

Provably Constant-time Planning and Replanning for Real-time Grasping Objects off a Conveyor Belt

Fahad Islam¹, Oren Salzman², Aditya Agarwal¹, Maxim Likhachev¹

¹The Robotics Institute, Carnegie Mellon University

²Technion-Israel Institute of Technology

Abstract—In warehouse and manufacturing environments, manipulation platforms are frequently deployed at conveyor belts to perform pick and place tasks. Because objects on the conveyor belts are moving, robots have limited time to pick them up. This brings the requirement for fast and reliable motion planners that could provide provable real-time planning guarantees, which the existing algorithms do not provide. Besides the planning efficiency, the success of manipulation tasks relies heavily on the accuracy of the perception system which is often noisy, especially if the target objects are perceived from a distance. For fast moving conveyor belts, the robot cannot wait for a perfect estimate before it starts executing its motion. In order to be able to reach the object in time it must start moving early on (relying on the initial noisy estimates) and adjust its motion on-the-fly in response to the pose updates from perception. We propose an approach that meets these requirements by providing provable constant-time planning and replanning guarantees. We present it, give its analytical properties and show experimental analysis in simulation and on a real robot.

I. INTRODUCTION

Conveyor belts are widely used in automated distribution, warehousing, as well as for manufacturing and production facilities. In the modern times robotic manipulators are being deployed extensively at the conveyor belts for automation and faster operations [1]. In order to maintain a high-distribution throughput, manipulators must pick up moving objects without having to stop the conveyor for every grasp. In this work, we consider the problem of motion planning for grasping moving objects off a conveyor. An object in motion imposes a requirement that it should be picked up in a short window of time. The motion planner for the arm, therefore, must compute a path within a bounded time frame to be able to successfully perform this task.

Manipulation relies on high quality detection and localization of moving objects. When the object first enters the robot’s field of view, the initial perception estimates of the object’s pose are often inaccurate. Consider the example of an object (sugar box) moving along the conveyor towards the robot in Fig. 1, shown through an image sequence as captured by the robot’s Kinect camera in Fig. 2. The plot in Fig. 2 shows the variation of the error between the filtered input point cloud and a point cloud computed from the predicted pose from our ICP based perception strategy as the object gets closer to the camera. We observe that the error decreases as the object moves closer,



Fig. 1: A scene demonstrating the PR2 robot picking up a moving object (sugar box) off a conveyor belt.

indicating that the point clouds overlap more closely due to more accurate pose estimates closer to the camera.

However, if the robot waits too long to get an accurate estimate of the object pose, the delay in starting plan execution could cause the robot to miss the object. The likelihood of this occurring increases proportionately with the speed of the conveyor. Therefore, the robot should start executing a plan computed for the initial pose and as it gets better estimates, it should repeatedly replan for the new goals. However, for every replanning query, the time window for the pickup shrinks. This makes the planner’s job difficult to support real-time planning.

Furthermore, the planning problem is challenging because the motion planner has to account for the dynamic object and thus plan with time as one of the planning dimension. It should generate a valid trajectory that avoids collision with the environment around it and also with the target object to ensure that it does not damage or topple it during the grasp. Avoiding collisions with the object requires precise geometric collision checking between the object geometry and the geometry of the manipulator. The resulting complexity of the planning problem makes it infeasible to plan online for this task.

Motivated by these challenges, we propose an algorithm that leverages offline preprocessing to provide bounds on the planning time when the planner is invoked online. Our key insight is that in our domain the manipulation task is highly repetitive. Even for different object poses, the computed paths are quite similar and can be efficiently reused to speed up online planning. Based on this insight, we derive a method that precomputes a representative set of paths with some auxiliary




Fig. 2: (a) Depiction of an object moving along a conveyor towards the robot. (b) Pose error as a function of the distance from the conveyor’s start. Specifically we use ADD-S error [2].

datastructures offline and uses them online in a way that provides *constant-time planning guarantee*. Here, we assume that the geometric models of the target objects are known a priori. To the best of our knowledge, our approach is the first to provide constant-time planning guarantee on generating motions all the way to the goal for a dynamic environment.

We experimentally show that constant-time planning and re-planning capability is necessary for a successful conveyor pickup task. Specifically if we only perform one-time planning, (namely, either following the plan for the initial potentially inaccurate pose estimate or from a delayed but accurate pose estimate) the robot frequently fails to pick the object.

II. RELATED WORK

A. Motion planning for conveyor pickup task

Existing work on picking moving objects has focused on different aspects of the problem ranging from closed-loop controls to object perception and pose estimation, motion planning and others [3, 4, 5, 1]. Here, we focus on motion-planning related work. Time-configuration space representation was introduced to avoid moving obstacles [6, 7, 8]. Specifically in [8], a bidirectional sampling-based method with a time-configuration space representation was used to plan motions in dynamic environments to pickup moving objects. While their method showed real-time performance in complex tasks, it used fully-specified goals; namely knowing the time at which the object should be picked, which weakens the completeness guarantee. Furthermore their method is probabilistically complete and therefore, does not offer constant-time behavior. Graph-search based approaches have also been used for the motion-planning problem [9, 10]. The former uses a kinodynamic motion planner to smoothly pick up moving objects i.e., without an impactful contact. A heuristic search-based motion planner that plans with dynamics and could generate optimal trajectories with respect to the time of execution was used. While this planner provides strong optimality guarantees, it is not real-time and thus cannot be used online. The latter work demonstrated online real-time planning capability. The approach plans to a pregrasp pose with pure kinematic planning and relies on Cartesian-space controllers to perform the pick up. The usage of the Cartesian controller limits the types of objects that the robot can grasp.

B. Preprocessing-based planning

Preprocessing-based motion planners often prove beneficial for real-time planning. They analyze the configuration space offline to generate some auxiliary information that can be used online to speed up planning. Probably the best-known example is the Probabilistic Roadmap Method (PRM) [11] which precomputes a roadmap that can answer any query by connecting the start and goal configurations to the roadmap and then searching the roadmap. PRMs are fast to query yet they do not provide constant-time guarantees. Moreover, in our case, to account for a moving object, they would require edge re-evaluation which is often computationally expensive.

