lazy motion planning, but few integrate these approaches into a single algorithm. In this section we offer a comprehensive literature review of these works and situate our own work within the existing taxonomy of lazy and incremental search algorithms.

We classify planners according to their utilization in terms of three important properties:

1. *Lazy search*: The planner uses an admissible edgeweight heuristic to reduce the number of relatively expensive edge evaluations. Lazy planners vary primarily in the manner they select edges for full evaluation.

- 2. Bounded suboptimality: The planner returns a solution with cost within a specified multiplicative bound of the optimal cost. Given a bounded suboptimal planner, an anytime asymptotically optimal planner may be constructed by repeatedly solving the same query, while decreasing the suboptimality bound.
- 3. *Lifelong planning*: The planner reuses prior search results to accelerate subsequent planning queries. For sequences of similar queries, lifelong planning significantly reduces the total cost of solving all queries compared to replanning from scratch. We restrict the definition of lifelong planning to include only planners that can handle arbitrary changes to the underlying graph.

Figure 2 illustrates some of the most popular current planners according to the above classification. The following subsections give a brief summary each class of planners, describing the salient features of each algorithm.

Lazy Search

The earliest and most widely studied category of planners are based lazy search. Lazy edge evaluation as a form of heuristic for faster motion planning stems from the observation that in many practical settings, the time spent on edge evaluation dominates the rest of the computation time (Hauser, 2015). Thus, reducing the number of edge evaluations presents a natural opportunity to speed up planning. When inexpensive admissible heuristics that underestimate the cost of an edge are available, planning may proceed for a while using only these heuristics to guide the search, deferring evaluation until it is necessary to ensure correctness.

The Lazy Probabilistic Roadmap (LazyPRM) by Bohlin and Kavraki (2000) was the first algorithm to introduce laziness, initializing all edges by a heuristic instead of performing actual edge evaluations. Upon completing the graph search, the edges on the candidate solution path are evaluated, after which the graph is updated with their true edge costs, and the search begins anew. The first solution containing only already-evaluated edges is the optimal one.

The Lazy Shortest Path (LazySP) (Haghtalab et al., 2018) framework generalizes LazyPRM by explicitly introducing an *edge selector* function. Upon finding a new candidate solution, the user-specified edge selector function can choose any combination of edges to evaluate, including those not belonging to the candidate path. The authors in (Haghtalab et al., 2018) demonstrate this approach by incorporating prior information about the environment into the edge selector. Evaluating edges on a candidate path in order, from start to goal, as in LazyPRM, is called *forward selection*.

Both LazyPRM and LazySP algorithms search for a complete candidate path to the goal before any edge evaluation is performed, a behavior of *infinite lookahead*. Infinite lookahead is proven to minimize edge evaluations before an optimal path is found, but it requires a high search effort to generate many complete candidate paths. Lazy Receding-Horizon A* (LRA*) (Mandalika et al., 2018) introduces the more general notion of *n-step lookahead*, in which the search is limited to *n* edges beyond the current evaluated frontier, after which a cost-to-go heuristic is used



Figure 2. Taxonomy of relevant motion planning algorithms.

to estimate the remainder of the path cost. The parameter n allows the user to tune the computational effort spent on lazy search.

Generalized Lazy Search (GLS) (Mandalika et al., 2019) further generalizes LazySP by introducing user-specified *evaluation events*. In GLS the search continues until the conditions of a specified event are met, at which point the search immediately stops and edge evaluation commences. Infinite and *n*-step lookahead are examples of evaluation events. The user may also specify an arbitrary edge selector, as in LazySP. The theoretical properties of GLS do not explicitly depend on the underlying search algorithm. The new algorithms presented in this work are instances of the GLS framework, specifying particular underlying search algorithms and evaluation events that preserve correctness, while incorporating additional desirable properties.

The Advanced Informed Trees (AIT*) and Effort Informed Trees (EIT*) (Strub and Gammell, 2020b; Strub and Gammell, 2022) are two recent algorithms that also use lazy edge evaluations to build informed heuristics via a reverse search rooted at the goal vertex. The lazy reverse search gets incrementally informed by subsequent forward searches with actual edge evaluations. While AIT* and EIT* are comparable in spirit and performance to the other lazy algorithms presented here, they notably lack flexibility. Specifically, a reverse heuristic search tree must be connected to a forward search tree to be functional, which limits their applications to undirected planning problems only. In contrast, the GLS framework offers more flexibility in balancing computational effort between search and evaluation in both undirected and directed planning problems. Furthermore, GLS offers the potential to incorporate prior knowledge into both the search and heuristic estimation. Hence, in this work we limit our attention to GLS-based algorithms.

Lastly, it is worth noting that if multiple processors are available for planning, strategies that separate search and evaluation in parallel can significantly accelerate the planning process. The Massively Parallelized Lazy Planning (MPLP) by Mukherjee et al. (2022b) leverages asynchronous graph search, delegating edge evaluations to multiple processors for parallel execution. Similarly, Edge-Based Parallel A* for Slow Evaluations (ePA*SE) (Mukherjee et al., 2022a) parallelizes edge evaluations during the search by distributing them across multiple processors.

However, it is important to acknowledge that parallelized algorithms introduce inherent stochasticity and necessitate the implementation of either a locking mechanism or reliable communication with a termination detection mechanism (Kishimoto et al., 2013; Matocha and Camp, 1998), thereby increasing the overall complexity of the algorithms. In addition, the practical utilization of multiple processors for planning purposes may be limited in many robotic applications due to constraints on available computational resources. Parallel planning is work-inefficient as it speculatively evaluates extra edges, leading to faster goal discovery but wasting CPU resources, which can be costly when tasks compete for CPU time.

Given these considerations, for the remainder of our developments, we will specifically focus on algorithms designed to run on a single processor, acknowledging nonetheless the potential benefits of parallelization in certain problem instances.

Bounded Suboptimality

Strictly optimal motion planning is computationally difficult and may also be unnecessary in many problem instances. Algorithms with bounded suboptimality exploit this fact to achieve dramatic performance improvements by returning suboptimal solutions with cost within a constant factor of the optimal one. Note that, unlike lazy search, the heuristic in these algorithms estimates the cost-to-go rather than the cost of traversing an edge.

Weighted A^* (wA^*) (Pohl, 1970) inflates the given heuristic by a constant factor to bias the search. Increasing the inflation factor prioritizes expansion of vertices close to the goal, since the heuristic comprises a smaller proportion of their estimated solution cost. This modification lends the search a greedy character, leading to faster discovery of an initial solution. The returned solution suboptimality factor equals the heuristic inflation factor. Anytime Repairing A* (ARA*) (Likhachev et al., 2003) tracks and propagates only the inconsistencies introduced by changing the suboptimality bound in Weighted A*, resulting in a more efficient anytime implementation. Note that ARA* is an anytime algorithm but not a lifelong-planning algorithm as it does not handle general graph changes.

Focal search (A_{ε}^*) (Pearl and Kim, 1982) incorporates directly the use of an inadmissible heuristic. Without modification, A* always uses an admissible heuristic to decide vertex expansions, increasing the lower bound on possible solutions at every iteration until an optimal solution is found. A_{ε}^{*} , on the other hand, applies a suboptimality factor to the current A* lower bound, producing a range of solution costs which it considers in focus and a corresponding focal set of in-focus vertices. The inadmissible heuristic can be used to order the expansion of the focal set without violating the suboptimality bound. When the focal set is empty, A_{ε}^* expands the next vertex by an admissible heuristic, raising the lower bound and replenishing the focal set. Since inadmissible heuristics can be much better estimators of path cost compared to admissible heuristics, A_{ε}^{*} can find a solution with far fewer expansions than A^{*}.

Several extensions based on focal search have been explored. Cohen et al. (2018) introduce the Anytime Focal Search (AFS) framework that addresses considerations for transforming focal search into an efficient anytime algorithm, concluding that the use of bounded-cost rather than boundedsuboptimal subsearches yields a more efficient algorithm overall. Explicit Estimation Search (EES) (Thayer and Ruml, 2011) builds on focal search, using an additional heuristic to take the potential cost increment into consideration to remove the negative correlation between OPEN and FOCAL. Multi-Heuristic A* (MHA*) (Aine et al., 2016) provides a framework to combine multiple arbitrarily inadmissible heuristics with a single admissible heuristic for bounded suboptimal search and a finite number of expansions of each vertex. It should be noted that these algorithms solve a single-query problem, and they do not handle general graph changes.

Lifelong Planning

Incremental search algorithms solve a sequence of similar problems efficiently by reusing the previous search results to facilitate finding a new plan quickly (Ramalingam and Reps, 1996; Koenig et al., 2004; Aine and Likhachev, 2013; Aine and Likhachev, 2016). Instead of building a new search tree from scratch, these methods identify the portion of the previous search tree that is inconsistent with the graph changes. Expanding only the inconsistent vertices produces a new optimal solution and a tree consistent with the changed graph. These methods can often expand significantly fewer vertices compared to searching from scratch after each graph change.

The LPA* algorithm of Koenig et al. (2004) utilizes a consistent heuristic to restrict the repair of the search tree to only the relevant part for the current problem, making the search tree consistent with the relevant graph changes. LPA* chooses the inconsistent vertices from a priority queue

similar to the one used in A* search (Hart et al., 1968), such that only the optimal path candidates are chosen to be repaired, in a best-first search fashion. LPA* stores, for each vertex, two distinct cost-to-come values to identify cost inconsistencies, and finds the new optimal solution by propagating the cost-to-come inconsistency upon graph changes. LPA* is provably optimal and efficient, in the sense that no vertex is expanded more than twice; also, it does not make any limiting assumptions about the structure of the underlying graph. It can work with any type of graphs as well as any type of changes (e.g., edge addition/removal, weight increase/decrease). These theoretical properties have made the LPA* algorithm (and, in particular, its dynamic version, D*-Lite) the backbone to numerous applications where efficient replanning is imperative (Koenig and Likhachev, 2005; Arslan and Tsiotras, 2013; Gammell et al., 2015; Koenig and Likhachev, 2002).

Recently, the efficiency of LPA* has been further improved at the expense of optimality, by truncating the inconsistency propagation as early as the current solution is guaranteed to be bounded suboptimal (Aine and Likhachev, 2016). Finding the optimal solution exactly is usually a computationally expensive process, especially when there exists many good optimal path candidates (Pearl and Kim, 1982). If finding a good enough solution instead of the exact solution is satisfactory, then a significant amount of computations spent to find the best among a set of good solutions can be eliminated. Truncated Lifelong Planning A* (TLPA*) (Aine and Likhachev, 2016) achieves this by stopping the repair procedure of LPA* as soon as the current solution is guaranteed not to exceed the best possible solution in the current graph more than a given factor. Compared to LPA*, which uses a binary notion of change to propagate all changes to find the optimal solution regardless of their impact, TLPA* only expands the vertices with significant changes. Hence, TLPA* can expand much fewer vertices compared to LPA*, as it restricts repair to both the relevant and the significantly changed part of the tree.

Unfortunately, both of the LPA* and TLPA* algorithms are vertex-optimal, but they are indifferent to the number of edge evaluations. They often incur a significant amount of unnecessary edge evaluations, slowing down the replanning process. In this paper we address this issue by generalizing these algorithms within the lazy search framework. Before delving into this topic further, let us first discuss the two properties of LPA*/TLPA* that cause unnecessary edge evaluations.

LPA* needs to evaluate all the incident edges when expanding an inconsistent vertex to find a new optimal parent. Similar to A*, which expands the frontier vertices with the lowest cost estimates (the so-called f-value) as a best-first search to grow the optimal search tree, LPA*/TLPA* expands only the inconsistent vertices in a best-first manner to repair the optimal search tree. Hence, when a vertex is expanded to find the new optimal parent, all the values of the incident edges must have been known for the correctness of the algorithm. This A*-like propagation is often referred as a "zero-step lookahead" in the literature (Mandalika et al., 2018, 2019), where a no heuristic estimate of the edge value is used to prioritize the



Figure 3. The propagation of the LPA* search from (a) to (h) while it searches to find the shortest path from the start vertex (•) to the goal vertex (•) given the graph in a dynamic environment. The colored lines are the evaluated edges, where the bold edges are part of the current search tree and the dashed edges are not. The expanded vertices are shown with blue dots. The search begins by expanding the start vertex as shown in (b), and stops after expanding the goal vertex, as shown in (f). The optimal solution is found in the search tree in (g). All the incident edges of the expanded vertices are evaluated regardless of their relevancy to the current problem. After finding the optimal solution, the environment changes in the bottom left corner in (h), where the previous solution does not change. Evaluating this edge does not produce any cost-to-come inconsistencies which are relevant to the new problem instance.

next best vertex^{*}. Regardless of their potential to be a part of the optimal path, in LPA^{*} all the incident edges are evaluated upon expanding a vertex. Among those edges, some edges are not relevant to the current problem, i.e., they are not helpful in constructing a new optimal search tree, and hence evaluating them is not necessary. Figure 3(a)-(f) illustrates this point, where all the incident edges of expanded vertices are evaluated regardless of their relevancy to the current problem.

In order to identify the inconsistent part of the tree in the current graph, LPA* also needs to evaluate all the changed edges before repair propagation begins. When a new edge is evaluated, the vertex corresponding to the child (end) vertex of the edge is updated by assigning a new optimal parent and a new cost-to-come value. This vertex is then inserted in the priority queue if the cost-to-come is altered since last search (and is therefore inconsistent), so that the inconsistency information can propagate to its successor vertices. When the updated vertex in the search tree cannot possibly be on the optimal path, which can be determined using a consistent heuristic cost-to-go, then the vertex is not expanded, as the change cannot help find the optimal path in the current search, and the propagation stops. When the change is not relevant, LPA* never uses this new edge information to find a new optimal solution, and therefore it does not waste computational resources by evaluating the edge. Regardless of whether the propagation continues or stops, LPA* needs to evaluate all changed edges. When there are many changed edges, then evaluating those edges can result in significant computational overhead. Figure 3(h)illustrates this case, where the edge in the bottom left is evaluated as the obstacle disappears, although such an edge is not relevant to the new problem instance.

Our proposed algorithms eliminate these evaluations by restricting edge evaluations only to the relevant and significant part of the graph.

Problem Formulation

We first introduce the common variables and notation that will be used throughout the rest of the paper. Additional notation and definitions will appear in each section of the proposed algorithms, as they become relevant to describe each specific algorithm.

