Using Classical Planners for Tasks with Continuous Operators in Robotics

Siddharth Srivastava and Lorenzo Riano and Stuart Russell and Pieter Abbeel

Computer Science Division University of California, Berkeley Berkeley, CA 94720

Abstract

The need for high-level task planning in robotics is well understood. However, interfacing discrete planning with continuous actions often requires extensive engineering of the solution. For instance, picking up an object may require removing many others that obstruct it. Identifying the exact obstructions requires geometric reasoning which is prohibitively expensive to precompute, with results that are difficult to represent efficiently at the level of a discrete planner. We propose a new approach that utilizes representation techniques from first-order logic and provides a method for synchronizing between continuous and discrete planning layers. We evaluate the approach and illustrate its robustness through a number of experiments using a state-of-the-art robotics simulator, accomplishing a variety of challenging tasks like picking objects from cluttered environments, where the planner needs to figure out which other objects need to be moved first to be able to reach the target object, and laying out a table for dinner, where the planner figures out effective tray-loading, navigation and unloading strategies.

1 Introduction

A robot capable of achieving high-level goals was one of the original motivations behind the automated planning problem (Fikes and Nilsson 1971). Since this early work, numerous advances have been made in automated planning. Discrete planners are able to automatically compute effective problem-specific heuristics for solving planning problems specified implicitly in terms of parameterized operators with preconditions and effects (e.g., (Hoffmann and Nebel 2001; Helmert, Haslum, and Hoffmann 2007)). Continuous state space planners have also been developed for solving the lower level search problems in the configuration space of a robot to achieve desired motion trajectories (LaValle 2006). Techniques from both of these areas are required for a realworld robot to solve a high-level problem like preparing a table for dinner. However, using classical planners to solve planning problems encountered by a robot presents several fundamental challenges.

For instance, consider the task of picking up an object in blocks-world, a domain that is widely regarded as trivial for modern classical planners. If the table is cluttered, a continuous planner will fail to find a solution because other objects will have to be removed first (see Fig. 1). A classical planner could compute the sequence of operations required to clear



Figure 1: Example scenario with the target object in blue (L) and the desired solution after removing obstructing objects (R)

a path to the object, but only if it gets the set of obstructions. Superficially, this seems to be easily resolved by adding a high-level predicate, $obstructs(c, b_1, b_2)$ where c ranges over the robot configurations and b_i over objects. However, any such formulation leads to two major problems. First, the robot configuration is usually a high dimensional real-valued vector (11-dimensional for the PR2 using only one gripper). Representing this problem for a discrete planner requires a discretization of the possible configurations which in itself is infeasible for the high dimensional spaces involved. Moreover, a pre-processing step will need to compute the graspability and obstruction facts specific to the given scenario, for every combination of discretized configurations and objects used as arguments to the corresponding predicates. This will lead to planner input files with thousands of configuration symbols at the least, making planning infeasible. Even if all of these representational steps are carried out, it is difficult to specify in a discrete problem definition how obstructions change on the application of a pick and place operation, due to the spatial reasoning involved. Finally, one of the problems with fixing a discretization a priori is that it may miss the points required for a solution.

These issues go beyond pick and place tasks and stem from the following fundamental problems: how to enable classical planners to efficiently utilize information provided by a continuous planner, and how to utilize them in situations where facts and operator effects over continuous variables are not available a priori. Our main contribution is a new approach addressing these problems. We show that this approach can solve number of challenging robot planning tasks.

Overview of Our Approach We assume that every high-

level action corresponds to a continuous implementation. The continuous implementation may use a variety of techniques for accomplishing the post-conditions of the highlevel action. If the preconditions of the high-level action are satisfied, a trajectory is guaranteed to be found by the continuous implementation. However, as discussed in the introduction, the preconditions of high-level actions may involve spatial reasoning and are generally not computable a priori. Thus, the high-level planner may be wrong in selecting an action for execution at a certain step. If a solution trajectory cannot be found, the lower level returns the violated preconditions to the high-level planning layer, which incorporates the new facts and computes a new high-level solution plan. Throughout this paper, we will use the terms "discrete planning" interchangeably with high-level planning and "lower level" with continuous planning and execution.