A provably constant-time planner was recently proposed in [12]. Given a start state and a goal region, it precomputes a compressed set of paths that can be utilized online to plan to any goal within the goal region in bounded time. As we will see, our approach while algorithmically different, draws some inspiration from this work. Both of the above two methods ([12, 11]) are mainly targetting pure kinematic planning and thus they cannot be used for the conveyor-planning problem which is dynamic in nature.

Another family of preprocessing-based planners utilizes previous experiences to speed up the search [13, 14, 15]. Experience graphs [15], provide speed up in planning times for repetitive tasks by trying to reuse previous experiences. These methods are also augmented with sparsification techniques (see e.g., [16, 17]) to reduce the memory footprint of the algorithm. Unfortunately, none of the mentioned algorithms provide fixed planning-time guarantees that we strive for in our application.

C. Online replanning and real time planning

The conveyor-planning problem can be modelled as a Moving Target Search problem (MTS) which is a widely-studied topic in the graph search-based planning literature [18, 19, 20, 21]. These approaches interleave planning and execution incrementally and update the heuristic values of the state space to improve the distance estimates to the moving target. Unfortunately, in high-dimensional planning problems, this process is computationally expensive which is why these approaches are typically used for two-dimensional grid problem such as those encountered in video games. More generally, real-time planning is widely considered in the search community (see, e.g., [22, 23, 24]). However, as mentioned, these works are typically applicable to low-dimensional search spaces.

III. PROBLEM DEFINITION

Our system is comprised of a robot manipulator \mathcal{R} , a conveyor belt \mathcal{B} moving at some known velocity, a set of known objects \mathcal{O} that need to be grasped and a perception system \mathcal{P} that is able to estimate the type of object and its location on \mathcal{B} .

Given a pose g of an object $o \in \mathcal{O}$, our task is to plan the motion of \mathcal{R} such that it grasps o from \mathcal{B} at some future time. Unfortunately, the perception system \mathcal{P} may give inaccurate object poses. Thus, the pose g will be updated by \mathcal{P} as \mathcal{R}

is executing its motion. To allow for \mathcal{R} to move towards the updated pose in real time, we introduce the additional requirement that planning should be done within a user-specified time bound T_{bound} . For ease of exposition, when we say that we plan to a pose g of o that is given by \mathcal{P} , we mean that we plan the motion of \mathcal{R} such that it will be able to pick o from \mathcal{B} at some future time. This is explained in detail in Sec. V and in Fig. 7.

We denote by G^{full} the discrete set of initial object poses on \mathcal{B} that \mathcal{P} can perceive. Finally, we assume that \mathcal{R} has an initial state s_{home} corresponding to the time $t = 0$ from which it starts planning to grasp any object.

Roughly speaking, the objective, following the set of assumptions we will shortly state, is to enable planning and replanning to any goal pose $g \in G^{\text{full}}$ in bounded time T_{bound} regardless of \mathcal{R} 's current state. To formalize this idea, let us introduce the notion of *reachable* and *covered* states:

Definition 1. A goal pose $g \in G^{\text{full}}$ is said to be *reachable* from a state s if there exists a path from s to g and it can be computed in finite time.

Definition 2. A reachable pose $g \in G^{\text{full}}$ is said to be *covered* by a state s if the planner can find a path from s to g within time T_{bound} .

Thus, we wish to build a system such that for any state s that the system can be in and every reachable goal pose $g \in G^{\text{full}}$ from s updated by \mathcal{P} , g is covered by s .

We are now ready to state the assumptions for which we can solve the problem defined.

- A1** There exists a replan cutoff time $t = t_{\text{rc}}$, after which the planner does not replan and continues to execute the last planned path.
- A2** For any time $0 \leq t \leq t_{\text{rc}}$, the environment is static. Namely, objects on \mathcal{B} cannot collide with \mathcal{R} during that time.
- A3** The pose estimation error of \mathcal{P} is bounded by a distance $\varepsilon_{\mathcal{P}}$.

Assumptions **A1-A2** enforce a requirement that \mathcal{P} must converge on an accurate estimate g before t_{rc} , and until t_{rc} , o is guaranteed to be at a safe distance from \mathcal{R} . Assumption **A3** bounds the maximum error of \mathcal{P} and is explained in detail in Sec. V and in Fig. 7

IV. ALGORITHMIC FRAMEWORK

Our approach for constant-time planning relies on a *preprocessing* stage that allows to efficiently compute paths in a *query* stage to any goal (under Assumptions **A1-A3**). Before we describe our approach, we start by describing a naïve method that solves the aforementioned problem but requires a prohibitive amount of memory. This can be seen as a warmup




Fig. 3: The figures show paths discretized from timesteps t_0 to t_{rc} with steps of size δ_t . (a) At $t = 0$, the algorithm computes n_{goal} paths, that is from s_{home} to every $g \in G^{\text{full}}$. (b) At $t = \delta_t$, the algorithm computes n_{goal}^2 paths, that is from all n_{goal} replanable states at $t = \delta_t$ to every $g \in G^{\text{full}}$ (here we only show paths from three states). Thus, the number of paths increases exponentially at every timestep.

before describing our algorithm which exhibits the same traits but does so in a memory-efficient manner.

A. Straw man approach

We first compute from s_{home} a path π_g to every reachable $g \in G^{\text{full}}$. These paths can be stored in a lookup (hash) table which can be queried in constant time (assuming perfect hashing [25]). Thus, all goals are covered by s_{home} and this allows us to start executing a path once \mathcal{P} gives its initial pose estimate. However, we need to account for pose update while executing π_g . Following **A1** and **A2**, this only needs to be done up until time t_{rc} . Thus, we discretize each path uniformly with resolution δ_t . We call all states that are less than t_{rc} time from s_{home} *replorable states*.