Lazy Edge Weight Function

Let G = (V, E) be a directed graph with vertex set V and edge set E. For a vertex $v \in V$, we denote the set of predecessor vertices of v with pred(v) and the set of its successor vertices with succ(v). For each edge $e \in E$, a weight function $w: E \to (0, \infty]$ assigns a positive real number, including infinity, to this edge, e.g., the distance to traverse this edge, and infinity if traversing the edge is infeasible. We denote an admissible heuristic weight function with $\widehat{w}: E \to (0, \infty)$, which assigns to an edge a non-overestimating positive real number such that $\widehat{w}(e) \leq$ w(e) for all $e \in E$. We assume that evaluating the true weight w is computationally expensive, but the heuristic edge \widehat{w} -value is relatively easy to compute. Let $E_{\text{eval}} \subseteq E$ be the subset of the edge set E, whose actual w-values have been computed in the current graph. We introduce a lazy weight function $\overline{w}: E \to (0, \infty]$ which assigns to an edge its admissible heuristic weight \widehat{w} before evaluation and its true weight w after evaluation, that is,

$$\overline{w}(e) := \begin{cases} w(e), & \text{if } e \in E_{\text{eval}}, \\ \widehat{w}(e), & \text{otherwise.} \end{cases}$$
(1)

^{*}This terminology is different from the one-step lookahead used in LPA* (Koenig et al., 2004), which refers to the rhs-value of an inconsistent vertex being one-step better informed than its g-value with respect to the current graph.

Path Planning Problem

Define a path $\pi = (v_1, v_2, \ldots, v_m)$ on the graph G = (V, E) as an ordered set of distinct vertices $v_i \in V$, $i = 1, \ldots, m$ such that, $v_1 = v_s$ the start vertex and $v_m = v_g$ the goal vertex, and such that for any two consecutive vertices v_i, v_{i+1} , there exists an edge $e = (v_i, v_{i+1}) \in E$. Throughout this paper, we will interchangeably denote a path as the sequence of such edges. With some abuse of notation, we denote the cost of a path as $w(\pi) := \sum_{e \in \pi} w(e)$. Likewise, we denote $\overline{w}(\pi) := \sum_{e \in \pi} \overline{w}(e)$ for the lazy cost estimate of the path π . Let $v_s, v_g \in V$ be the start and goal vertices, respectively, and define a path-planning problem as the problem of finding a path given a tuple $P = (G, w, v_s, v_g)$. Let Π be the set of all finite-cost paths connecting v_s to v_g in G. Then, the shortest path-planning problem seeks to solve

$$\pi^* := \operatorname*{argmin}_{\pi \in \Pi} w(\pi). \tag{2}$$

Given a tuple $P^{\varepsilon} = (G, w, v_s, v_g, \varepsilon)$ with $\varepsilon > 1$, a boundedsuboptimal path-planning problem is the problem of finding a path π^{ε} such that $w(\pi^{\varepsilon}) \leq \varepsilon w(\pi^*)$. In this case, we say that the path π^{ε} is ε -bounded.

Replanning Problem

The optimal replanning problem is the problem of finding the shortest paths $(\pi_i^*)_{i\in\mathbb{N}}$, given a sequence of pathplanning problem instances $(P_i)_{i \in \mathbb{N}}$, where each $P_i =$ $(G_i, w_i, v_{\mathbf{s}_i}, v_{\mathbf{g}_i})$ is a path-planning problem, possibly with different graph, weight, start and goal vertices. The stationary optimal replanning problem finds the shortest paths given a sequence of path-planning problems $(P_i)_{i \in \mathbb{N}}$ where the start and goal vertices remain the same, i.e., $v_{s_i} = v_s$ and $v_{g_i} = v_g$ for all $i \in \mathbb{N}$. The stationary boundedsuboptimal replanning problem solves bounded-suboptimal path-planning problems given a sequence of of path-planning problems $(P_i^{\varepsilon})_{i \in \mathbb{N}}$, where each $P_i^{\varepsilon} = (G_i, w_i, v_s, v_g, \varepsilon_i)$ has the same start and goal vertices. The non-stationary optimal replanning problem is to find the shortest paths given a sequence of path-planning problems with different start vertices but with the same goal vertex. Similarly, the nonstationary bounded-suboptimal replanning problem is to find a bounded suboptimal path given a sequence of pathplanning problems with different start vertices but with the same goal vertex.

Lifelong Lazy Search

In this section we present our first proposed lazy incremental search algorithm, namely, Lifelong-GLS (L-GLS) (Lim et al., 2022) that combines the vertex efficiency of LPA* with the edge efficiency of GLS to solve a stationary optimal replanning problem. L-GLS returns the shortest path given the current graph, regardless of changes in either the graph topology or the edge values. L-GLS maintains a lazy LPA* search tree to update the inconsistencies that arise both from graph changes as well as from edge value discrepancies between the heuristic weight and the actual weight. Unlike LPA* however, L-GLS restricts the edge evaluations only to the optimal path candidates, and unlike GLS, L-GLS uses previous search results to find a new optimal path.

The lazy LPA* search tree is identical to the standard LPA* search tree (Koenig et al., 2004), except that lazy LPA* uses the lazy weight function \overline{w} instead of the actual weight function w.

Lazy LPA* Search Tree

Each vertex of the lazy search tree corresponds to a unique vertex in G, storing the two cost-to-come values to reach that vertex from the start vertex, namely, the g-value and the rhs-value. Similarly to LPA*, these two cost-to-come values are used to identify the inconsistent vertices of the search tree. The g-value is the accumulated cost-to-come by traversing the previous search tree, whereas the rhs-value is the cost-to-come based on the g-value of the predecessor and the \overline{w} -value of the current edge. Hence, the rhs-value is potentially better informed than the g-value, and it is defined as follows:

$$rhs(v) := \begin{cases} 0, & \text{if } v = v_{\text{s}}, \\ \min_{u \in pred(v)}(g(u) + \overline{w}(u, v)), & \text{otherwise.} \end{cases}$$
(3)

A vertex v with g(v) = rhs(v) is called consistent, otherwise it is called inconsistent. An inconsistent vertex is locally overconsistent if g(v) > rhs(v) and locally underconsistent if g(v) < rhs(v). Additionally, the *rhs*value minimizing the predecessor of v is stored as a backpointer, denoted with

$$bp(v) := \underset{u \in pred(v)}{\operatorname{argmin}} (g(u) + \overline{w}(u, v)). \tag{4}$$

Hence, the subpath from v_s to v is readily retrieved by following the backpointers from v to v_s . Note that during implementation, if equation (3) is undefined in case $g(u) + \overline{w}(u, v)$ equals to ∞ for all predecessors of v, we assign rhs(v) to ∞ . In that case, we also assign bp(v) to a null pointer.

A queue Q prioritizes the inconsistent vertices using the key

$$k(v) = [\min(g(v), rhs(v)) + h(v); \min(g(v), rhs(v))],$$
(5)

with lexicographic ordering, where h(v) is a consistent heuristic cost-to-go from v to v_g . Note that during implementation, if both key values are infinity, we assign the minimum value of the key to infinity.

Lifelong-GLS

The proposed algorithm, Lifelong-GLS (L-GLS), consists of two loops: the inner loop and the outer loop. The inner loop is the main search loop which guarantees to return the shortest path in the current graph upon termination. The outer loop updates the current graph heuristically to reflect any external graph changes. The edge evaluations in the inner loop may induce internal changes to the graph. Both external and internal changes are efficiently repaired by a lazy LPA* search tree.

In the inner loop, the lazy LPA* search tree updates the new shortest path from v_s toward v_g in the current graph *G* based on the previous search results. The lazily evaluated LPA* search tree uses the lazy estimates of the edge values when it propagates the inconsistencies to find the shortest subpath to the goal in the current graph. The first unevaluated edge on the shortest subpath returned by the lazy LPA* is then evaluated. If the evaluation results in an inconsistency, then the lazy LPA* search tree is updated and returns the next best subpath for evaluation. If all the edges on the current shortest path to the goal returned by the lazy LPA* have already been evaluated, then L-GLS has found the optimal solution and exits the inner loop.

In the outer loop, L-GLS waits for graph changes. When the edges of the graph G change, L-GLS assigns admissible heuristic values to the changed edges instead of assigning the true edge values by evaluation. Then, the inner loop begins again to search for the new optimal path. As long as the heuristic edge values do not overestimate the true edge values, L-GLS finds the new optimal path, a result we prove in the Section "Analysis of the L-GLS Algorithm" below. By assigning admissible heuristic values to the changed edges, only a subset of the changed edges that could potentially be on the shortest path in the current graph are actually evaluated.

Details of the Algorithm and Main Procedures

In this section we describe step-by-step the procedures used in L-GLS in greater detail. Before the first search begins, all g-values of the vertices are initialized with ∞ , similar to the regular LPA*, and all lazy estimates of edge values are assigned with admissible heuristic values. The first search begins by setting $rhs(v_s) = 0$ and by inserting v_s in the priority queue Q. In the main search loop (Line 35-39 of Algorithm 1) the lazy LPA* search tree is grown with COMPUTESHORTESTPATH(EVENT) until an EVENT is triggered by the expansion of a leaf vertex which just became consistent upon this expansion (Line 16 of Algorithm 1). EVENT is a binary function that returns true by checking whether the current expansion of a vertex satisfies a userdefined condition, for example, the path to the goal is found or the path of a certain depth is found (see Algorithm 2). When EVENT returns true, the subpath to this leaf vertex which triggered the EVENT is returned for evaluation (Line 37 of Algorithm 1). Then, EVALUATEEDGES evaluates the unevaluated edges along this subpath and updates the lazy estimates with their true weights. If the evaluation of an edge results in a different value than the previous lazy estimate, then EVALUATEEDGES returns the edge for the lazy LPA* to update this change accordingly by UPDATEVERTEX. The inconsistency is propagated by the lazy LPA* again until the next time the EVENT is triggered. If the path to the goal is found, and all the edges along this path have been evaluated, then the path is the optimal path in the current graph. This procedure repeats again when the graph changes.

The procedure UPDATEVERTEX is identical to that of the regular LPA*. The only difference is that when UPDATEVERTEX(v) is called, the rhs-value of the vertex v is updated based on the lazy estimate of the incident edge values. This is done to avoid edge evaluations of irrelevant incident edges of v. When a minimizing predecessor is found lazily, then the vertex assigns its backpointer to this predecessor. Finally, the key of this vertex is updated with CALCULATEKEY to be prioritized in the queue Q.

The choice of the EVENT function determines the balance between vertex expansions (Line 13 of Algorithm 1) and edge evaluations (Line 37 of Algorithm 1), as in the GLS

Algorithm 1 Lifelong-GLS (G, v_s, v_g)

1: **procedure** CALCULATEKEY(v) **return** 2: $[\min(g(v), rhs(v)) + h(v); \min(g(v), rhs(v))];$

```
3: procedure UPDATEVERTEX(v)
```

```
4: if v \neq v_s then
```

```
5: bp(v) \leftarrow \operatorname{argmin}_{u \in pred(v)}(g(u) + \overline{w}(u, v));
```

```
6: rhs(v) \leftarrow g(bp(v)) + \overline{w}(bp(v), v);
```

7: **if** $v \in Q$ **then** Q.**REMOVE**(v);

```
8: if g(v) \neq rhs(v) then \triangleright only insert inconsistent v
9: Q.INSERT((v, CALCULATEKEY(v)));
```

10: **procedure** COMPUTESHORTESTPATH(EVENT)

while Q.TOPKEY \prec CALCULATEKEY (v_g) or 11: 12: $g(v_{\rm g}) \neq rhs(v_{\rm g})$ do $u \leftarrow Q.POP();$ 13: \triangleright expand u 14: if g(u) > rhs(u) then \triangleright overconsistent u $g(u) \leftarrow rhs(u);$ 15: **if** EVENT(u) is triggered **then** \triangleright stop search 16: 17: **return** path from $v_{\rm s}$ to u; for all $v \in succ(u)$ do UPDATEVERTEX(v); 18: 19: else \triangleright underconsistent u $g(u) \leftarrow \infty;$ 20: 21: for all $v \in succ(u) \cup \{u\}$ do UPDATEVERTEX(v); 22: **procedure** EVALUATEEDGES($\overline{\pi}$) 23: 24: for each $e \in \overline{\pi}$ do if $e \notin E_{eval}$ then 25: $\overline{w}(e) \leftarrow w(e);$ ▷ update with true weight 26. $E_{\text{eval}} \leftarrow E_{\text{eval}} \cup \{e\};$ 27: if $\overline{w}(e) \neq \widehat{w}(e)$ then return *e*; 28: 29: procedure MAIN()

30: **for all** $e \in E$ **do** $\overline{w}(e) \leftarrow \widehat{w}(e)$;

31: **for all** $v \in V$ **do** $g(v) \leftarrow \infty$; $rhs(v) \leftarrow \infty$;

32: $Q \leftarrow \varnothing; E_{\text{eval}} \leftarrow \varnothing; rhs(v_{\text{s}}) \leftarrow 0;$

33:	UPDATEVERTEX (v_{s}) ;	▷ initialize search
34:	while true do	

```
repeat
                                                                         ⊳ inner loop
35:
                       \overline{\pi} \leftarrow \text{COMPUTESHORTESTPATH}(\text{EVENT});
36:
37:
                       (u, v) \leftarrow \text{EVALUATEEDGES}(\overline{\pi});
                       UPDATEVERTEX(v);
38:
                                                                     ▷ repair search
                 until v_{\rm g} \in \overline{\pi} and \overline{\pi} \subseteq E_{\rm eval}
39:
                 Wait for changes in E;
40:
                 L \leftarrow the set of edges that changed;
41:
                 for all e = (u, v) \in L do
42:
                       \overline{w}(e) \leftarrow \widehat{w}(e);
                                                   ▷ assign heuristic estimate
43:
44:
                       E_{\text{eval}} \leftarrow E_{\text{eval}} \setminus \{e\};
```

```
45: UPDATEVERTEX(v);
```

Algorithm 2 Candidate EVENT Definitions (Mandalika et al., 2019)

1: **procedure** SHORTESTPATH(v)

2: **if** $v = v_{g}$ **then return** true;

3: **procedure** CONSTANTDEPTH(v, depth α)

4: $\overline{\pi} \leftarrow \text{path from } v_{s} \text{ to } v;$

- 5: $\alpha_v \leftarrow$ number of unevaluated edges in $\overline{\pi}$;
- 6: **if** $\alpha_v = \alpha$ **or** $v = v_g$ **then return** true;

framework. For example, if one chooses the SHORTESTPATH as the EVENT, then the algorithm becomes a lifelong version of LazySP (Dellin and Srinivasa, 2016). That is, the lazy LPA* repairs the inconsistent part of the tree all the way up to the goal vertex, and then returns the shortest path to the goal for evaluation. This minimizes the number of edge evaluations of the inner loop. On the other hand, if one chooses the CONSTANTDEPTH of GLS as the EVENT, then the algorithm becomes a lifelong version of LRA* (Mandalika et al., 2018). The number of tree repair steps (vertex expansions) of the lazy LPA* is reduced, since the inconsistency propagation is restricted not to exceed a certain depth before evaluating the edges. This comes at the expense of possibly more edge evaluations. Some candidate EVENT definitions of GLS (Mandalika et al., 2019) are reproduced in Algorithm 2.