While this overall approach is similar to re-planning (Talamadupula et al. 2010; Yoon, Fern, and Givan 2007), our focus is on the problem of continuous operator representations for high-level, discrete planning and communication between the continuous and discrete planning levels. The discrete problem is specified by replacing the domains of real-valued variables (required to express action preconditions and effects, even at the high-level) with finite sets of Skolem symbols. This translation is described in Sec. 2.1 The lower layer maintains an interpretation of these symbols (e.g., specific robot configurations). Failure messages from the lower level specify the failed preconditions of actions in the form of literals over such symbols. This interface as well as the details of planning methods used is described in Sec. 3 If the high-level planning layer finds the problem to be unsolvable, the lower level computes a new interpretation and the planning process resumes. This way the highlevel planning process never has to deal with the continuous variables and the lower level does not have to deal with discrete task sequencing. The resulting approach allows easy modeling of various tasks. We present many experimental results in a simulated execution scenario developed using OpenRave (Diankov 2010) for the PR2 robot in Sec.4 A summary of related work is presented in Sec. 5

2 Problem Formulation

We assume that a propositional framework like typed PDDL (Fox and Long 2003) is used for specifying highlevel planning domains and problems. We formalize the representation briefly as follows. A *planning domain* $\langle \mathcal{R}, \mathcal{A} \rangle$ consists of a set of relation symbols \mathcal{R} with their arities and type signatures, and a set \mathcal{A} of actions. Let the set of literals and atoms constructed using a set of relations \mathcal{R} and a set of typed variables or constants V be $lit(\mathcal{R}, V)$ and $atoms(\mathcal{R}, V)$ respectively. An action is defined as $\langle Param, Pre, Eff \rangle$, where *Param* is a sequence of typed parameter variables; *Pre* is a first-order logic formula and *Eff* is a conjuction, both over $lit(\mathcal{R}, V)$ with only the variables in *Param* being free. With some syntactic limitations, such expressions can be represented in the PDDL input language used by most classical planners.

Given a finite set U of typed constants (informally constituting the "objects" in a problem), a state is an assignment of truth values to each atom in $atoms(\mathcal{R}, U)$; for compactness, we use the closed world assumption to represent states as sets of true atoms. We define ground actions as $a(\bar{c})$, where \bar{c} denotes a mapping from the parameters of a to appropriately typed constants from U. The preconditions and effects of a ground action $a(\bar{c})$ are obtained using the variable assignments in \bar{c} in the usual way.

We consider full observability and deterministic action effects, so that a ground action ga is applicable in a state s iff all of its preconditions are true in s.

Finally, a planning problem consists of a planning domain $\langle \mathcal{R}, A \rangle$, a set U of typed constants, an initial state s_0 and a set $g \subseteq atoms(\mathcal{R}, U)$ denoting the goal condition. A solution to a planning problem is a sequence of ground actions $ga_1, \ldots ga_k$ such that ga_i is applicable in $ga_{i-1}(\cdots ga_1(s_0)\cdots)$ and $ga_k(\cdots ga_1(s_0)\cdots)$ satisfies g.

2.1 Domain Transformation

We introduce the central principle behind using Skolem symbols with an example. Consider an action like grasping in a pick and place domain. Its preconditions include real-valued vectors and preconditions require spatial reasoning. For clarity in this description, we consider a situation where objects are not stacked and assume that the target location where an object has to be placed is clear. The high-level grasp action can be specified as follows: $grasp(c, obj_1)$

precon graspable(c, obj_1) \land robotAt(c) $\land \forall obj_2 \neg obstructs(c, obj_2, obj_1)$ effect in-gripper(obj_1)

In this specification, the variable c represents robot configurations. In order to motivate our approach, consider the effect of this action in a framework like situation calculus (Levesque, Pirri, and Reiter 1998) but using a timestep rather than a situation to represent the fluents:

 $\begin{array}{ll} \forall t, obj_1 \forall c(graspable(c, obj_1, t) & \wedge & robotAt(c, t) & \wedge \\ \forall obj_2 \neg obstructs(c, obj_2, obj_1, t) & \rightarrow & inGripper(obj_1, t + 1)) \end{array}$

Note that this is not the complete *successor state axiom* for *in-gripper*, which will also have to include other actions that affect it and default conditions under which it doesn't change across timesteps. However, this implication is sufficient to illustrate the main representational device we will use. We first use the clearer, logically equivalent form:

 $\begin{array}{lll} \forall t, obj_1(\exists c(graspable(c, obj_1, t) \land robotAt(c, t) \land \\ \forall obj_2 \neg obstructs(c, obj_2, obj_1, t)) \rightarrow inGripper(obj_1, t) \end{array}$

which asserts more clearly that *any* value of c that satisfies the preconditions allows us to achieve the postcondition. We can now use the standard technique of Skolemization to replace occurrences of c with a function of obj_1, t : $\forall t, obj_1((graspable(gp(obj_1, t), obj_1, t))$