Next, for every replorable state along each path π_g , we compute a new path to all goals. This will ensure that all goals are covered by all replorable states. Namely, it will allow to immediately start executing a new path once the goal location is updated by \mathcal{P} . Unfortunately, \mathcal{P} may update the goal location more than once. Thus, this process needs to be performed recursively for the new paths as well.

The outcome of the preprocessing stage is a set of precomputed collision-free paths starting at states that are at most t_{rc} from s_{home} and end at goal states. The paths are stored in a lookup table $\mathcal{M} : S \times G^{\text{full}} \rightarrow \{\pi_1, \pi_2, \dots\}$ that can be queried in $O(1) (< T_{\text{bound}})$ time to find a path from any given $s \in S$ to $g \in G^{\text{full}}$.

In the query stage we obtain an estimation g_1 of the goal pose by \mathcal{P} . The algorithm then retrieves the path $\pi_1(s_{\text{home}}, g_1)$ (from s_{home} to g_1) from \mathcal{M} and the robot starts executing $\pi_1(s_{\text{home}}, g_1)$. For every new estimation g_i of the goal pose obtained from \mathcal{P} while the system is executing path $\pi_{i-1}(s_{i-1}, g_{i-1})$, the algorithm retrieves from \mathcal{M} the path $\pi_i(s_i, g_i)$ from the first state s_i along $\pi_{i-1}(s_{i-1}, g_{i-1})$ that is least T_{bound} away from s_{i-1} . The robot \mathcal{R} will then start executing $\pi_i(s_i, g_i)$ once it reaches s_i .

Clearly, every state is covered by this brute-force approach, however it requires a massive amount of memory. Let $n_{\text{goal}} = |G^{\text{full}}|$ be the number of goals and ℓ be the number of states between s_{home} and the state that is t_{rc} time away. This approach

requires precomputing and storing $O(n_{\text{goal}}^{\ell})$ paths which is clearly infeasible (see Fig. 3). In the next sections, we show how we can dramatically reduce the memory footprint of the approach without compromising on the system’s capabilities.

B. Algorithmic approach

While the straw man algorithm presented allows for planning to any goal pose $g \in G^{\text{full}}$ in bounded time T_{bound} , its memory footprint is prohibitively large. We suggest to reduce the memory footprint by building on the observation that many paths to close-by goals traverse very similar parts of the configurations space.

The key idea of our approach is that instead of computing (and storing) paths to all reachable goals in G^{full} , we compute a relatively small subset of so-called “root paths” that can be reused in such a way that we can still cover G^{full} fully. Namely, at query time, we can reuse these paths to plan to any $g \in G^{\text{full}}$ within T_{bound} . The idea is illustrated in Fig. 4.

First, we compute a set of root paths $\{\Pi_1, \dots, \Pi_k\}$ from s_{home} to cover G^{full} by s_{home} (here we will have that $k \ll n_{\text{goal}}$) Next, the algorithm recursively computes for all replanable states along these root paths, additional root paths so that their reachable goals are also covered. During this process, additional root paths are computed only when the already existing set of root paths does not provide enough guidance to the search to cover G^{full} i.e. to be able to compute a path to any $g \in G^{\text{full}}$ within T_{bound} . The remainder of this section formalizes these ideas.

C. Algorithmic building blocks

We start by introducing the algorithmic building blocks that we use. Specifically, we start by describing the motion planner that is used to compute the root paths and then continue to describe how they can be used as *experiences* to efficiently compute paths to other goals.

1) *Motion planner*: We use a heuristic search-based planning approach with motion primitives (see, e.g. [26, 27, 28]) as it allows for deterministic planning time which is key in our domain. Moreover, such planners can easily handle under-defined goals as we have in our setting—we define a goal as a grasp pose for the goal object while the planning dimension includes the DoFs of the robot as well as the time dimension.

State space and graph construction. We define a state s as a pair (q, t) where $q = (\theta_1, \dots, \theta_n)$ is a configuration represented by the joint angles for an n -DOF robot arm (in our setting $n = 7$) and t is the time associated with q . Given a state s we define two types of motion primitives which are short kinodynamically feasible motions that \mathcal{R} can execute.

The first type of motion primitives are predefined primitives. These are small individual joint movements in either direction as well as *wait* actions. For each motion primitive, we compute its duration by using a nominal constant velocity profile for the joint that is moved.

The second type of primitives are dynamic primitives. They are generated by the search only at the states that represent the arm configurations where the end effector is close to the object. These primitives correspond to the actual grasping of the object while it is moving. The dynamic primitives are generated by using a Jacobian pseudo inverse-based control law similar to what [9] used. The velocity of the end effector is computed such that the end-effector minimizes the distance to the grasp pose. Once the gripper encloses the object, it moves along with the object until the gripper is closed.

Motion planner. The states and the transitions implicitly define a graph $\mathcal{G} = (S, E)$ where S is the set of all states and E is the set of all transitions defined by the motion primitives. We use Weighted A* (wA*) [29] to find a path in \mathcal{G} from a given state s to a goal g . wA* is a suboptimal heuristic search algorithm that allows a tradeoff between optimality and greediness by inflating the heuristic function by a given weight w . The search is guided by an efficient and fast-to-compute heuristic function which in our case has two components. The first component drives the search to intercept the object at the right time and the second component guides the search to correct the orientation of the end effector as it approaches the object. Mathematically, our heuristic function is given by

$$h(s, g) = \max(\lambda \cdot t(s, g), \text{ANGLEDIFF}(s, g)).$$

Here, $t(s, g)$ is the expected time to intercept the object which can be analytically computed from the velocities and positions of the target object and the end-effector and $\text{ANGLEDIFF}(s, g)$ gives the magnitude of angular difference between the end-effector’s current pose and target pose. The coefficient λ is used as a weighting factor.