Note that the lazy LPA* algorithm maintained under the L-GLS algorithm is almost identical to the regular LPA* algorithm, except for three points. First, the procedure UPDAT-EVERTEX of L-GLS updates an inconsistent vertex with respect to the lazy estimates of the incident edges instead of their actual values. Second, COMPUTESHORTESTPATH is identical to that of LPA* in the way it expands the inconsistent vertices of the lowest key first, that is, when it expands an inconsistent vertex, it makes an underconsistent vertex overconsistent and an overconsistent vertex consistent. The difference is that when the overconsistent vertices are expanded, the EVENT checks whether to continue or stop propagating the inconsistency information to the successor vertices. Finally, when the graph changes, L-GLS updates the changed edge values with admissible heuristic values lazily instead of evaluating them to find the exact values. Hence, the lazy LPA* inherits all the theoretical properties of the regular LPA*, albeit with respect to a different weight, namely, \overline{w} rather than w. This becomes useful when proving the correctness of the algorithm.

The L-GLS algorithm is different from the GLS algorithm in the following points. First, L-GLS stores the previous search results to propagate any inconsistencies efficiently in dynamic graphs, whereas GLS is explicitly designed for a shortest path planning problem in a static graph. L-GLS can possibly evaluate much fewer edges compared to the GLS from scratch, since the search tree of L-GLS is better informed than that of GLS. Second, the exact values for all feasible edges are known a priori in the GLS framework, that is, the heuristic estimates of all feasible edges are accurate. The edge evaluation in GLS only reveals a binary trait of the edge, that is, whether the edge is feasible or not, rather than its exact cost. This is relaxed in L-GLS, such that the edge costs can vary upon evaluation. This relaxation is important in problem domains where obtaining an accurate heuristic edge cost may be difficult. As long as the heuristic edge cost does not overestimate the actual edge cost, L-GLS finds the optimal solution in the current graph.

Note that detecting edge changes in the outer loop (Line 41 of Algorithm 1) may require edge evaluations unless the changes are explicitly provided by the user. However, in our proposed framework, the aim of detecting changes is to restore the admissibility of the lazy edge weight rather than obtaining the exact weight. Hence, we incorporate following strategies to defer full edge evaluations.

First, we only check previously evaluated edges for changes, as previously unevaluated edges already have an admissible lazy edge weight. Moreover, the number of evaluated edges is kept minimal due to the algorithm's structure.

Second, we initially check the end vertices of an edge for changes. Detecting changes in a vertex implies changes in the incident edges, allowing us to identify edge changes without performing a complete edge evaluation.

Furthermore, any necessary edge evaluations made during the change detection process can be directly added to the set of evaluated edges ($E_{\rm eval}$), avoiding redundant computation during the subsequent inner loop search.

Note that any alternative approach for detecting edge changes in the outer loop (Lines 40-45 of Algorithm 1) is permissible as long as the edge estimate update remains admissible and any inconsistent vertex is updated with UPDATEVERTEX.

Illustrative Example

We visualize the search and evaluation process of L-GLS with infinite-lookahead and compare the results to those of LPA* in a dynamic 2D environment, as depicted in Figure 4. An implicit graph is first constructed whose vertices are arranged in a grid and edges are defined for two vertices within a certain distance. The same graph topology is used throughout this example, and only some of the edge weights are changed because of environment changes. Each algorithm finds the shortest path within the graph using the previous search tree.

In the first search, LPA* is equivalent to A*, and L-GLS is equivalent to GLS (See Figure 4.a). LPA* evaluated 390 edges and expanded 45 vertices, whereas L-GLS and GLS both evaluated 61 edges and expanded 314 vertices.

After the first search, only a small part of the environment changes (see Figure 4.b), opening a shorter passage to the goal. LPA* evaluated 18 edges corresponding to the change, then expanded 4 inconsistent vertices to find the shortest path in the current graph. L-GLS evaluated 4 edges that belong to the new shortest path to the goal, and expanded 4 inconsistent vertices. The GLS evaluated 7 edges and expanded 6 inconsistent vertices.

When the environment changed in the irrelevant region (see Figure 4.c), LPA* evaluated 153 edges corresponding to this environment change, but did not expand any vertices, as they were irrelevant to the current search. L-GLS did not do any additional operations to find the shortest path, since the path was already optimal. GLS was identical to the previous search with 7 edge evaluations and 6 vertex expansions.

Finally, the environment changed back to that of the first search episode with the addition of a new obstacle in the irrelevant region. The GLS search was identical to the first search episode with 61 edge evaluations and 314 vertex expansions. On the other hand, L-GLS evaluated only 11 edges and expanded 83 vertices. This is because the majority of the relevant edges were already evaluated during the previous searches, and the majority of the relevant vertices were already consistent. Similarly, LPA* expanded a fewer number of vertices and evaluated a fewer number of edges compared to the first search episode with 273 edge evaluations and 9 vertex expansions, since it utilized



Figure 4. LPA*(top row) and L-GLS with ∞ -lookahead (bottom row) search results to find the shortest path from start vertex(\bullet) to goal vertex(\blacksquare) per environment change, from left to right: (a) first search, (b) second search, (c) third search, and (d) final search. Lines(\nearrow) are the evaluated edges, and dots ($\bullet \bullet$) are the expanded vertices during the current search. Bold lines(\checkmark) are the edges belonging to the current search tree. Blue and red colors represent free space and obstacles, respectively.

the previous search results. The results are summarized in Table 1.

Table 1. Number of edge evaluations and number of vertex expansions for different planners over four consecutive search queries in a dynamic environment of Figure 4.

	LPA*	GLS	L-GLS
First Query			
# Edge Evaluation	390	61	61
# Vertex Expansion	45	314	314
Second Query			
# Edge Evaluation	18	7	4
# Vertex Expansion	4	6	4
ii voiton Enpansion	•	Ũ	•
Third Query			
# Edge Evaluation	153	7	0
# Vertex Expansion	0	6	0
Fourth Quary			
# Edge Evaluation	273	61	11
# Vertex Expansion	0	314	83
	,	514	0.5
Total			
# Edge Evaluation	834	136	76
# Vertex Expansion	81	640	401

Analysis of the L-GLS Algorithm

We next present some of the properties of L-GLS to provide further insights how the algorithm works. We also prove the completeness and correctness of the algorithm, based on the inherited properties from both the LPA* and the GLS algorithms. First, let us state two facts that are invariant during the main search loop.

Invariant 1. The lazy estimate of an edge never overestimates the true edge value, that is, $\overline{w} \leq w$.

Proof. Since $\overline{w}(e) = w(e)$ for all $e \in E_{\text{eval}}$, and $\overline{w}(e) = \widehat{w}(e) \leq w(e)$ for all $e \notin E_{\text{eval}}$, it follows that $\overline{w}(e) \leq w(e)$ for all $e \in E$.

The next invariant property shows that when the lazy LPA* returns the shortest subpath to the goal, then this subpath is optimal. This follows from the theoretical properties of LPA*, which are similar to A*.

Invariant 2. The returned subpath $\overline{\pi}$ from v_s to v of COMPUTESHORTESTPATH(EVENT) is optimal with respect to \overline{w} , that is, $\overline{\pi} = \operatorname{argmin}_{\pi \in \Pi_v} \overline{w}(\pi)$, where Π_v is the set of paths from v_s to v.

Proof. COMPUTESHORTESTPATH with an EVENT returns the path $\overline{\pi}$ from v_s to v, when the triggering vertex vis expanded. Right before the expansion, v was locally overconsistent (Line 14 of Algorithm 1). Theorem 6 of LPA* (Koenig et al., 2004) states that whenever COMPUTESHORTESTPATH selects a locally overconsistent vertex for expansion, then the g-value of v is optimal with respect to \overline{w} .

Next, we show the completeness and correctness of the inner loop of L-GLS. The first theorem is due to the completeness of GLS (Mandalika et al., 2019), which we restate here.

Theorem 3. Let EVENT be a function that, upon halting, ensures that there is at least one unevaluated edge on the current shortest path or that the goal is reached. Then, the inner loop (Line 35-39) of L-GLS implemented with EVENT on a finite graph terminates.

Proof. Suppose the path to the goal has not been evaluated, such that COMPUTESHORTESTPATH(EVENT) returns at least one unevaluated edge to evaluate. Since there is a finite number of edges, the inner loop will eventually terminate.

Theorem 4. *L-GLS finds the shortest path with respect to the current graph when the inner loop (Line 35-39) terminates.*

Proof. Let π^* be the optimal path with respect to w in the current graph, that is, $w(\pi^*) = \min_{\pi \in \Pi} w(\pi)$, where Π is the set of all paths from v_s to v_g . L-GLS terminates its innerloop when $v_g \in \overline{\pi}$ and $\overline{\pi} \subseteq E_{\text{eval}}$, where $\overline{\pi}$ is the output subpath of COMPUTESHORTESTPATH(EVENT). Then, we have

$$\overline{w}(\overline{\pi}) = \sum_{e \in \overline{\pi}} \overline{w}(e) \le \sum_{e' \in \pi^*} \overline{w}(e') \le \sum_{e' \in \pi^*} w(e') = w(\pi^*),$$
(6)

where the first inequality holds by Invariant 2, and the second inequality follows by Invariant 1. Hence, $\overline{w}(\overline{\pi}) \leq w(\pi^*)$, and since $\overline{\pi} \subseteq E_{\text{eval}}$, we have $w(\overline{\pi}) = \overline{w}(\overline{\pi}) \leq w(\pi^*)$. But $w(\pi^*) \leq w(\overline{\pi})$, since π^* is the optimal path. Therefore, $\overline{\pi}$ must be the optimal path with respect to w.

Bounded Suboptimal Lazy Search

The lazy search framework relies on using an admissible heuristic edge evaluation, and its efficiency depends on the accuracy of the heuristic edge values. In other words, the lazy search works best when the error between the actual edge value and the heuristic edge value is minimal. As in any heuristic search, the smaller the gap between the heuristic and the actual value, the faster the search. Similarly to A*, as a perfect heuristic cost-to-go leads to the minimal search (i.e., the minimum number of vertex expansions) (Hart et al., 1968), a perfect heuristic edge value leads L-GLS to the minimal search and evaluation. However, obtaining an accurate and admissible heuristic can be as difficult as solving the original problem itself.

Finding an inadmissible heuristic, on the other hand, is relatively easy, which then often provides a more informed guidance in many problem domains (Pohl, 1970; Likhachev et al., 2008; Barer et al., 2014). One way is to inflate the heuristic value by multiplying with some constant factor greater than 1. The inflation of the heuristic cost-to-go allows a goal-oriented greedy search, and it proves to be faster in many problem domains, as a bounded solution can be found rather quickly at the expense of optimality (Pohl, 1970; Pearl and Kim, 1982; Aine et al., 2016; Likhachev et al., 2003).

Unfortunately, the L-GLS algorithm cannot be directly used with inadmissible edge heuristics, as it will harm the efficiency of the algorithm. If the heuristic estimate is not admissible and there is a discrepancy between the actual value and the heuristic value, then the entire subtree must be repaired upon evaluation to find another admissible path, making the lazy search inefficient. This excessive search is necessary to find the optimal path. Truncation (Aine and Likhachev, 2016) can relax this excessive search, as it restricts replanning to only the part of the tree needed to maintain a bounded suboptimal path. Incorporating truncation in the lazy search framework is the main topic of this section.

To this end, we present a bounded version of L-GLS, called B-LGLS, that uses a lazy version of Truncated LPA* (TLPA*) (Aine and Likhachev, 2016) in place of lazy LPA* to handle more general cases of heuristic edge weight functions. In fact, this extension is beneficial in two

different aspects. First, when an accurate heuristic is difficult to compute, or a more informed heuristic is available but not necessarily admissible, B-LGLS can handle inadmissible heuristics with a guaranteed solution quality. Second, when the change in the environment is not significant, such that complete and thorough replanning may not produce a significantly better quality solution, B-LGLS can reduce the search effort.

In essence, B-LGLS uses inflation of the heuristic edge weight and truncation of the inconsistency propagation. Truncation stops the propagation of inconsistencies of the lazy search tree, and inflation restricts the evaluation of non-promising edges. Both inflation and truncation are two techniques to make the search itself lazy, in the sense that they restrict the repair of the tree to only the part where it is necessary to guarantee that the current solution is boundedsuboptimal.

Next, we first describe the core elements of B-LGLS, namely the lazy-TLPA* search, and then we provide a complete description of B-LGLS which builds on the lazy-TLPA*.

Lazy-TLPA* Search Tree

The lazy-TLPA* search tree is identical to TLPA* except that lazy-TLPA* uses a lazy weight function \widetilde{w} instead of the actual weight function w. The lazy weight function \widetilde{w} : $E \to (0, \infty]$ assigns to an edge its inflated heuristic weight $\varepsilon_1 \widehat{w}$ before evaluation and its true weight w after evaluation, that is,

$$\widetilde{w}(e) := \begin{cases} w(e), & \text{if } e \in E_{\text{eval}}, \\ \varepsilon_1 \widehat{w}(e), & \text{otherwise,} \end{cases}$$
(7)

for some constant factor $\varepsilon_1 \geq 1$, where $\widehat{w} \leq w$ is an admissible heuristic weight, and E_{eval} , as before, is the subset of the edge set E whose actual w-values have been computed in the current graph. We call ε_1 the inflation factor. Hence, the theoretical properties of TLPA* hold exactly the same for the lazy-TLPA*, except that now they hold with respect to \widetilde{w} instead of the actual w.

Similarly to lazy-LPA*, the rhs-value of a vertex v is defined as follows:

$$rhs(v) := \begin{cases} 0, & \text{if } v = v_{\text{s}}, \\ \min_{u \in pred(v)}(g(u) + \widetilde{w}(u, v)), & \text{otherwise.} \end{cases}$$
(8)

Also, the rhs-value minimizing the predecessor of v is stored as a backpointer, denoted with

$$bp(v) := \operatorname*{argmin}_{u \in pred(v)} (g(u) + \widetilde{w}(u, v)). \tag{9}$$

Next, we briefly describe the fundamental properties of TLPA*. A full description can be found in Aine and Likhachev (2016).

TLPA* is identical to LPA*, except that it uses two truncation rules to stop the best-first inconsistency propagation of the LPA* search. Recall that LPA* uses the two cost-to-come values, the g-value, and rhs-value to identify inconsistent vertices. In addition to these, TLPA* uses an additional cost-to-come value, denoted by g^{π} -value, which assigns to a vertex the cost-to-come along the current path following its backpointer to v_s . The g^{π} -value of a vertex reflects the cost-to-come on the current tree with or without repair, whereas the rhs-value of a vertex reflects the cost-to-come after the repair. Hence, the g^{π} -value may be different than the rhs-value. Also, the g^{π} -value may be different than the g-value, the previous cost-to-come value without any repair. The g^{π} -value reflects the in-process cost-to-come value during repair propagation. TLPA* uses the g^{π} -value to decide whether an inconsistency should be further propagated or not when an inconsistent vertex from the priority queue is chosen for expansion. The decision is based on the two truncation rules.

The first truncation rule applies to all inconsistent vertices selected for expansion. When an inconsistent vertex v with the lowest key value from the priority queue is selected for expansion, TLPA* first checks whether the current solution cost is already within the bound even without any further repair. That is, given some truncation factor $\varepsilon_2 \ge 1$, TLPA* checks whether the inequality

$$g^{\pi}(v_{g}) \le \varepsilon_{2}(\min\{g(v), rhs(v)\} + h(v)), \qquad (10)$$

holds for the vertex v selected for expansion. If so, then TLPA* stops expanding any remaining inconsistent vertices and returns an existing bounded suboptimal solution. Since the lowest key value of the priority queue is a lower bound on the optimal solution cost of the current graph, i.e., $\min[g(v), rhs(v)] + h(v) \le g^*(v_g)$, once the inequality (10) is satisfied, then the current solution cost without further repair does not exceed the optimal solution by more than ε_2 , that is, $g^{\pi}(v_g) \le \varepsilon_2 g^*(v_g)$ holds. Thus, truncating the propagation still guarantees that the current solution is within the desired suboptimality bound.

The second truncation rule applies to an underconsistent vertex (i.e., a vertex v with g(v) < rhs(v)) selected for expansion by checking whether

$$g^{\pi}(v) + h(v) \le \varepsilon_2(g(v) + h(v)). \tag{11}$$

If this inequality holds, then the underconsistent vertex is put in the Truncated list instead of the regular priority queue, so that the propagation of inconsistency constrained to this vertex stops. Since g(v) is the previous shortest path cost and h(v) is a consistent heuristic cost-to-go, g(v) + h(v)is a lower bound on the solution cost constrained to pass through v. Thus, if $g^{\pi}(v) + h(v) \leq \varepsilon_2(g(v) + h(v))$, then any vertex v' that uses g(v) to compute its rhs(v') will not underestimate the actual solution cost by more than an ε_2 factor. Truncating this inconsistency propagation guarantees that the solution is within the desired bound as the old path has not deteriorated beyond this bound.

The rest of the TLPA* algorithm is similar to LPA*. It uses the same CALCULATEKEY procedure to compute the key values of inconsistent vertices as in LPA*, and it prioritizes inconsistent vertices according to the lexicographical ordering of the key values. TLPA* uses the same UPDATEVERTEX procedure to update the *rhs*-value, except that TLPA* does not insert an inconsistent vertex into the priority queue if the vertex is already truncated, i.e., is in the Truncated list. The main propagation loop, the procedure COMPUTESHORTESTPATH of LPA*, except that COMPUTEPATH computes the g^{π} -value of a goal and an underconsistent vertex and applies the two truncation rules

mentioned above. The auxiliary routines of TLPA*, namely, computing the g^{π} -value and obtaining the path, are provided in Algorithm 4. For now, line 19 in Algorithm 4 can be ignored. It can be shown that TLPA* guarantees the following properties:

- Bounded Suboptimality: When the COMPUTEPATH function returns the path constructed using OBTAIN-PATH(v_g), the path has cost less than or equal to $\varepsilon_2 g^*(v_g)$ for a chosen $\varepsilon_2 \geq 1$.
- Efficiency: In COMPUTEPATH, no vertex is expanded more than twice.

The Bounded L-GLS Algorithm

Bounded L-GLS (B-LGLS) is identical to L-GLS, except that the heuristic search tree is repaired with the lazy-TLPA* instead of the lazy-LPA*, and the heuristic edge value is inflated. Like L-GLS, B-LGLS contains two loops: an inner loop and an outer loop. In the inner loop, the inconsistency between the heuristic search tree and the actual path is repaired using the lazy-TLPA* to produce a bounded suboptimal path candidate. In the outer loop, any changes are updated lazily with inflated heuristic edge values. The B-LGLS algorithm is presented in Algorithm 3, where the differences between B-LGLS and L-GLS are colored in blue.

Details of the Algorithm and Main Procedures. In this section we describe step-by-step the procedures of B-LGLS in greater detail. Before the first search begins, all the costto-come values are initialized with ∞ , and all the edges are assigned with an inflated heuristic value. The search begins by assigning the *rhs*-value and the g^{π} -value of the start vertex v_s to 0, and inserting v_s in the priority queue via UPDATEVERTEX. In the main search loop (Lines 42-49 of Algorithm 3), the lazy TLPA* search tree is grown with COMPUTEPATH(EVENT) until the first truncation rule is satisfied or an EVENT is triggered. Then, EVALUATEEDGES evaluates the unevaluated edges along the subpath and updates their lazy estimate. The inconsistency induced by the edge evaluation is propagated by the lazy TLPA* until next time the first truncation rule applies, or an EVENT is triggered when expanding an overconsistent vertex. If the path to the goal is found, and all the edges along this path are evaluated, then the path is a bounded suboptimal solution. This procedure repeats again when the graph changes.

When the graph changes, the weights of the changed edges are updated with inflated heuristic values, and the end vertices of the updated edges are put in the priority queue via UPDATEVERTEX. The procedure UPDATEVERTEX is identical to that of LPA*, except that it uses the lazy edge estimate \tilde{w} to update the *rhs*-value and its *bp*, and also it does not put a vertex in the priority queue if that vertex has been already truncated.

Illustrative Example

We visualize the differences between LPA*, TLPA*, L-GLS, and B-LGLS for two consecutive planning problems in a 2D environment shown in Figure 5. The top row shows the first problem instance and the bottom row shows the second problem instance with a new opening in the middle of the map.



Figure 5. First search (top row) and second search (bottom row) to find a bounded suboptimal path from start vertex (•) to goal vertex (•) per environment change with $\varepsilon_1 = 1.2$ and $\varepsilon_2 = 1.2$, from left to right: (a) LPA*, (b) TLPA*, (c) L-GLS, and (d) B-LGLS. Lines (\checkmark) are the evaluated edges during the current search. Blue and red colors represent free space and obstacles, respectively. Bold lines (\checkmark) are the edges belonging to the current search tree and the path is drawn with the boldest lines (\checkmark).

Table 2. Number of edge evaluations, number of vertex expansions, and solution length for different planners over two consecutive search queries in a dynamic environment of Figure 5.

	LPA*	TLPA*(1.44)	GLS	L-GLS	B-LGLS(1.2,1.2)
First Query in scene 1					
# Edge Evaluation	826	826	66	66	68
# Vertex Expansion	58	58	2000	2000	604
Solution Length	0.924	0.924	0.924	0.924	0.956
Second Query in scene 2					
# Edge Evaluation	267	93	35	14	0
# Vertex Expansion	23	1	871	315	1
Solution Length	0.84	0.924	0.84	0.84	0.956
Total					
# Edge Evaluation	1093	919	101	80	68
# Vertex Expansion	81	59	2871	2315	605

Each problem instance was solved using a graph with a fixed topology whose vertices are sampled with a Halton sequence (Halton, 1964) and an edge is defined for two vertices within a certain distance. The vertices are shown with black dots, and the evaluated edges are shown in either blue (feasible) or red (infeasible) lines. Unevaluated edges which belong to the lazy search trees are drawn in black thin lines in Figures 5(c) and 5(d). Unused edges are not drawn. The final search trees grown from the start vertex (green dot) in the lower left toward the goal vertex (magenta dot) in the upper right are drawn with bold blue lines, and the resulting path is marked with thick blue lines, which is best viewed in color. From left to right, the search trees of LPA*, TLPA*, L-GLS, and B-LGLS of each instance are shown.

LPA* and L-GLS find the optimal paths for both of the problem instances, whereas TLPA* and B-LGLS do not waste computational resources to replan, as the current solution is guaranteed to be within a user-given parameter of $\varepsilon_1\varepsilon_2 = 1.44$. Compared to TLPA*, B-LGLS finds the same solution with a significant fewer number of edge evaluations. Compared to L-GLS, B-LGLS saves both edge evaluations and vertex expansions. These results are summarized in Table 2.

Analysis of the Bounded L-GLS Algorithm

The completeness of B-LGLS can be proven using the same argument as in Theorem 3. Next, we prove that B-LGLS returns a bounded suboptimal solution. First, let us consider two useful invariants.

Invariant 5. The lazy estimate of an edge never overestimates the true edge value by more than a factor of ε_1 , that is, $\tilde{w} \leq \varepsilon_1 w$.

Proof. Since $\widetilde{w}(e) = w(e)$ for all $e \in E_{\text{eval}}$, and $\widetilde{w}(e) = \varepsilon_1 \widehat{w}(e) \le \varepsilon_1 w(e)$ for all $e \notin E_{\text{eval}}$, it follows that $\widetilde{w}(e) \le \varepsilon_1 w(e)$ for all $e \in E$.

Invariant 6. When COMPUTEPATH of TLPA* selects a locally overconsistent vertex v for expansion, the path

Algorithm 3 Bounded L-GLS (G, v_s, v_g)

```
1: procedure CALCULATEKEY(v) return
          [\min(g(v), rhs(v)) + h(v); \min(g(v), rhs(v))];
 2:
 3: procedure UPDATEVERTEX(v)
 4:
          if v \neq v_{\rm s} then
                \begin{array}{l} bp(v) \leftarrow \operatorname{argmin}_{u \in pred(v)}(g(u) + \widetilde{w}(u,v));\\ rhs(v) \leftarrow g(bp(v)) + \widetilde{w}(bp(v),v); \end{array}
 5:
 6:
 7:
           if v \in Q then Q.REMOVE(v);
 8:
           if g(v) \neq rhs(v) and v \notin Truncated then
 9:
                Q.INSERT((v, CALCULATEKEY(v)));
10: procedure COMPUTEPATH(EVENT)
           while Q. TOPKEY \prec CALCULATEKEY(v_g) or
11:
12: g(v_{\rm g}) \neq rhs(v_{\rm g}) do
13:
                u \leftarrow Q.POP();
                                                                              \triangleright expand u
                g^{\pi}(v_{g}) \leftarrow \text{COMPUTEGPI}(v_{g});
14:
15:
                if g^{\pi}(v_{g}) \leq \varepsilon_{2}(\min\{g(u), rhs(u)\} + h(u)) then
16:
                     return OBTAINPATH(v_g);
                                                                 \triangleright 1st truncation rule
                if g(u) > rhs(u) then
17:
18:
                     g(u) \leftarrow rhs(u);
19:
                     if EVENT(u) is triggered then
20
                          return OBTAINPATH(u);
21:
                     for all v \in succ(u) do UPDATEVERTEX(v);
22:
                else
                      g^{\pi}(u) \leftarrow \text{COMPUTEGPI}(u);
23:
24:
                     if g^{\pi}(u) + h(u) \leq \varepsilon_2(g(u) + h(u)) then
25:
                          Truncated.INSERT(u);
                                                               ▷ 2nd truncation rule
26:
                     else
27:
                          g(u) \leftarrow \infty;
28:
                          for all v \in succ(u) \cup \{u\} do
                               UPDATEVERTEX(v);
29:
30:
     procedure EVALUATEEDGES(\overline{\pi})
31:
           for each e \in \overline{\pi} do
32:
                if e \notin E_{eval} then
33:
                      \widetilde{w}(e) \leftarrow w(e); E_{\text{eval}} \leftarrow E_{\text{eval}} \cup \{e\};
                     if \widetilde{w}(e) \neq \widehat{w}(e) then return e;
34:
35: procedure MAIN()
36:
           for all e \in E do \widetilde{w}(e) \leftarrow \varepsilon_1 \widehat{w}(e);
           for all v \in V do g(v) \leftarrow \infty; rhs(v) \leftarrow \infty;
37:
           Q \leftarrow \varnothing; Truncated \leftarrow \varnothing;
38:
           E_{\text{eval}} \leftarrow \varnothing; rhs(v_{\text{s}}) \leftarrow 0; g^{\pi}(v_{\text{s}}) \leftarrow 0;
39:
           UPDATEVERTEX(v_s);
                                                                    ▷ initialize search
40:
           while true do
41:
42:
                repeat
                     \overline{\pi} \leftarrow \text{COMPUTEPATH}(\text{EVENT});
43.
44:
                     (u, v) \leftarrow \text{EvaluateEdges}(\overline{\pi});
45:
                      UPDATEVERTEX(v);
46:
                     for all s \in Truncated do
                                                                     ▷ clear Truncated
47.
                          Truncated.REMOVE(s);
                          g^{\pi}(s) \leftarrow \infty;
48:
49:
                until v_{\rm g} \in \overline{\pi} and \overline{\pi} \subseteq E_{\rm eval}
                Wait for changes in E;
50:
51:
                L \leftarrow the set of edges that changed;
                for all e = (u, v) \in L do
52:
53:
                      \widetilde{w}(e) \leftarrow \varepsilon_1 \widehat{w}(e);
                      E_{\text{eval}} \leftarrow E_{\text{eval}} \setminus \{e\};
54:
                     UPDATEVERTEX(v);
55:
```

constrained to pass through v does not overestimate the optimal solution cost by more than an ε_2 factor, that is, $g^{\pi}(v) + h(v) \leq \varepsilon_2 g^*(v_g)$.