2) *Planning with Experience Reuse*: We now show how previously-computed paths which we named as root paths can be reused as experiences in our framework. Given a heuristic function h we define for a root path Π and a goal $g \in G^{\text{full}}$ the *shortcut* state $s_{\text{sc}}(\Pi, g)$ as the state that is closest to g with respect h . Namely,

$$s_{\text{sc}}(\Pi, g) := \arg \min_{s_i \in \Pi} h(s_i, g).$$

Now, when searching for a path to a goal $g \in G^{\text{full}}$ using root path Π as an experience, we add $s_{\text{sc}}(\Pi, g)$ as a successor for any state along Π (subject to the constraint that the path along Π to s_{sc} is collision free). In this manner we reuse previous experience to quickly reach a state close to the g .

D. Algorithmic details

We are finally ready to describe our algorithm describing first the preprocessing phase and then the query phase.

1) *Preprocessing*: Our preprocessing stage starts by sampling a goal $g_1 \in G^{\text{full}}$ and computing a root path Π_1 from s_{home} to g_1 . We then associate with Π_1 the set of goals $G_1 \subset G^{\text{full}}$ such that Π_1 can be used as an experience in reaching any




Fig. 4: First step of the preprocessing stage. (a) A goal g_1 is sampled and the root path Π_1 is computed between s_{home} and g_1 . (b) The set $G_1 \subset G^{\text{full}}$ of all states that can use Π_1 as an experience is computed and associated with Π_1 . (c) The goal region covered by four root paths from s_{home} after the first step of the preprocessing stage terminates.

Algorithm 1 Plan Root Paths

```

1: procedure PLANROOTPATHS( $s_{\text{start}}, G^{\text{uncov}}$ )
2:    $\Psi_{s_{\text{start}}} \leftarrow \emptyset$   $\triangleright$  a list of pairs  $(\Pi_i, G_i)$ 
3:    $G_{s_{\text{start}}}^{\text{uncov}} \leftarrow \emptyset$ ;  $i = 0$ 
4:   while  $G^{\text{uncov}} \neq \emptyset$  do  $\triangleright$  until all reachable goals are covered
5:      $g_i \leftarrow \text{SAMPLEGOAL}(G^{\text{uncov}})$ 
6:      $G^{\text{uncov}} \leftarrow G^{\text{uncov}} \setminus \{g_i\}$ 
7:     if  $\Pi_i \leftarrow \text{PLANROOTPATH}(s_{\text{start}}, g_i)$  then  $\triangleright$  planner succeeded
8:        $G_i \leftarrow \{g_i\}$   $\triangleright$  goals reachable
9:       for each  $g_j \in G^{\text{uncov}}$  do
10:        if  $\pi_j \leftarrow \text{PLANPATHWITHEXPERIENCE}(s_{\text{start}}, g_j, \Pi_i)$  then
11:           $G_i \leftarrow G_i \cup \{g_j\}$ 
12:           $G^{\text{uncov}} \leftarrow G^{\text{uncov}} \setminus \{g_j\}$ 
13:        $\Psi_{s_{\text{start}}} \leftarrow \Psi_{s_{\text{start}}} \cup \{(\Pi_i, G_i)\}$ ;  $i \leftarrow i + 1$ 
14:     else
15:        $G_{s_{\text{start}}}^{\text{uncov}} \leftarrow G_{s_{\text{start}}}^{\text{uncov}} \cup \{g_i\}$   $\triangleright$  goals unreachable
16:   return  $\Psi_{s_{\text{start}}}, G_{s_{\text{start}}}^{\text{uncov}}$ 

```

$g_j \in G_1$ within T_{bound}^1 . Thus, all goals in G_1 are covered by s_{home} . We then repeat this process but instead of sampling a goal from G^{full} , we sample from $G^{\text{full}} \setminus G_1$, thereby removing covered goals from G^{full} in every iteration. At the end of this step, we obtain a set of root paths. Each root path Π_i is associated with a goal set $G_i \subseteq G^{\text{full}}$ such that (i) Π_i can be used as an experience for planning to any $g_j \in G_i$ in T_{bound} and (ii) $\bigcup_i G_i = \text{REACHABLE}(s_{\text{home}}, G^{\text{full}})$ (i.e all reachable goals for s_{home} in G^{full}). Alg. 1 details this step (when called with arguments $(s_{\text{home}}, G^{\text{full}})$). It also returns a set of unreachable goals that are left uncovered. The process is illustrated in Fig. 4.

So far we explained the algorithm for one-time planning when the robot is at s_{home} ($t = 0$); we now need to allow for efficient replanning for any state s between $t = 0$ to t_{rc} . In order to do so, we iterate through all the states on these root paths and add additional root paths so that these states also cover their respective reachable goals. This has to be done recursively since newly-added paths generate new states which the robot may have to replan from. The complete process is detailed in Alg. 2. The PREPROCESS procedure takes in a state s_{start} , the goal region that it has to cover G^{uncov} and region that it already has covered G^{cov} . Initially PREPROCESS is called with arguments $(s_{\text{home}}, G^{\text{full}}, \emptyset)$ and it runs recursively until no state

¹In practice, to account for other query phase operations, such as hash table lookups etc., a slightly smaller time than T_{bound} is provided to the experience-based planner, to ensure that the overall query time is bounded by T_{bound} .

Algorithm 2 Preprocess

```

1: procedure TRYLATCHING( $s, \Psi_{s_{\text{home}}}, G^{\text{uncov}}, G^{\text{cov}}$ )
2:   for each  $(\Pi_i, G_i) \in \Psi_{s_{\text{home}}}$  do
3:     if  $\text{CANLATCH}(s, \Pi_i)$  then
4:        $G^{\text{uncov}} \leftarrow G^{\text{uncov}} \setminus G_i$ 
5:        $G^{\text{cov}} \leftarrow G^{\text{cov}} \cup G_i$ 
6:   return  $G^{\text{uncov}}, G^{\text{cov}}$ 
7: procedure PREPROCESS( $s_{\text{start}}, G^{\text{uncov}}, G^{\text{cov}}$ )
8:    $\Psi_{s_{\text{start}}}, G_{s_{\text{start}}}^{\text{uncov}} \leftarrow \text{PLANROOTPATHS}(s_{\text{start}}, G^{\text{uncov}})$ 
9:   if  $s_{\text{start}} = s_{\text{home}}$  then  $\Psi_{s_{\text{home}}} = \Psi_{s_{\text{start}}}$ 
10:   $G_{s_{\text{start}}}^{\text{cov}} \leftarrow G^{\text{cov}} \cup (G^{\text{uncov}} \setminus G_{s_{\text{start}}}^{\text{uncov}})$ 
11:  if  $t(s_{\text{start}}) \leq t_{\text{rc}}$  then
12:    for each  $(\Pi_i, G_i) \in \Psi_{s_{\text{start}}}$  do
13:       $G_i^{\text{cov}} \leftarrow G_i$ ;  $G_i^{\text{uncov}} \leftarrow G_{s_{\text{start}}}^{\text{cov}} \setminus G_i$ 
14:      for each  $s \in \Pi_i$  (from last to first) do  $\triangleright$  states up to  $t_{\text{rc}}$ 
15:         $G_i^{\text{uncov}}, G_i^{\text{cov}} \leftarrow \text{TRYLATCHING}(s, \Psi_{s_{\text{home}}}, G_i^{\text{uncov}}, G_i^{\text{cov}})$ 
16:        if  $G_i^{\text{uncov}} = \emptyset$  then
17:          break
18:         $G_i^{\text{uncov}}, G_i^{\text{cov}} \leftarrow \text{PREPROCESS}(s, G_i^{\text{uncov}}, G_i^{\text{cov}})$ 
19:        if  $G_i^{\text{uncov}} = \emptyset$  then
20:          break
21:  return  $G_{s_{\text{start}}}^{\text{uncov}}, G_{s_{\text{start}}}^{\text{cov}}$ 

```

is left with uncovered reachable goals.

At a high level, the algorithm iterates through each root path Π_i (loop at line 12) and for each state $s \in \Pi_i$ (loop at line 14) the algorithm calls itself recursively (line 18). The algorithm terminates when all states cover their reachable goals. The pseudocode in blue constitute an additional optimization step which we call ‘‘latching’’ and is explained later in Sec. IV-D3.

In order to minimize the required computation, the algorithm leverages two key observations:

- O1** If a goal is not reachable from a state $s \in \Pi$, it is not reachable from all the states after it on Π .
- O2** If a goal is covered by a state $s \in \Pi$, it is also covered by all states preceding it on Π .

We use **O1** to initialize the uncovered set for any state; instead of attempting to cover the entire G^{full} for each replanable state s , the algorithm only attempts to cover the goals that could be reachable from s , thereby saving computation. **O2** is used by iterating backwards on each root path (loop at line 14) and for each state on the root path only considering the goals that are left uncovered by the states that appear on the path after it.

Algorithm 3 Query

Inputs: $\mathcal{M}, s_{\text{home}}$

- 1: **procedure** PLANPATHBYLATCHING(s_{start}, g)
- 2: **if** $\Pi_{\text{home}} \leftarrow \mathcal{M}(s_{\text{home}}, g)$ **exists** **then** \triangleright lookup root path
- 3: **if** CANLATCH(s, Π_{home}) **then**
- 4: $\pi_{\text{home}} \leftarrow$ PLANPATHWITHEXPERIENCE($s_{\text{start}}, g, \Pi_{\text{home}}$)
- 5: $\pi \leftarrow$ MERGEPATHSBYLATCHING($\pi_{\text{curr}}, \pi_{\text{home}}, s$)
- 6: **return** π
- 7: **return failure**

- 8: **procedure** QUERY($g, \pi_{\text{curr}}, s_{\text{start}}$)
- 9: **for each** $s \in \pi_{\text{curr}}$ (from last to s_{start}) **do** \triangleright states up to t_{rc}
- 10: **if** $\Pi_{\text{next}} \leftarrow \mathcal{M}(s, g)$ **exists** **then** \triangleright lookup root path
- 11: $\pi_{\text{next}} \leftarrow$ PLANPATHWITHEXPERIENCE($s_{\text{start}}, g, \Pi_{\text{next}}$)
- 12: $\pi \leftarrow$ MERGEPATHS($\pi_{\text{curr}}, \pi_{\text{next}}, s$)
- 13: **return** π
- 14: **if** $\pi \leftarrow$ PLANPATHBYLATCHING(s_{start}, g) **successful** **then**
- 15: **return** π
- 16: **return failure** \triangleright goal is not reachable

Specifically, **O2** is used to have a single set of uncovered goals G_i^{uncov} for all states that appear on Π_i instead of having individual sets for each state and the goals that each $s \in \Pi_i$ covers in every iteration of loop **14** are removed from G_i^{uncov} . **O1** is used to initialize G_i^{uncov} (in line **13**). Namely, it is initialized not by the entire G_i^{full} but by the set of goals covered by s_{start} . G_i is excluded since it is already covered via Π_i . The iteration completes either when all goals in G_i^{uncov} are covered (line **19**) or the loop backtracks to s_{start} . The process is illustrated in Fig. 5

Thus, as the outcome of the preprocessing stage a map $\mathcal{M} : S \times G^{\text{full}} \rightarrow \{\Pi_1, \Pi_2, \dots\}$ is constructed that can be looked up to find which root path can be used as an experience to plan to a goal g from a state s within T_{bound} .

2) *Query*: Alg. 3 describes the query phase of our algorithm. Again, the lines in blue correspond to the blue pseudocode in Alg. 2 for the additional optimization step which is explained in Sec. IV-D3. Assume that the robot was at a state s_{curr} while executing a path π_{curr} when it receives a pose update g from the perception system. Alg. 3 will be called for a state s_{start} that is T_{bound} ahead of s_{curr} along π_{curr} , allowing the algorithm to return a plan before the robot reaches s_{start} .