Proof. Suppose a locally overconsistent vertex v is selected for expansion, that is, $g(v) \ge rhs(v)$. From Lemma 6 of Aine and Likhachev (2016) we have that $rhs(v) + h(v) \le g^*(v_g)$ and from Lemma 7 of Aine and Likhachev (2016)

Algorithm 4 Lazy Truncated LPA*: auxiliary routines: (Aine and Likhachev, 2016)

```
1: procedure COMPUTEGPI(v)
          cost \leftarrow 0; v' \leftarrow v;
 2:
 3:
          visited \leftarrow \emptyset;
 4:
          while v' \neq v_s do
                if v' \in visited or bp(v') = null then
 5:
 6:
                     q^{\pi}(v) \leftarrow \infty; return
 7:
                else
                    insert v' in visited:
 8.
 9:
                     cost \leftarrow cost + \widetilde{w}(bp(v'), v'); v' \leftarrow bp(v');
10:
          return q^{\pi}(v) \leftarrow \text{cost};
11: procedure OBTAINPATH(v)
12:
          \pi \leftarrow v
13:
          while v \neq v_s do
                if bp(v) \in Truncated then
14:
15:
                     \pi \leftarrow \pi \cup \pi(bp(v)); return \pi;
16:
                \pi \leftarrow \pi \cup bp(v); v \leftarrow bp(v)
          return \pi;
17:
```

we have that $g^{\pi}(v) + h(v) \leq \varepsilon_2(rhs(v) + h(v))$. Hence, the result $g^{\pi}(v) + h(v) \leq \varepsilon_2 g^*(v_g)$ follows.

Now we are ready to prove that B-LGLS is correct, that is, it returns a bounded suboptimal path when the inner loop terminates. This is stated in the next theorem.

Theorem 7. *B-LGLS finds a path that is not longer than the shortest path by more than a factor of* $\varepsilon_1 \varepsilon_2$ *with respect to the current graph when the inner loop (Lines* 42-49) *terminates.*

Proof. Let π^* be the optimal path with respect to w in the current graph, that is, $w(\pi^*) = \min_{\pi \in \Pi} w(\pi)$, where Π is the set of all paths from v_s to v_g . B-LGLS terminates its inner-loop when $v_g \in \overline{\pi}$ and $\overline{\pi} \subseteq E_{\text{eval}}$, where $\overline{\pi}$ is the output subpath of COMPUTEPATH(EVENT). Then, we have

$$\widetilde{w}(\overline{\pi}) = \sum_{e \in \overline{\pi}} \widetilde{w}(e) \le \varepsilon_2 \sum_{e' \in \pi^*} \widetilde{w}(e')$$

$$\le \varepsilon_2 \sum_{e' \in \pi^*} \varepsilon_1 w(e') = \varepsilon_1 \varepsilon_2 w(\pi^*),$$
(12)

where the first inequality holds by Invariant 6, and the second inequality follows by Invariant 5. Hence, $\tilde{w}(\bar{\pi}) \leq \varepsilon_1 \varepsilon_2 w(\pi^*)$ and since $\bar{\pi} \subseteq E_{\text{eval}}$, we have $w(\bar{\pi}) = \tilde{w}(\bar{\pi}) \leq \varepsilon_1 \varepsilon_2 w(\pi^*)$. Therefore, the path cost of $\bar{\pi}$ is not greater than the optimal path cost by a factor $\varepsilon_1 \varepsilon_2$.

Generalized Dynamic Search (GD*)

Mobile robots usually operate in environments that are only partially known or continuously changing. Therefore, they need to quickly replan, as the environment changes or when previously unknown parts of it become known. D*-Lite is an algorithm that is based on similar ideas as LPA* and uses previous search information to quickly replan while the robot is moving. Hence, D*-Lite has been used in many practical applications. It uses the same underlying tree as LPA* to repeatedly compute the shortest path to the goal. Compared to LPA*, the D*-Lite algorithm is primarily aimed to handle environment changes while the robot moves along the previously calculated optimal solution.

In this section, we detail a new algorithm, called Generalized D* (GD*), which generalizes D*-Lite (Koenig and Likhachev, 2005) in the lazy search framework.

The Generalized D* (GD*) algorithm consists of two loops: an inner loop that conducts the actual search given the current graph, and the outer loop that is mainly responsible for perceiving the environment changes and for updating the graph. Note that, similarly to D*-Lite, the search direction in the GD* algorithm is reversed, and the search is started from the goal vertex, progressing towards the start vertex. Thus, the rhs(v) and g(v) values are estimated distances calculated from the vertex v to the goal vertex. Similarly, the heuristic values are computed from the start vertex instead of the goal.

In the inner loop, the candidate shortest path is computed from the goal vertex to the start vertex based on the current heuristic weight values of the GD* tree. Once a candidate solution is found, the first unevaluated edge along the shortest subpath is then evaluated. If the evaluation results in an inconsistency (e.g., collision with an obstacle), then GD* gets updated and returns a new candidate shortest path. Finally, if all the edges along the candidate path are evaluated and they produce no inconsistency, the candidate shortest path is the actual shortest path with respect to the current graph with the true edge values.

Once the shortest path is found, the robot executes the optimal plan and moves along a small segment of the shortest path in the outer loop. The algorithm then scans for graph changes. In case a change is perceived, instead of evaluating the changed edges, the algorithm assigns lazy estimates using the admissible heuristic function for those edges. This ensures that the lazy estimate of the path cost does not overestimate the optimal path cost. Then the inner loop finds the new optimal path with respect to the updated graph. Therefore, only a subset of the changed edges that could be part of the shortest path are evaluated. Algorithm 5 lays out the main procedures for GD*.

Details of the GD* Algorithm

The procedures for the GD* are very similar to the ones of LPA* described in Section "Lifelong Lazy Search" and are summarized briefly below. Given a graph, the g-values of all the vertices are set to ∞ and the lazy edge values are set via an admissible heuristic. As the robot moves along the path, the priorities in the queue must be recalculated to maintain the correct order for expansion. D*-Lite avoids this repeated reordering by adjusting the key values using an additional km variable (Line 2). The km variable alters the priorities of the new vertices when they are added to the queue. Secondly, when a vertex is selected for expansion, its key is recomputed (Line 15) and is expanded only if its key value is indeed a lower bound of the queue. Otherwise, the vertex gets reinserted into the queue with the recomputed key value. Note that the km value is incremented each time the robot moves (Line 49).

The search is initialized by setting the *rhs*-value of the goal vertex to zero and inserting it in the priority queue. Furthermore, the km value is set to zero since the robot has not moved yet. The main search loop of the algorithm (Line 43) expands the search tree via the procedure COMPUTESHORTESTPATH. COMPUTESHORTESTPATH continues expanding the search tree via expanding the frontier vertices until an event is

Algorithm 5 GD* (G, v_s, v_q)

procedure CALCULATEKEY(v) return
$[\min(q(v), rhs(v)) + h(v, v_s) + km, \min(q(v), rhs(v))]$
procedure UPDATEVERTEX(v)
if $v \neq v_a$ then
$bp(v) \leftarrow \arg\min_{u \in succ(v)} g(u) + \overline{w}(u, v);$
$rhs(v) \leftarrow g(bp(v)) + \overline{w}(bp(v), v);$
if $v \in Q$ then Q.Remove(v);
if $a(v) \neq rhs(v)$ then
O.Insert($(v, CALCULATEKEY(v))$);
procedure COMPUTE SHORTESTPATH(event)
while Q . TOPKEY < CALCULATEKEY (v_a) OR
$a(v_a) \neq rhs(v_a)$ do
$k_{old} \leftarrow Q.$ TOPKEY;
$\mathbf{u} \leftarrow Q.POP();$
if $k_{old} < \text{CALCULATEKEY}(u)$ then
Q.INSERT(u, CALCULATEKEY(u));
else if $g(u) > rhs(u)$ then
$g(u) \leftarrow rhs(u);$
if $EVENT(u)$ is triggered then
return path from v_g to u
for all $v \in Pred(u)$ do UPDATEVERTEX (v) ;
else
$a(u) \leftarrow \infty$:
for all $v \in Pred(u) \cup u$ do
UPDATEVERTEX(v);
procedure EVALUATEEDGES($\overline{\pi}$)
for all $e \in \overline{\pi}$ do
if $e \notin E_{eval}$ then
$\overline{w}(e) \leftarrow w(e);$
$E_{eval} \leftarrow E_{eval} \cup \{e\};$
if $\overline{w}(e) \neq \widehat{w}(e)$ then return e
procedure MAIN()
for all $e \in E$ do $\overline{w}(e) \leftarrow \widehat{w}(e)$
for all $v \in V$ do $q(v) \leftarrow \infty$; $rhs(v) \leftarrow \infty$;
$km \leftarrow 0$:
$v_{\text{last}} = v_{\text{current}};$
$E_{eval} \leftarrow \emptyset;$
$Q \leftarrow \varnothing;$
$rhs(v_g) = 0;$
UPDATEVERTEX (v_g) ;
while $v_{\text{current}} \neq v_g$ do
repeat
$\overline{\pi} \leftarrow \text{COMPUTESHORTESTPATH}(event);$
$(u, v) \leftarrow \text{EVALUATEEDGES}(\overline{\pi});$
UPDATE VERTEX (v) ;
until $v_s \in \overline{\pi}$ and $\overline{\pi} \subseteq E_{eval}$
$v_{\text{current}} \leftarrow bp(v_{\text{current}});$
NOVE to v_{current} ;
$\kappa m \leftarrow \kappa m + n(v_{\text{current}}, v_{\text{last}});$
$v_{\text{last}} \leftarrow v_{\text{current}},$
for all $e = (u, u)$ with changed adde cost de
$\overline{w}(\rho) \leftarrow \widehat{w}(\rho)$.
$E_{consl} \leftarrow E_{consl} \setminus \{e\}$
UPDATEVERTEX(v):

triggered after that particular frontier vertex has just become consistent. Similarly to L-GLS (Algorithm 1), the edges of the subpath from the goal to this frontier vertex are returned for evaluation (Line 44).

Similarly to L-GLS (Algorithm 1), all the evaluated edges are then evaluated, and all lazy estimates get updated with their true weight values in EVALUATEEDGES. If the



Figure 6. D*-Lite (top row) and GD* with ∞ -lookahead (bottom row) search results to find the shortest path from start vertex (\bullet) to goal vertex (\bullet) per environment change, from left to right: (a) first search, (b) environment change, (c) second search. The current robot position is shown with \bullet . Lines (\nearrow) are the evaluated edges during the current search. Bold lines(\checkmark) are the edges belonging to the current search tree. Blue and red represents free space and obstacle, respectively.

evaluation results in a different weight value compared to the previous lazy estimate for a particular edge, then this change is reflected via UPDATEVERTEX (Line 45). The vertex that caused the inconsistency is then updated and the changes from the inconsistency are handled by propagating them through the lazy D*-Lite tree COMPUTESHORTESTPATH until interrupted by a triggered EVENT (Line 19). Once a path is found, the goal and all the edges are evaluated, and the path found will be the optimal path to the goal, given the current graph. Similarly to D*-Lite, the robot then progresses towards the goal by traversing the optimal path until it reaches the successor (Line 48) of the current vertex (v_{current}). The km-value also gets updated via the heuristic to reflect the new start position for the algorithm. The algorithm then scans the graph for changes and repeats the same procedures in case of a change in the environment.

The procedure UPDATEVERTEX is similar to that of the regular D*-Lite (Koenig and Likhachev, 2005), the main difference being that the lazy estimate of the edge value (\overline{w}) is used for updating the *rhs*-value of the current vertex. Notice that this is done to avoid evaluation of the irrelevant edges. The backpointer of the current vertex is set to the minimizing successor using \overline{w} . As the last step, and if the vertex is inconsistent, the priority of the current vertex is updated based on the newly calculated key value using CALCULATEKEY.

Note that similarly to the L-GLS algorithm (Algorithm 1) the choice of the EVENT function determines the tradeoff between edge evaluations and vertex expansions, as explained in Section "Lookahead Variation" later on. Also, the procedures of the GD* algorithm are identical to the ones of the D*-Lite algorithm except for three places. First, the GD* algorithm uses the lazy weight function instead of the actual weight function during the UPDATEVERTEX procedure. Second, the GD* algorithm assigns admissible heuristic weights to the edges instead of the actual weights when their values are changed (Line 53). Lastly, when a vertex is made overconsistent in COMPUTESHORTESTPATH, the algorithm checks whether EVENT is triggered to continue or stop propagating the inconsistency information to the predecessors (Line 19).

Table 3. Number of edge evaluations and number of vertexexpansions for different planners over two consecutive searchqueries in the dynamic environment of Figure 6.

	D*-Lite	GD*
First Query		
# Edge Evaluation	1255	154
# Vertex Expansion	113	3305
Second Query		
# Edge Evaluation	82	4
# Vertex Expansion	7	18
Total		
# Edge Evaluation	1347	158
# Vertex Expansion	120	3323

Illustrative Example

We visualize the progression of the D*-Lite and the GD* algorithms with an infinite-step lookahead, as shown in Figure 6. Note that the same implicit graph is used for both algorithms with the only change being the edge weights due to environment changes. Similarly to LPA*, both algorithms find the shortest path during replanning using the previous search tree information.

The first search of D*-Lite is equivalent to the A* search with reversed order (swapping start and goal). The D*-Lite algorithm evaluates 1255 edges and expands 113 vertices, whereas GD* only evaluates 154 edges and expands 3305

vertices. After finding the optimal path, the agent traverses the first segment of the found path, at which point an opening appears, making the path no longer optimal. The environment change causes both D*-Lite and GD* to find the optimal path via replanning. Using the previous search tree information, the D*-Lite algorithm expands 7 vertices and evaluates 82 edges when replanning, while GD* only evaluates 4 edges and expands 18 vertices. Notice that the GD* algorithm finds the same optimal solution with almost an order of magnitude fewer edge evaluations for this simple 2D example. These results are summarized in Table 3.

Generalized Bounded Dynamic Search (B-GD*)

As mentioned earlier, using an inflated heuristic can help find an initial suboptimal solution to the goal faster. Note that this solution is suboptimal due to the heuristic function being inadmissible. Similarly to the B-LGLS algorithm, in this section we combine the idea of an inflated heuristic with truncation rules to introduce a bounded version of the Generalized D* (B-GD*). We will first discuss the lazy version of Truncated D*-Lite (TD*) (Aine and Likhachev, 2016), and then include the inflated heuristic to obtain the bounded lazy D*.