Alg. 2 assures that there exists one state on π_{curr} between s_{start} and the state at t_{rc} that covers g . Therefore, we iterate over each $s \in \pi_{\text{curr}}$ backwards (similar to Alg. 2) between s_{start} and the state at t_{rc} and find the one that covers g by querying \mathcal{M} . Once found, we use the corresponding root path Π_{next} as an experience to plan the path π_{next} from s to g . Finally the paths π_{curr} and π_{next} are merged together with s being the transitioning state to return the final path π .

3) *Latching: Reusing Root Paths*: We introduce an additional step called ‘‘Latching’’ to minimize the number of root paths computed in Alg. 2. With latching, the algorithm tries to reuse previously-computed root paths as much as possible using special motion primitives that allow transitions from one root path to another. The primitive is computed from a state $s \in \Pi_i$

to $s' \in \Pi_j$ such that $t(s') = t(s) + \delta_t$ by simple linear interpolation while ensuring that kinodynamic constraints of the robot are satisfied. Specifically, given the nominal joint velocities of the robot, if s' can be reached from s in time δ_t while respecting the kinematic and collision constraints, then the transition is allowed.

In Alg. 2, before calling the PREPROCESS procedure for a state, the algorithm removes the set of goals that can be covered via latching, thereby reducing the number of goals that need to be covered by the PREPROCESS procedure. Correspondingly, in Alg. 3, an additional procedure is called to check if the path can be found via latching. These additions in the two pseudocodes are shown in blue. An iteration of the complete algorithm with latching is illustrated in Fig 6.

E. Theoretical guarantees

Lemma 1 (Completeness). *For a robot state s and a goal g , if g is reachable from s and $t(s) \leq t_{rc}$, the algorithm is guaranteed to find a path from s to g .*

Proof (Sketch): In order to prove it we show that (1) if g is reachable from s , it is covered by s in Alg. 2 and (2) if g is covered by s , Alg. 3 is guaranteed to return a path.

Alg. 2 starts by computing a set of root paths from s_{home} that ensures that it covers all of its reachable goals. It then iterates over all states on these paths and adds additional root paths ensuring that these states also cover their reachable goals. It does it recursively until no state before t_{rc} is left with uncovered goals. Therefore, it is ensured that any reachable g is covered by s , provided that $t(s) \leq t_{rc}$.

Alg. 2 covers g via at least one state between s and the state at t_{rc} (inclusively) (loop at line **14**). In query phase, Alg. 3 iterates through all states between s and the state at t_{rc} (inclusively) to identify the one that covers g (loop at line **9**). Since g is covered by at least one of these states by Alg. 2, Alg. 3 is guaranteed to find a path from s to g . ■

Lemma 2 (Constant-time complexity). *Let s be a replanable state and g a goal provided by \mathcal{P} . If g is reachable, the planner is guaranteed to provide a solution in constant time.*

Proof: We have to show that the query stage (Alg. 3) has a constant-time complexity. The number of times the algorithm queries \mathcal{M} which is $O(1)$ operation in case of perfect hashing is bounded by $l = t_{rc}/\delta_t$ which is the maximum number of time steps from $t = 0$ to t_{rc} . The number of times the algorithm will attempt to latch on to a root path (namely, a call to CANLATCH which is a constant-time operation) is also bounded by l . Finally, Alg. 3 calls the PLAN method only once. Since the state that it is called for covers g , meaning that the planner can find a path from it to g within T_{bound} , the computation time is constant. Hence the overall complexity of Alg. 3 is $O(1)$. ■




Fig. 5: Preprocess loop for Π_1 without latching. (a) Initially the state s covers G_1 via Π_1 . (b) New root paths are computed from s to cover remaining uncovered region. (c) This process is repeated by backtracking along the root path. (d) Outcome of a preprocessing step for one path: G^{full} is covered either by using Π_1 as an experience or by using newly-computed root paths.




Fig. 6: Preprocess loop for Π_1 with latching. (a) The algorithm starts by trying to latch on to every other root path; for successful latches, the corresponding goals are removed from uncovered region. (b) New root paths are computed from s to cover remaining uncovered region. (c) This process is repeated by backtracking along the root path. (d) Outcome of a preprocessing step: G^{full} is covered either by using Π_1 as an experience, latching on to Π_2 , Π_3 or Π_4 (at different time steps) or by using newly-computed root paths.

V. EVALUATION

We evaluated our algorithm in simulation and on a real robot. The conveyor speed that we used for all of our results is 0.2m/s . The experiments video can be found at <https://youtu.be/iLVPBWxa5b8>. We used Willow Garage’s PR2 robot in our experiments using its 7-DOF arm. The additional time dimension makes the planning problem eight dimensional.

A. Experimental setup

1) *Sense-plan-act cycle*: As object o moves along \mathcal{B} , we use the Brute Force ICP pose estimation baseline proposed in [30] to obtain its 3-Dof pose for each captured input point cloud. The plan is computed for the sensed object pose projected forward by T_{bound} time, giving the planner T_{bound} time to plan. If the plan comes in earlier, the robot waits until the object reaches the projected pose, before executing the plan to ensure that timing is tracked properly.

2) *Goal region specification*: To define the set of all goal poses G^{full} , we need to detail our system setup, depicted in Fig. 7. The conveyor belt \mathcal{B} moves along the x -axis from left to right. We pick a fixed x -value termed x_{exec} , such that when the incoming o reaches x_{exec} as per the perception information, at that point we start execution.

Recall that a pose of an object o is a three dimensional point (x, y, θ) corresponding to the (x, y) location of o and to its orientation (yaw angle) along \mathcal{B} . G^{full} contains a fine discretization of all possible y and θ values and x values in $[x_{\text{exec}} - 2\varepsilon_p, x_{\text{exec}} + 2\varepsilon_p]$. We select G^{full} such that $\varepsilon_p = 2.5\text{cm}$, making the goal region 10cm long along x -axis. Its dimension along y -axis is 20cm , equal to the width of the \mathcal{B} . The discretization in x, y and θ is 1.0cm and 10 degrees respectively.