Generalized TD*

The truncation rules described for the Lazy-TLPA* search tree search can be easily applied to D*-Lite. Using truncation, TD* (Aine and Likhachev, 2016) terminates the solution early once the solution cost during the replanning is found to be within the chosen bounds of the optimal cost. Similarly to TLPA*, TD* having a reversed search direction, also uses an extra cost-to-come value $q^{\pi}(s)$ along with the current path to the goal vertex for each vertex $\pi(s)$, in order to check for possible truncation. The truncation conditions are checked in COMPUTESHORTESTPATH (Lines 20 and 31) for terminating the propagation of inconsistency of a vertex early. Furthermore, the vertices are only updated if they are not truncated in the UPDATEVERTEX. The auxiliary routines for TD* are identical to those of the TLPA* (Algorithm 4) except for the reversed search direction. The rest of the TD* algorithm procedures remain the same as D*-Lite.

Bounded Generalized D*

Bounded-GD* (B-GD*) is similar to the GD* except that a lazy TD* tree is used as the underlying search tree instead of lazy D*-Lite and that the heuristic edge values are inflated. Algorithm 6 outlines the B- GD* with the differences with the GD* version (Algorithm 5) outlined in blue. Note that the procedure EVALUATEEDGES for the B-GD* algorithm is the same as the one in the B-LGLS (Algorithm 3) and hence is skipped.

Analysis of the Bounded GD* Algorithm

As TD* inherits the same theoretical properties of TLPA*, B-GD* inherits the same theoretical properties as B-LGLS. Namely, the final solution is bounded suboptimal, and a vertex is expanded at most twice. These properties can be proved similarly to the ones of the TLPA* with minor modifications, thus we omit the details.

```
Algorithm 6 Bounded-GD* (G, v_s, v_a) Subroutines
 1: procedure CALCULATEKEY(v)
         return [\min(g(v), rhs(v)) + h(v, v_s) + km,
 2.
 3: \min(q(v), rhs(v))]
 4: procedure UPDATEVERTEX(v)
         if v \neq v_g then
 5:
 6:
              bp(v) \leftarrow \arg\min_{u \in succ(v)} g(u) + \overline{w}(u, v);
 7:
              rhs(v) \leftarrow g(bp(v)) + \overline{w}(bp(v), v);
         if v \in Q then Q.REMOVE(v)
 8:
         if g(v) \neq rhs(v) and v \notin Truncated then
 9:
10:
              Q.Insert((v, CALCULATEKEY(v)));
11: procedure COMPUTESHORTESTPATH(event)
         while Q.TOPKEY<CALCULATEKEY(v_g) OR
12:
13: g(v_g) \neq rhs(v_g) do
14:
              k_{old} \leftarrow Q.TOPKEY;
              \mathbf{u} \leftarrow Q.POP();
15:
16:
              if k_{old} < CALCULATEKEY(u) then
17:
                   Q.INSERT(u, CALCULATEKEY(u));
18:
              else
19:
                   g^{\pi}(v_s) \leftarrow \text{COMPUTE}\text{GPI}(v_s);
20:
                   if g^{\pi}(v_s) \leq \varepsilon_2(\min\{g(u), rhs(u)\} + h(u)) then
                        return OBTAINPATH(v_s)
21:
22:
                   if q(u) > rhs(u) then
                        g(u) \leftarrow rhs(u);
23:
24:
                        if EVENT(u) is triggered then
25:
                            return path from v_g to u
                        for all v \in Pred(u) do
26:
                            UPDATEVERTEX(v);
27:
28:
                   else
                        g^{\pi}(u) \leftarrow \text{COMPUTEGPI}(u);
29.
30.
                       if g^{\pi}(u) + h(u) \leq
31: \varepsilon_2(min\{g(u), rhs(u)\} + h(u)) then
32:
                            Truncated.INSERT(u);
33:
                        else
                             q(u) \leftarrow \infty;
34:
                            for all v \in Pred(u) \cup u do
35:
36:
                                 UPDATEVERTEX(v);
37: procedure MAIN()
         for all e \in E do \overline{w}(e) \leftarrow \varepsilon_1 \widehat{w}(e);
38:
39:
         for all v \in V do g(v) \leftarrow \infty; rhs(v) \leftarrow \infty;
         Q \leftarrow \varnothing; Truncated \leftarrow \varnothing;
40:
         km \leftarrow 0; v_{\text{last}} \leftarrow v_{\text{current}};
41:
          E_{eval} \leftarrow \emptyset; rhs(v_g) \leftarrow 0; g^{\pi}(v_g) \leftarrow 0;
42:
         UPDATEVERTEX(v_a);
43:
44:
         while v_{\text{current}} \neq v_g do
45:
              repeat
46:
                   \overline{\pi} \leftarrow \text{COMPUTESHORTESTPATH}(event);
47:
                   (u, v) \leftarrow \text{EVALUATEEDGES}(\overline{\pi});
48:
                   UPDATEVERTEX(v);
                   for all s \in Truncated do
49:
50:
                        Truncated.REMOVE(s);
51:
                       g^{\pi}(s) \leftarrow \infty;
              until v_s \in \overline{\pi} and \overline{\pi} \subseteq E_{eval};
52:
53.
              v_{\text{current}} \leftarrow bp(v_{\text{current}}); Move to v_{\text{current}};
54:
              km \leftarrow km + h(v_{\text{current}}, v_{\text{last}});
              v_{\text{last}} \leftarrow v_{\text{current}};
55.
56:
              Perceive the edge changes in E;
```

```
58:
                            \overline{w}(e) \leftarrow \varepsilon_1 \widehat{w}(e);
                            E_{eval} \leftarrow E_{eval} \setminus \{e\};
59·
60:
                            UPDATEVERTEX(v);
```

Illustrative Example

57:

Using the same underlying graph, we visualize the results for D*-Lite, TD*, GD*, and B-GD* for two consecutive

for all e = (u, v) with changed edge cost do



Figure 7. First search (top row) and second search (bottom row) to find a bounded suboptimal path from start vertex (•) to goal vertex (•) per environment change with $\varepsilon_1 = 1.2$ and $\varepsilon_2 = 1.2$, from left to right: (a) D*-Lite, (b) TD*, (c) GD*, and (d) Bounded B-GD*. The current robot position is shown with •. Lines (\checkmark) are the evaluated edges during the current search. Bold lines (\checkmark) are the edges belonging to the current search tree. Blue and red represents free space and obstacle, respectively.

Table 4. Number of edge evaluations, number of vertex expansions, and solution length for different planners over two consecutive search queries in a dynamic environment of Figure 7.

	D*-Lite	TD*(1.2)	GD*	B-GD*(1.2,1.2)
Planning in scene 1				
# Edge Evaluation	832	832	67	78
# Vertex Expansion	64	64	2622	449
Solution Length	0.878	0.878	0.878	0.907
Replanning in scene 2 # Edge Evaluation # Vertex Expansion Solution Length	267 22 0.767	117 4 0.857	15 226 0.767	0 0 0.879
Total # Edge Evaluation # Vertex Expansion	1099 86	949 68	82 2848	78 449

planning steps in the 2D environment in Figure 7. The top row shows the first search. The bottom row shows the second search after the environment changes, which occurs when the robot traverses one segment of the optimal path from the first plan. Both D*-Lite and GD* find the optimal solution, whereas the TD* and B-GD* find a bounded suboptimal path, saving computational resources. Moreover, the GD* and B-GD* algorithms save computational resources via reducing unnecessary edge evaluations compared to the D*-Lite and TD* algorithms. The performance comparison in terms of the number of edge evaluations and vertex expansions for the algorithms are summarized in Table 4.

Experimental Results

In this section we compare the algorithms developed in the previous sections (L-GLS, B-LGLS, GD*, B-GD*) to existing incremental search and lazy search algorithms using various examples. This section is divided into two different parts. The first part tests the algorithms in a stationary replanning problem instance, while the second part compares the various algorithms in a non-stationary replanning problem instance. We also implement the proposed algorithms on a scaled ground vehicle platform to show their experimental performance for robot navigation problems.

Stationary Query Experiment (L-GLS, B-LGLS)

In this section we present numerical results comparing our lazy incremental search algorithms with LPA*, TLPA*, and GLS for a stationary replanning experiment, in which the start and goal of the planning instances are fixed, but either the environment or the graph approximation of the environment changes over time. We demonstrate the efficiency of lazy incremental search algorithms in two different scenarios: first, a scenario in which the path planning problem is solved consecutively in a dynamic environment with a fixed-topology graph, and a second scenario in which graph approximations of a static environment become denser over time via incremental sampling. We consider three different environments for benchmark testing: The Piano Movers' problem in SE(2), a manipulation problem in \mathbb{R}^7 using a PR2 robot (a mobile robot with 7DoF arm), and a non-holonomic motion planning problem in SE(2) using the RACECAR/J platform robot[†]. All the algorithm implementations were done in C++, and the experiments were run on a 2.20 GHz Intel(R) Core(TM) i7-8750H CPU machine with 15.5GB of RAM within a Docker container running Ubuntu 18.04 LTS. The RACECAR/J experiment was run onboard using an NVIDIA Jetson TX2. The source code is available at https://github.com/DCSLgatech/lazy_incremental_search.

Dynamic Environment and Static Graph. We first consider scenarios where the environment is dynamic, and we have available a topologically fixed graph approximating the environment. The search is performed on a pre-built graph with uniformly distributed vertices in the configuration space, in which an edge is added between two vertices if they belong to the first k-nearest neighbors. The graph topology does not change throughout the experiment, and only the edge values change due to underlying environment changes. We repeat the experiment 100 times with randomly sampled start and goal pairs, comparing 5 different planners: LPA*, TLPA*, GLS, L-GLS, and B-LGLS. We report the average number of edge evaluations, the average number of vertex expansions, and the average total runtime of these planners.

Piano Movers' Problem: The purpose of the first experiment is to find the shortest path of a piano in a compact and dynamic apartment environment. For each experiment, we find the shortest paths from the Apartment scenario in OMPL (Sucan et al., 2012) from a start configuration to a goal configuration without colliding with the moving obstacles (see Figure 8). There were six consecutive searches in the environment, where the first search was on scene 1 (Figure 8(a)), the second search was on scene 2 (Figure 8(b)), and the third search was on scene 3 (Figure 8(c)), and the fourth, the fifth, and the sixth searches were again repeated on scene 1, scene 2, and scene 3, respectively. The search was performed on a randomly generated graph with 5,000 randomly sampled vertices and edges connecting two vertices if they belong to the first 20-nearest vertices, analogous to k-PRM (Kavraki et al., 1996) with k = 20. The vertices were sampled from SE(2) bounded by the map size. Figure 8 shows an example of this experiment.

We compare the search results, namely, the number of edge evaluations, the number of vertex expansions, and the total runtime for 5 different algorithms: LPA*, TLPA* (ε_2 =2), GLS (1-step lookahead), L-GLS (1-step lookahead), and B-LGLS (1-step lookahead), $\varepsilon_1 = \sqrt{2}$, $\varepsilon_2 = \sqrt{2}$). We repeated the experiment 100 times with randomly sampled start and goal pairs with 100 different graphs and report the average values. The results are summarized in Figure 9.

The average number of edge evaluations for the regular incremental search algorithms, namely, LPA* and TLPA*, are similar across the search instances. TLPA* reduces the number of vertex expansions in subsequent planning instances compared to LPA*, as it halts the cost inconsistency propagation as early as the current solution

is guaranteed to be bounded suboptimal. Lazy search algorithms including GLS, L-GLS, and B-LGLS reduce edge evaluations by a significant amount compared to the regular incremental search algorithms, and yet these algorithms expand more vertices to repair the cost inconsistencies between the heuristic edge value and the actual value. L-GLS uses fewer edge evaluations and vertex expansions compared to GLS, as it reuses the previous search results. Although Bounded-LGLS uses more vertex expansions in the first search, it uses both fewer edge evaluations and vertex expansions in subsequent planning instances compared to L-GLS, as it halts as soon as the current solution is guaranteed to be bounded suboptimal.

PR2 Robot: The purpose of the second experiment is to find a path for the 7DoF right arm of a PR2 robot from a start configuration to a goal configuration without collision in a dynamic environment where the obstacle moves, as shown in Figure 10.

There were six consecutive searches in the environment, where the first search was on scene 1 (Figure 10(a)), the second search was on scene 2 (Figure 10(b)), and the remaining searches were repeated alternating between the two scenes. We repeated the experiment 100 times with random start and goal pairs and 100 random graphs.

For each experiment, similarly to the Piano Movers' problem, the search was performed on a randomly generated graph with 5,000 vertices. The vertices were randomly uniformly sampled in \mathbb{R}^7 , bounded by the PR2 arm's joint-angle bounds. We implemented LPA*, TLPA*, GLS, L-GLS, and B-LGLS as an OMPL Planner (§ucan et al., 2012) with the MoveIt! interface by Coleman et al. (2014).

We compared the search results, namely, the number of edge evaluations and the number of vertex expansions for 5 different algorithms: LPA*, TLPA* (ε_2 =2), GLS (1-step lookahead), L-GLS (1-step lookahead), and B-LGLS (1-step lookahead), $\varepsilon_1 = \sqrt{2}$, $\varepsilon_2 = \sqrt{2}$). The number of edge evaluations, the number of vertex expansions, and the runtime were recorded for each planning instance, and we took the averages of them over 100 random experiments. The results are summarized in Figure 11.

A similar trend to the Piano Movers' problem is observed for this PR2 robotic arm experiment as well. The regular incremental search algorithms, namely, LPA* and TLPA*, incur unnecessary edge evaluations compared to the lazy search algorithms. LPA* and TLPA* resulted in the same amount of edge evaluations, but TLPA* saves a few more vertex expansions compared to LPA* in consecutive searches. The lazy search algorithms, namely, GLS, L-GLS, and B-LGLS save a significant amount of edge evaluations at the expense of more search. The lazy incremental search algorithms, L-GLS and B-LGLS save both edge evaluations and vertex expansions in consecutive searches compared to GLS, as they use the previous search results. B-LGLS further saves a number of edge evaluations and vertex expansions compared to L-GLS in consecutive searches as it halts the search as soon as the current solution is guaranteed to be bounded suboptimal.

RacecarJ Experimental Platform: The purpose of the third experiment is to find a dynamically feasible path for a

thttps://racecar.mit.edu/



Figure 8. The shortest paths of the Piano Movers' problems in a dynamic environment.



Figure 9. The average number of edge evaluations, the average number of vertex expansions, and the average total runtime of 100 random experiments for the Piano Movers' problem.



Figure 10. The shortest paths of the right arm of PR2 robot for the same query in dynamic environment.