In the example depicted in Fig. 7, the thick and the thin




Fig. 7: A depiction of G^{full} -specification on a conveyor belt (overhead view) and perception noise handling.

solid rectangles show the ground truth and estimated poses, respectively at two time instances in the life time of the object. The first plan is generated for the pose shown at x_{exec} . During execution, the robot receives an improved estimate and has to replan for it. At this point we back project this new estimate in time using the known speed of the conveyor and the time duration between the two estimates. This back-projected pose (shown as the dotted rectangle) is then picked as the new goal for replanning. Recall that under the assumption **A3** the back projected pose will always lie inside G^{full} .

B. Results

The preprocessing stage (i.e. running Alg. 2) took roughly 3.5 hours and the memory footprint following this stage was less than 20Mb. This supports our intuition that the domain allows for efficient compression of a massive amount of paths in a reasonable amount of preprocessing time. In all experiments, we used $t_{\text{rc}} = 3.5\text{s}$ and $\delta_t = 0.5\text{s}$.

1) *Real-robot experiments*: To show the necessity of real-time replanning in response to perception updates, we performed three types of experiments, (E1) using our approach to replan every time new object pose estimate arrives, (E2) single-shot planning based on the first object pose estimate (E3) single-shot planning using the late (more accurate) pose estimate. For

	Our Method	wA*			E-Graph			RRT		
	$T_b = 0.2$	$T_b = 0.5$	$T_b = 1.0$	$T_b = 2.0$	$T_b = 0.5$	$T_b = 1.0$	$T_b = 2.0$	$T_b = 0.5$	$T_b = 1.0$	$T_b = 2.0$
Pickup success [%]	92.0	0.0	0.0	18.0	0.0	0.0	80.0	0.0	0.0	18.0
Planning success [%]	94.7	4.0	17.0	19.0	31.0	80.0	90.0	12.0	9.0	13.0
Planning time [s]	0.069	0.433	0.628	0.824	0.283	0.419	0.311	0.279	0.252	0.197
Planning cycles	3	2	2	2	2	2	2	2	2	2
Path cost [s]	10.11	8.19	8.28	7.60	8.54	8.22	7.90	9.68	8.96	8.04

TABLE I: Simulation results. Here T_b denotes the (possibly arbitrary) timebound that the algorithm uses. Note that for our method $T_b = T_{\text{bound}}$ is the time bound that the algorithm is ensured to compute a plan.

	Success rate	Accuracy of \mathcal{P} [%]	Success rate (Accurate \mathcal{P})	Success rate (Inaccurate \mathcal{P})
E1	69.23	42.31	83.33	57.14
E2	16.00	24.00	66.67	0.00
E3	34.61	34.62	55.56	23.53

TABLE II: Real-robot experiments. Success rate for the three experiments (E1—our method, E2—First-pose planning and E3—Best-pose planning).

each set of experiments, we determined the pickup success rate to grasp the moving object (sugar box) off \mathcal{B} . In addition, we report on \mathcal{P} 's success rate by observing the overlap between the point cloud of the object's 3D model transformed by the predicted pose (that was used for planning) and the filtered input point cloud containing points belonging to the object. A high (low) overlap corresponds in an accurate (inaccurate) pose estimate. We use the same strategy to determine the range for which \mathcal{P} 's estimates are accurate and use it to determine the time for the best-pose planning. Further, for each method, we determine the pickup success rate given that the \mathcal{P} 's estimate was or wasn't accurate.

The experimental results are shown in Table II. Our method achieves the highest overall pickup success rate on the robot by a large margin, indicating the importance of continuous replanning with multiple pose estimates. First-pose planning has the least overall success rate due to inaccuracy of pose estimates when the object is far from the robot's camera. Best-pose planning performs better overall than the first pose strategy, since it uses accurate pose estimates, received when the object is close to the robot. However it often fails even when perception is accurate, since a large number of goals are unreachable due to limited time remaining to grasp the object when it is closer to the robot.

2) *Simulation experiments*: We simulated the real world scenario to evaluate our method against other baselines. We compared our method with wA* [29], E-graph [15] and RRT [31]. For wA* and E-graph we use the same graph representation as our method. For E-graph we precompute five paths to randomly-selected goals in G^{full} . We adapt the RRT algorithm to account for the under-defined goals. To do so, we sample pre-grasp poses along the conveyor and compute IK solutions for them to get a set of goal configurations for goal biasing. When a newly-added node falls within a threshold distance from the object, we use the same dynamic primitive that we use in the search-based methods to add the final grasping maneuver. If the primitive succeeds, we return

success. We also allow wait actions at the pre-grasp locations.

For any planner to be used in our system, we need to endow it with a (possibly arbitrary) planning time bound to compute the future location of the object from which the new execution will start. If the planner fails to generate the plan within this time, then the robot misses the object for that cycle and such cases are recorded as failures. We label a run as a success if the planner successfully replans once after the object crosses the 1.0m mark. The mark is the mean of accurate perception range that was determined experimentally and used in the robot experiments as described in Section V-B1. The key takeaway from our experiments (Table I) is that having a known time bound on the query time is vital to the success of the conveyor pickup task.

Our method shows the highest pickup success rate, planning success rate (success rate over all planning queries) and an order of magnitude lower planning times compared to the other methods. The planning success rate being lower than 100% can be attributed to the fact that some goals were unreachable during the runs. We tested the other methods with several different time bounds. After our approach E-graph performed decently well. RRT suffers from the fact that the goal is under-defined and sampling based planners typically require a goal bias in the configuration space. Another important highlight of the experiments is the number of planning cycles over the lifetime of an object. While the other approaches could replan at most twice, our method was able to replan thrice due to fast planning times.

VI. CONCLUSION

To summarize, we developed a provably constant-time planning and replanning algorithm that can be used to grasp fast moving objects off conveyor belts and evaluated it in simulation and in the real world on the PR2 robot. Through this work, we advocate the need for algorithms that guarantee (small) constant-time planning for time critical applications, such as the conveyor pickup task, which are often encountered in warehouse and manufacturing environments.

REFERENCES

- [1] Yizhe Zhang, Lianjun Li, Michael Ripperger, Jorge Nicho, Malathi Veeraraghavan, and Andrea Fumagalli. [Gilbreth: A conveyor-belt based pick-and-sort industrial robotics application](#). In *IRC*, pages 17–24, 2018.

- [2] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. [Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes.](#) In *ACCV*, pages 548–562, 2012.
- [3] Peter K Allen, Aleksandar Timcenko, Billibon Yoshimi, and Paul Michelman. [Automated tracking and grasping of a moving object with a robotic hand-eye system.](#) *TRA*, 9(2):152–165, 1993.
- [4] Shuai D Han, Si Wei Feng, and Jingjin Yu. [Toward Fast and Optimal Robotic Pick-and-Place on a Moving Conveyor.](#) *RAL*, 2019.
- [5] Denis Stogl, Daniel Zumkeller, Stefan Escalda Navarro, Alexander Heilig, and Björn Hein. [Tracking, reconstruction and grasping of unknown rotationally symmetrical objects from a conveyor belt.](#) In *EFTA*, pages 1–8. IEEE, 2017.
- [6] Thierry Fraichard. [Dynamic trajectory planning with dynamic constraints: A ‘state-time space’ approach.](#) In *Proceedings of 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’93)*, volume 2, pages 1393–1400. IEEE, 1993.
- [7] Massimo Cefalo, Giuseppe Oriolo, and Marilena Vendittelli. [Task-constrained motion planning with moving obstacles.](#) In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5758–5763. IEEE, 2013.
- [8] Yiming Yang, Wolfgang Merkt, Vladimir Ivan, and Sethu Vijayakumar. [Planning in time-configuration space for efficient pick-and-place in non-static environments with temporal constraints.](#) In *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, pages 1–9. IEEE, 2018.
- [9] Arjun Menon, Benjamin Cohen, and Maxim Likhachev. [Motion planning for smooth pickup of moving objects.](#) In *ICRA*, pages 453–460, 2014.
- [10] Anthony Cowley, Benjamin Cohen, William Marshall, Camillo J Taylor, and Maxim Likhachev. [Perception and motion planning for pick-and-place of dynamic objects.](#) In *IROS*, pages 816–823, 2013.
- [11] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. [Probabilistic roadmaps for path planning in high-dimensional configuration spaces.](#) *RAL*, 12(4):566–580, 1996.
- [12] Fahad Islam, Oren Salzman, and Maxim Likhachev. [Provable Indefinite-Horizon Real-Time Planning for Repetitive Tasks.](#) In *ICAPS*, pages 716–724, 2019.
- [13] Dmitry Berenson, Pieter Abbeel, and Ken Goldberg. [A robot path planning framework that learns from experience.](#) In *ICRA*, pages 3671–3678, 2012.
- [14] David Coleman, Ioan A Şucan, Mark Moll, Kei Okada, and Nikolaus Correll. [Experience-based planning with sparse roadmap spanners.](#) In *ICRA*, pages 900–905, 2015.
- [15] Michael Phillips, Benjamin J. Cohen, Sachin Chitta, and Maxim Likhachev. [E-Graphs: Bootstrapping Planning with Experience Graphs.](#) In *RSS*, 2012.
- [16] Andrew Dobson and Kostas E. Bekris. [Sparse roadmap spanners for asymptotically near-optimal motion planning.](#) *IJRR*, 33(1):18–47, 2014.
- [17] Oren Salzman, Doron Shaharabani, Pankaj K. Agarwal, and Dan Halperin. [Sparsification of motion-planning roadmaps by edge contraction.](#) *IJRR*, 33(14):1711–1725, 2014.
- [18] Toru Ishida and Richard E Korf. [Moving Target Search.](#) In *IJCAI*, volume 91, pages 204–210, 1991.
- [19] Toru Ishida and Richard E. Korf. [Moving-target search: A real-time search for changing goals.](#) *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(6):609–619, 1995.
- [20] Sven Koenig, Maxim Likhachev, and Xiaoxun Sun. [Speeding up moving-target search.](#) In *AAMAS*, pages 1–8, 2007.
- [21] Xiaoxun Sun, William Yeoh, and Sven Koenig. [Moving target D* lite.](#) In *AAMAS*, pages 67–74, 2010.
- [22] Sven Koenig and Maxim Likhachev. [Real-time adaptive A*.](#) In *AAMAS*, pages 281–288, 2006.
- [23] Sven Koenig and Xiaoxun Sun. [Comparing real-time and incremental heuristic search for real-time situated agents.](#) *JAAMAS*, 18(3):313–341, 2009.
- [24] Richard E Korf. [Real-time heuristic search.](#) *Artificial intelligence*, 42(2-3):189–211, 1990.
- [25] Zbigniew J Czech, George Havas, and Bohdan S Majewski. [Perfect hashing.](#) *Theoretical Computer Science*, 182(1-2):1–143, 1997.
- [26] Benjamin J. Cohen, Sachin Chitta, and Maxim Likhachev. [Search-based planning for manipulation with motion primitives.](#) In *ICRA*, pages 2902–2908, 2010.
- [27] Benjamin J. Cohen, Gokul Subramania, Sachin Chitta, and Maxim Likhachev. [Planning for Manipulation with Adaptive Motion Primitives.](#) In *ICRA*, pages 5478–5485, 2011.
- [28] Maxim Likhachev and Dave Ferguson. [Planning Long Dynamically Feasible Maneuvers for Autonomous Vehicles.](#) *IJRR*, 28(8):933–945, 2009.
- [29] Ira Pohl. [.](#) *Artificial intelligence*, 1(3-4):193–204, 1970.
- [30] Venkatraman Narayanan and Maxim Likhachev. [PERCH: Perception via search for multi-object recognition and localization.](#) In *ICRA*, pages 5052–5059, 2016.
- [31] Steven M LaValle. [Rapidly-exploring random trees: A new tool for path planning.](#) 1998.