Figure 11. The average number of edge evaluations, the average number of vertex expansions, and the average total runtime taken over 100 random experiments for the PR2 robotic arm problem.

non-holonomic vehicle from a start configuration to a goal configuration without collision in a dynamic environment where the obstacle moves, as shown in Figure 12.

There were six consecutive searches in the environment, where after each search, a randomly chosen 10 percent of the edges changed. Figure 12 shows an example of consecutive searches.

The search was performed on a randomly sampled graph with 500 vertices where the edges were defined for two vertices if they belong to the first 10-nearest vertices. The vertices were sampled using a uniform distribution in \mathbb{R}^3 bounded by the map size, and the edges were connected using Reeds-Shepp curves (Reeds and Shepp, 1990). The length of the Reeds-Shepp curve connecting the two end vertices is the edge weight. We implemented LPA*, TLPA*,



(a) Scene 1

(b) Scene 2

(c) Scene 3





Figure 13. The average number of edge evaluations, the average number of vertex expansions, and the average total runtime for six consecutive searches for the non-holonomic RacecarJ problem.

GLS, L-GLS, and B-LGLS as an OMPL Planner (Sucan et al., 2012) on board a RacecarJ 1/10 scale robotic vehicle equipped with an NVIDIA Jetson TX2 board (Figure 14).



Figure 14. The RacecarJ platform (Karaman et al., 2017).

We compared the search results, namely the number of edge evaluations and the number of vertex expansions for 5 different algorithms: LPA*, TLPA* ($\varepsilon_2=2$), GLS (infinite-step lookahead), L-GLS (infinite-step lookahead), and B-LGLS (infinite-step lookahead, $\varepsilon_1 = \sqrt{2}$, $\varepsilon_2 = \sqrt{2}$). The number of edge evaluations, the number of vertex expansions, and the runtime were recorded for each planning instance, and the average values over 100 random experiments are plotted in Figure 13.

A similar trend as before is observed in this experiment as well. The regular incremental search algorithms, namely, LPA* and TLPA*, incur unnecessary edge evaluations compared to the lazy search algorithms. The lazy search versions of the algorithms, namely, GLS, L-GLS, and B-LGLS save a significant number of edge evaluations at the expense of more search. The lazy incremental search algorithms, L-GLS and B-LGLS save computation time in consecutive searches compared to GLS, as they make use of previous search results. B-LGLS further reduces the number of edge evaluations and vertex expansions compared to L-GLS in consecutive searches, as it halts the search as soon as the current solution is guaranteed to be bounded suboptimal.

Static Environment and Dynamic Graph. In this section, we consider scenarios where the graph representation of a static environment becomes denser over time via incremental sampling, a common strategy of many recent sampling-based planning algorithms (Karaman and Frazzoli, 2011; Arslan and Tsiotras, 2013; Janson et al., 2015; Gammell et al., 2015). A batch of 100 random uniform samples was added per instance, with rejection of infeasible samples due to obstacles. Edges were then added for the k-nearest neighbors. An improved solution is sought incrementally for each instance based on the previous search result[‡].

For each search instance, the time to evaluate the edges, the time to expand the vertices, and the total time to find a new solution was recorded. Each experiment was repeated 50 times, and the average and the standard deviation of the accumulated time of edge evaluations, the accumulated time of vertex expansions, and the solution length as a function of

[‡]Technically, the required k-value for asymptotic optimality would increase beyond any fixed constant value as the number of samples increases. However, in practice, we found that k=20 is sufficiently large to observe asymptotically optimal solutions within the scope of our experiments.

time were obtained. Only experiments with a success ratio exceeding 80% in finding a solution are displayed.

First, we study the effect of the three parameters of B-LGLS, namely the lookahead value, the inflation factor, and the truncation factor.

Lookahead Variation: The lookahead value determines how far the heuristic search tree is repaired before evaluating the edges along the path. Hence, a larger lookahead value delays edge evaluation, but the search tree may become inconsistent often when there is a discrepancy between the heuristic value and the actual value, requiring more vertex expansions. On the other hand, a smaller lookahead value reduces the number of vertex expansions at the expense of more edge evaluations. In all of the Piano Movers' problem instances, the PR2 robotic arm, and the nonholonomic racecar problems, this trend is clear. Figures 15, 16, and 17 show the total edge evaluation time and the total vertex expansion time accumulated as the number of vertices in the current graph for the Piano Movers' problem, the PR2 robotic arm problem, and the non-holonomic racecar problem, respectively. The larger the lookahead value, the lower the accumulated time to evaluate edges. The smaller the lookahead value, the lower the accumulated time to expand vertices. The trade-off between the edge evaluation and the vertex expansion determines the solution convergence rate.

Effect of Inflation Factor: We also investigated the effect of the inflation factor on the solution quality and speed. We inflated the heuristic edge value of the unevaluated edges by a multiplicative factor. The inflation factor has two effects. First, it makes the search greedy, resulting in finding the first solution faster. Second, it defers the evaluation of newly added edges unless such edges are promising to improve the current solution significantly. Although a high inflation factor makes it faster to find a first solution, it may not converge asymptotically to the optimal solution if the current solution is already within the given bound. To address this issue, we dynamically varied the inflation factor as a function of the number of vertices in the graph. We chose $\varepsilon_1 = 1 + 5/q$, where q is the number of vertices (Strub and Gammell, 2020a). Having a dynamic inflation factor finds a first solution fast, and then it asymptotically converges to the optimal solution as ε_1 converges to 1. Figures 18, 19, and 20 show the total edge evaluation time and the total vertex expansion time accumulated as the number of vertices in the current graph for the Piano Movers' problem, the PR2 robotic arm problem, and the non-holonomic racecar problem, respectively. Having an inflation factor greater than 1 reduced the edge evaluation time, but it also incurred more vertex expansion time, generally resulting in slower convergence.

Truncation Factor Variation: The truncation factor determines when to truncate the inconsistency propagation given the changes in the graph. It truncates the inconsistency propagation as soon as the current solution is guaranteed to be bounded suboptimal, saving computational resources until the samples are drawn, which could potentially improve the current solution significantly. Truncation saves both edge evaluation and vertex expansion times; however, as a result of ignoring samples that do not improve the solution by a certain margin, the solution may not converge to the optimal solution. To address this issue, we dynamically reduce the truncation factor as a function of the number of vertices, i.e., $\varepsilon_2 = 1 + 10/q$, where q is the number of vertices. As ε_2 approaches 1, the solution asymptotically converges to the optimal solution. Figures 21, 22, and 23 show the total edge evaluation time and the total vertex expansion time accumulated as the number of vertices in the current graph for the Piano Movers' problem, the PR2 Robot problem, and the non-holonomic RacecarJ problem, respectively. In all three experiments, having a truncation factor greater than 1 helped find the first solution faster. Having a dynamic truncation factor improved the convergence in the Piano Movers' problem and the PR2 Robot arm problem.

Comparison between LPA*, TLPA*, L-GLS, and B-LGLS: Having examined the effects of the lookahead, the inflation factor, and the truncation factor, in this section we compare the performance of our proposed lazy incremental search algorithms to their non-lazy incremental search counterparts. We compared LPA*, TLPA*, L-GLS, and B-LGLS using the same experiments. All of these incremental search algorithms improve the solution as the graph becomes denser as we add 100 samples per iteration. Figures 24, 25 and 26 show the total time to evaluate edges, the total time to expand vertices, and the solution length as a function of time, respectively, for the Piano Movers' problem, the PR2 Robot problem, and the non-holonomic RacecarJ problem.

LPA* and TLPA* do not use heuristic edge values when they propagate the cost inconsistencies, incurring a significant number of edge evaluations. The proposed lazy incremental search algorithms save time when evaluating edges per search iteration of more than one order of magnitude for the Piano Movers' problem and more than two orders of magnitude for the PR2 Robot arm problem and the non-holonomic RacecarJ problem. These savings resulted in a faster solution convergence rate. B-LGLS finds the first solution faster than L-GLS, guided by the inflated heuristic edge weight. As the inflation factor and the truncation factor asymptotically approaches 1, B-LGLS also asymptotically converges to the optimal solution.

Our experimental findings suggest that enhancing the convergence rate of asymptotically optimal sampling-based motion planners is achievable by reducing unnecessary edge evaluations and focusing on repairing only the promising paths. Contemporary state-of-the-art sampling-based motion planners, including ABIT*, AIT*, and EIT* (Strub and Gammell, 2020a; Strub and Gammell, 2022), also leverage these concepts to improve convergence rates, incorporating additional optimization techniques such as informed sampling and bidirectional search. Notably, we refrain from directly comparing our approach to those sampling-based motion planners, which are optimized to solve single-problem instances.

For the remaining part of the paper, we describe nonstationary versions of the proposed lazy incremental search algorithms for replanning when the start vertex changes. We use D*-Lite and Truncated D*-Lite algorithms to find a solution quickly for a mobile robot with expensive edge evaluations and compare the results.

Non-Stationary Query Experiment

In this section, we present numerical simulations of our lazy algorithms and compare them to D*-Lite and TD*. Similar



Figure 15. Lookahead variation for the Piano Movers' problem with $\varepsilon_1 = 1$ and $\varepsilon_2 = 1$.



Figure 16. Lookahead variation for the PR2 Robot problem with $\varepsilon_1 = 1$ and $\varepsilon_2 = 1$.





Figure 17. Lookahead variation for the RacecarJ problem with $\varepsilon_1 = 1$ and $\varepsilon_2 = 1$.

s

120

100

80

60

40

20

(b) Vertex Expansion

Total Vertex Expansion Time

 $\varepsilon_1 = 1.0$

 $\varepsilon_1=1.1$

 $\varepsilon_1 = 1.2$

 $\varepsilon_1 = 1.3$

 $\varepsilon_1 = dynami$

3000 3500 4000 Number of Samples



Figure 18. Inflation factor variation for the Piano Movers' problem.

3000 3500 4000 Number of Samples 4500

to the stationary replanning problem instances, we conduct the experiments in the PR2 Robot simulation environment, the Piano Movers problem, and the RacecarJ environment. For the first two experiments, we generated a random graph with 5,000 vertices each connected to 20 nearest neighbors, and solved the motion planning problem for 100 randomly sampled start and goal configurations similar to the previous stationary case. However, we now change the environment while the robot is moving towards the goal on the previously found solution path, and thus force the robot to replan.

(c) Solution Length

 $\varepsilon_1 = 1.0$

 $\varepsilon_1 = 1.1$

 $\varepsilon_1 = 1.2$

 $\varepsilon_1 = 1.3$

 $\varepsilon_1 = dynan$

s 45

40

Time

Evaluation T

Fotal

Edge E 20

(a) Edge Evaluation



Figure 19. Inflation factor variation for the PR2 Robot problem.





(a) Edge Evaluation

Figure 20. Inflation factor variation for the RacecarJ problem.





Figure 21. Truncation factor variation for the Piano Movers' problem.



Figure 22. Truncation factor variation for the PR2 Robot problem.

Comparison of the Algorithms. For the first experiment, we find the shortest path for a random start and goal configuration in scene 1 (Figure 8(a)) of the Piano Movers' problem. Once an optimal path is computed, the piano is moved towards the goal over one segment along the optimal

path when the environment is changed to the one in scene 2 (Figure 8(b)). The piano again traverses one segment of the newly found path and the environment is changed again into the one in the scene 3 (Figure 8(c)), forcing the planner to

s

Evaluation Time

Total Edge

0.6



Figure 23. Truncation factor variation for the RacecarJ problem.



Figure 24. Comparison between LPA*, TLPA*, B-LGLS for the Piano Movers' problem.



Figure 25. Comparison between LPA*, TLPA*, B-LGLS for the PR2 robotic arm problem.



Figure 26. Comparison between LPA*, TLPA*, B-LGLS for the RacecarJ problem.

find a new solution. We cycle through the environments once more to obtain a total of six consecutive searches.

We compared the results for four different algorithms, namely, D*-Lite, GD* (1-step lookahead), TD* ($\varepsilon_2 = 2$), and B-GD* (1-step lookahead, $\varepsilon_1 = \sqrt{2}, \varepsilon_2 = \sqrt{2}$) for 100

randomly sampled start and goal configurations. The average results from the performance comparison of the different algorithms in terms of the number of edge evaluations, vertex expansions, and runtime for the piano movers problem is shown in Figure 27. As shown in Figure 27, the proposed



Figure 27. The average number of edge evaluations, the average number of vertex expansions, and the average total runtime taken over 100 random experiments in the Piano Movers' environment.



Figure 28. The average number of edge evaluations, the average number of vertex expansions, and the average total runtime taken over 100 random experiments in the PR2 Robot environment.



Figure 29. The average number of edge evaluations, the average number of vertex expansions, and the average total runtime taken over 100 random experiments for the RacecarJ environment.

lazy algorithms significantly reduce the number of edge evaluations. Furthermore, the bounded algorithms generally reduce the number of vertex expansions via early termination in case a bounded suboptimal solution is found. Thus, the average runtime is also reduced for the generalized algorithms.

A similar experiment for the PR2 environment was conducted. Once an optimal path is computed for the environment in scene 1 (Figure 10(a)) the robot moves towards the goal traversing one segment along the optimal path. The environment is then changed to the one shown in scene 2 (Figure 10(b)), and thus the algorithms replan to find a new optimal path to the goal. We cycle through the environments twice more to obtain a total of six consecutive searches. We conducted the experiment for 100 different randomly generated start and goal configurations

for the robot. The comparison of the average performance results solved by the four different algorithms, namely, D*, GD* (1-step lookahead), TD* ($\varepsilon_2 = 2$), and B-GD* (1-step lookahead, $\varepsilon_1 = \sqrt{2}, \varepsilon_2 = \sqrt{2}$) is shown in Figure 28.

It can be seen that the lazy algorithms (GD* and B-GD*) significantly reduce the number of edge evaluations, and hence the total runtime. The bounded version of the algorithm (B-GD*) further reduces the total runtime via terminating upon finding a bounded suboptimal solution. Note that for the PR2 experiment, the GD* algorithm results in an approximately 7X speed-up to find the first solution.

Lastly, a similar experiment was conducted in the RacecarJ environment (Figure 12). We constructed the graph via 500 randomly sampled vertices, each having edges connected to the 10 nearest neighbors via Reeds-Shepp curves similar to the stationary case. Once an optimal path is

found, the RacecarJ traverses one segment along the optimal path and we randomly change 10 percent of the edges in the graph simulating a dynamic environment. We repeat this to obtain a total of six consecutive searches. We gathered the experimental data using 100 randomly generated start and goal configurations.

The comparison of the average performance results solved by the four different algorithms, namely, D*, GD* (infinitestep lookahead), TD* ($\varepsilon_2 = 2$), and B-GD* (infinite-step lookahead, $\varepsilon_1 = \sqrt{2}, \varepsilon_2 = \sqrt{2}$) is shown in Figure 29. As expected, the generalized algorithms, namely, GD* and B-GD* significantly reduce the number of edge evaluations in the graph in all six consecutive searches with the trade-off of increasing the number of vertex expansions (conducting more search). Note that the runtime is also significantly reduced for the first search in this environment.

Parametric Study. The choice of the lookahead value can have a significant effect on the performance of the algorithms. We demonstrate this for the dynamic algorithms for some sample experimental cases. The comparison between 1-step lookahead and infinite-step lookahead for the PR2 experiments are shown in Figure 30. As shown in Figure 30, an infinite-step lookahead reduces edge evaluations, while a 1-step lookahead minimizes vertex expansions within the lazy framework. For this particular problem, the infinite-step lookahead is about 5X faster than the 1-step lookahead. However, the optimal lookahead value is dependent on the problem.



Figure 30. Performance comparison for the extreme lookahead values $(1-\text{step and }\infty)$.

Figure 31 demonstrates the effect of different values of the truncation and inflation factors in the PR2 simulation environment. As expected, increasing the bounds on the suboptimal solution generally leads to the reduction of both the number of vertex expansions and edge evaluations, resulting in reduced total solution time.

Conclusion

We have presented a novel replanning framework that combines the vertex efficiency of incremental search methods with the edge efficiency of lazy search methods. Within the proposed lazy incremental search framework, we have presented four different algorithms: L-GLS, B-LGLS, GD*, and B-GD*. L-GLS and B-LGLS solve a sequence of planning problems with fixed start and goal vertices using the previous search results to efficiently restrict unnecessary edge evaluations. GD* and B-GD* solve a sequence of planning problems with moving start vertices, generalizing D*-Lite and TD* within the lazy search framework. We



Figure 31. Performance comparison for the different truncation (ε_2) and inflation factors (ε_1) .

have proven that these algorithms are complete and correct in returning a bounded suboptimal solution given a graph change. In our numerical experiments, it was shown that our generalization significantly improves the performance of the classical incremental search algorithms via reducing unnecessary edge evaluations, speeding up replanning. The proposed improvements enable classical algorithms to be applied in a broader range of applications where edge evaluation is expensive.

Acknowledgements

This work has been supported by ARL under DCIST CRA W911NF-17-2-0181 and SARA CRA W911NF-20-2-0095, and by NSF under award IIS-2008686.

References

- Aine S and Likhachev M (2013) Aytime truncated d*: Anytime replanning with truncation. In: *International Symposium on Combinatorial Search*, volume 4. AAAI Press, pp. 2–10.
- Aine S and Likhachev M (2016) Truncated incremental search. Artificial Intelligence 234: 49 – 77. DOI:https://doi.org/10. 1016/j.artint.2016.01.009.
- Aine S, Swaminathan S, Narayanan V, Hwang V and Likhachev M (2016) Multi-Heuristic A*. *The International Journal of Robotics Research* 35: 224 – 243.
- Arslan O and Tsiotras P (2013) Use of relaxation methods in sampling-based algorithms for optimal motion planning. In: *IEEE International Conference on Robotics and Automation*. Karlsrühe, Germany, pp. 2421–2428.
- Barer M, Sharon G, Stern R and Felner A (2014) Suboptimal variants of the conflict-based search algorithm for the multiagent pathfinding problem. In: Seventh Annual Symposium on Combinatorial Search. Prague, Czech Republic, pp. 19–27.
- Bohlin R and Kavraki LE (2000) Path planning using lazy PRM. In: *IEEE International Conference on Robotics and Automation*, volume 1. San Francisco, CA, pp. 521–528.
- Boyarski E, Felner A, Harabor D, Stuckey P, Cohen L, Li J and Koenig S (2020) Iterative-Deepening Conflict-Based Search. In: *International Joint Conference on Artificial Intelligence*, volume 4. Yokohama, Japan, pp. 4084–4090. DOI:10.24963/ ijcai.2020/565.
- Boyarski E, Felner A, Stern R, Sharon G, Tolpin D, Betzalel O and Shimony E (2015) ICBS: Improved conflict-based search algorithm for multi-agent pathfinding. In: *Proceedings of the International Conference on Artificial Intelligence*. Buenos Aires, Argentina, pp. 740–746.

- Cohen B, Phillips M and Likhachev M (2014) Planning single-arm manipulations with n-arm robots. In: *Proceedings of Robotics: Science and Systems*. Berkeley, CA, pp. 226–227.
- Cohen L, Greco M, Ma H, Hernández C, Felner A, Kumar TKS and Koenig S (2018) Anytime focal search with applications.
 In: *International Joint Conference on Artificial Intelligence*. Stockholm, Sweden, pp. 1434–1441.
- Coleman D, Şucan IA, Chitta S and Correll N (2014) Reducing the barrier to entry of complex robotic software: a MoveIt! case study. *Journal of Software Engineering for Robotics* 5(1): 3–16. DOI:10.6092/JOSER_2014_05_01_p3.
- Dellin CM and Srinivasa SS (2016) A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors. In: *Proceedings of the International Conference on Automated Planning and Scheduling*. London, UK, pp. 459–467.
- Felner A, Li J, Boyarski E, Ma H, Cohen L, Kumar TKS and Koenig S (2018) Adding heuristics to conflict-based search for multi-agent path finding. In: *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28. Delft, The Netherlands, pp. 83–87.
- Gammell JD, Srinivasa SS and Barfoot TD (2015) Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In: *IEEE International Conference on Robotics and Automation.* Seattle, WA, pp. 3067–3074.
- Haghtalab N, Mackenzie S, Procaccia AD, Salzman O and Siddhartha SS (2018) The provable virtue of laziness in motion planning. In: *Proceedings of the 28th International Conference on Automated Planning and Scheduling*. Delft, The Netherlands, pp. 6161–6165.
- Halton JH (1964) Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the ACM* 7(12): 701–702. DOI:10.1145/355588.365104.
- Hart PE, Nilsson NJ and Raphael B (1968) A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2): 100– 107. DOI:10.1109/TSSC.1968.300136.
- Hauser K (2015) Lazy collision checking in asymptotically-optimal motion planning. In: *IEEE International Conference on Robotics and Automation*. Seattle, WA, pp. 2951–2957. DOI: 10.1109/ICRA.2015.7139603.
- Janson L, Schmerling E, Clark A and Pavone M (2015) Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *The International Journal of Robotics Research* 34(7): 883–921.
- Karaman S, Anders A, Boulet M, Connor J, Gregson K, Guerra W, Guldner O, Mohamoud M, Plancher B, Shin R and Vivilecchia J (2017) Project-based, collaborative, algorithmic robotics for high school students: Programming self-driving race cars at mit. In: *IEEE Integrated STEM Education Conference*. Princeton, NJ, pp. 195–203. DOI:10.1109/ISECon.2017. 7910242.
- Karaman S and Frazzoli E (2010) Optimal kinodynamic motion planning using incremental sampling-based methods. In: *Proceedings of the 49th IEEE Conference on Decision and Control.* Atlanta,GA, pp. 7681–7687. DOI:10.1109/CDC. 2010.5717430.
- Karaman S and Frazzoli E (2011) Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research* 30(7): 846–894.

- Kavraki LE, Svestka P, Latombe J and Overmars MH (1996) Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* 12(4): 566–580. DOI:10.1109/70.508439.
- Kishimoto A, Fukunaga A and Botea A (2013) Evaluation of a simple, scalable, parallel best-first search strategy. *Artificial Intelligence* 195: 222–248.
- Koenig S and Likhachev M (2002) D* lite. In: Eighteenth National Conference on Artificial Intelligence. Edmonton, Canada. ISBN 0262511290, p. 476–483.
- Koenig S and Likhachev M (2005) Fast replanning for navigation in unknown terrain. *IEEE Transactions on Robotics* 21(3): 354– 363. DOI:10.1109/TRO.2004.838026.
- Koenig S, Likhachev M and Furcy D (2004) Lifelong planning A*. Artificial Intelligence 155(1): 93 – 146. DOI:https://doi.org/10. 1016/j.artint.2003.12.001.
- Kuwata Y, Teo J, Fiore G, Karaman S, Frazzoli E and How JP (2009) Real-time motion planning with applications to autonomous urban driving. *IEEE Transactions on Control Systems Technology* 17(5): 1105–1118.
- Lavalle SM (1998) Rapidly-exploring random trees: A new tool for path planning. Technical report, Computer Science Department, Iowa State University.
- Li J, Gange G, Harabor D, Stuckey PJ, Ma H and Koenig S (2020) New techniques for pairwise symmetry breaking in multi-agent path finding. *Proceedings of the International Conference on Automated Planning and Scheduling* 30(1): 193–201.
- Likhachev M, Ferguson D, Gordon G, Stentz A and Thrun S (2008) Anytime search in dynamic graphs. *Artificial Intelligence* 172(14): 1613 – 1643. DOI:https://doi.org/10.1016/j.artint. 2007.11.009.
- Likhachev M, Gordon GJ and Thrun S (2003) ARA*: Anytime a* with provable bounds on sub-optimality. In: *Neural Information Processing Systems*. Vancouver, Canada.
- Lim J, Srinivasa SS and Tsiotras P (2022) Lazy Lifelong Planning for Efficient Replanning in Graphs with Expensive Edge Evaluation. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Kyoto, Japan, pp. 8778–8783.
- Liu S, Mohta K, Atanasov N and Kumar V (2018) Towards searchbased motion planning for micro aerial vehicles. *arXiv e-prints* : arXiv:1810.03071.
- Mandalika A, Choudhury S, Salzman O and Srinivasa SS (2019) Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles. In: *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29. Berkeley, CA, pp. 745–753.
- Mandalika A, Salzman O and Srinivasa SS (2018) Lazy receding horizon A* for efficient path planning in graphs with expensive-to-evaluate edges. In: *Proceedings of the International Conference on Automated Planning and Scheduling*. Delft, Netherlands, pp. 476–484.
- Matocha J and Camp T (1998) A taxonomy of distributed termination detection algorithms. *Journal of Systems* and Software 43(3): 207–221. DOI:https://doi.org/10.1016/ S0164-1212(98)10034-1.
- Mukherjee S, Aine S and Likhachev M (2022a) ePA*SE: Edgebased parallel A* for slow evaluations. In: *International Symposium on Combinatorial Search*, volume 15. AAAI Press, pp. 136–144.

- Mukherjee S, Aine S and Likhachev M (2022b) MPLP: Massively parallelized lazy planning. *IEEE Robotics and Automation Letters* 7(3): 6067–6074.
- Pearl J and Kim JH (1982) Studies in semi-admissible heuristics. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 4(4): 392–399. DOI:10.1109/TPAMI. 1982.4767270.
- Pohl I (1970) Heuristic search viewed as path finding in a graph. Artificial Intelligence 1(3): 193–204. DOI:https://doi.org/10. 1016/0004-3702(70)90007-X.
- Ramalingam G and Reps T (1996) An incremental algorithm for a generalization of the shortest-path problem. *Journal of Algorithms* 21: 267–305.
- Reeds JA and Shepp LA (1990) Optimal paths for a car that goes both forwards and backwards. *Pacific Journal of Mathematics* 145(2): 367 – 393.
- Sharon G, Stern R, Felner A and Sturtevant NR (2015) Conflictbased search for optimal multi-agent pathfinding. Artificial Intelligence 219: 40–66.
- Shome R, Solovey K, Dobson A, Halperin D and Bekris KE (2020) dRRT*: Scalable and informed asymptotically-optimal multirobot motion planning. *Autonomous Robots* 44(3): 443–467. DOI:10.1007/s10514-019-09832-9.
- Stentz A (1995) The focussed D* algorithm for real-time replanning. In: Proceedings of the International Joint Conference on Artificial Intelligence, volume 2. Montreal, Quebec, Canada, pp. 1652–1659.
- Stern R, Sturtevant NR, Felner A, Koenig S, Ma H, Walker TT, Li J, Atzmon D, Cohen L, Kumar TKS, Barták R and Boyarski E (2019) Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. In: *Twelfth Annual Symposium on Combinatorial Search*. Vienna, Austria, pp. 151–158.
- Strub MP and Gammell JD (2020a) Adaptively informed trees (AIT*): Fast asymptotically optimal path planning through adaptive heuristics. In: *IEEE International Conference on Robotics and Automation*. pp. 3191–3198. DOI:10.1109/ ICRA40945.2020.9197338.
- Strub MP and Gammell JD (2020b) Advanced BIT* (ABIT*): Sampling-based planning with advanced graph-search techniques. In: *IEEE International Conference on Robotics and Automation.* Paris, France, pp. 130–136.
- Strub MP and Gammell JD (2022) Adaptively informed trees (AIT*) and effort informed trees (EIT*): Asymmetric bidirectional sampling-based path planning. *The International Journal of Robotics Research* 41(4): 390–417. DOI:10.1177/ 02783649211069572.
- Şucan IA, Moll M and Kavraki LE (2012) The Open Motion Planning Library. *IEEE Robotics & Automation Magazine* 19(4): 72–82. DOI:10.1109/MRA.2012.2205651. https: //ompl.kavrakilab.org.
- Thayer JT and Ruml W (2011) Bounded suboptimal search: A direct approach using inadmissible estimates. In: *International Joint Conference on Artificial Intelligence*. Barcelona, Spain, pp. 674–679.
- Tordesillas J and How JP (2022) Mader: Trajectory planner in multiagent and dynamic environments. *IEEE Transactions on Robotics* 38(1): 463–476. DOI:10.1109/TRO.2021.3080235.
- Webb DJ and van den Berg J (2013) Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In: *IEEE International Conference on Robotics*

and Automation. Karlsrühe, Germany, pp. 5054–5061. DOI: 10.1109/ICRA.2013.6631299.

- Yu J and LaValle SM (2013) Structure and intractability of optimal multi-robot path planning on graphs. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Bellevue, WA, pp. 1443–1449.
- Zhou X, Wen X, Wang Z, Gao Y, Li H, Wang Q, Yang T, Lu H, Cao Y, Xu C and Gao F (2022) Swarm of micro flying robots in the wild. *Science Robotics* 7(66): eabm5954. DOI:10.1126/ scirobotics.abm5954. URL https://www.science. org/doi/abs/10.1126/scirobotics.abm5954.