$$\wedge robotAt(gp(obj_1), t)$$

 $\land \forall obj_2 \neg obstructs(gp(obj_1), obj_2, obj_1, t))$

 $\rightarrow inGripper(obj_1, t+1)$

where the Skolem function gp(x, t) is just a symbol of type *location* to the discrete planner. Intuitively, it represents a robot configuration corresponding to the grasping-pose of x. This representation will allow the discrete planner to treat the entire problem at a symbolic level, without the need for

creating a problem-specific discretization. A potential problem however, is that the Skolem functions will depend on t, or the current step in the plan. In practice however, at the discrete level, the time argument in a Skolem function $f(\bar{x}, t)$ can be ignored as long as it is possible to recompute (or reinterpret) f when an action's precondition is violated by its existing interpretation during the execution (as is the case in our implementation). We therefore drop the t argument from the Skolem functions in the rest of this paper.

The equivalence with an existential form as described above can be used for each action effect as long as the continuous variable being Skolemized is not free in the subformula on the right of the implication. For instance, we can add the effect that grasping an object removes all obstructions that it had created, regardless of robot configurations. Therefore, to represent the grasp operator for a discrete planning problem, rather than using a discretized space of configurations, we only need to add symbols of the form $f(\bar{o})$ in the planning problem specification, for each object argument tuple \bar{o} consisting of the original objects, or constants in the problem. Since many classical planners don't support functions, they can be reified as objects of the form $f_{-\bar{o}}$ with an associated set of always true relations, e.g. $is_{-}f(f_{-}o_{i}, o_{i})$. The discrete description of grasp thus becomes:

 $\begin{array}{ll} grasp(\ell, obj_1): \\ precon & is_gp(\ell, obj_1) \land graspable(\ell, obj_1) \\ \land robotAt(\ell) \\ \land \forall obj_2 \neg \ obstructs(\ell, obj_2, obj_1) \\ effect & in_gripper(obj_1) \\ \land \forall \ell_2, obj_3 \neg obstructs(\ell_2, obj_1, obj_3) \end{array}$

Here ℓ ranges over the finite set of constant symbols of the form gp_obj_i where obj_i are the original constant symbols in the problem. In this way, regardless of how many samples are used in the lower level process for interpreting these symbol, the discrete planner has a limited problem size to work with while computing the high-level plan.

Finally, consider the only remaining case for an action effect, when a continuous variable occurs freely in the subformula on the right, e.g. $\forall x, c(\varphi_{precon}(x, c) \rightarrow \varphi_{effect}(x, c))$. In this case we don't perform Skolemization. The symbol used for *c* in this case will be an action argument, and must range over the original objects in the domain or those already introduced via Skolemization.

Although this exhausts the set of actions commonly used in PDDL benchmark problems, the accurate description of an action may involve side-effects on symbols not used as its arguments. E.g., the putDown (obj_1, loc_1) action may deterministically introduce obstructions between a number of robot configurations and other objects. We don't encode such side effects in the high-level planning problem specification; these facts are discovered and reported by the lower level when they become relevant.

The putDown action has a similar specification, with the Skolem function $pdp(target_{loc})$ denoting the pose required for putting down an object *o* at location ℓ_1 when the robot is configuration ℓ_2 . $putDown(o, \ell_l, \ell_2)$:

precon is_pdp(ℓ_2, ℓ_1) \land poseForPlacing(ℓ_2, o, ℓ_1) InGripper(o) \land RobotAt(ℓ_2) effect \neg InGripper(o) \land At(o, ℓ_1)

The *moveto* (c_1, c_2) action changes the robot's configuration with the only precondition that it is at c_1 . We assume that the environment is not partitioned and the robot can move between any two collision-free areas. These three actions constitute the high-level pick and place domain.

To summarize, we transform the given planning domain with actions using continuous arguments by replacing occurrences of continuous variables with symbols representing Skolem function application terms whenever the continuous variable occurs in the precondition but not in the effects. Every continuous variable or the symbol replacing one, of type τ gets the new type τ_{sym} . Planning problems over the modified domains are defined by adding finite set of constants for each such τ_{sym} in addition to the constants representing problems in the original domain. The added constants denote function application terms, e.g. gp_obj17, for each physical object and Skolem function application. This increases the size of the input only linearly in the number of original objects if the Skolem functions are unary. Note that the set of Skolem functions itself is fixed by the domain and does not vary across problem instances. The initial state of the problem is described using facts over the set of declared objects, e.g. "is_gp(gp_obj17, obj17)" denoting that the location name gp_obj17 is a grasping pose for obj17 and "obstructs(gp_obj17, obj10, obj17)", denoting that *obj10* obstructs *obj17* when the robot is at *gp_obj17*.

In this formulation, the size of the input problem specification is independent of the sampling-based discretization: we do not need to represent sampled points from the domains of continuous variables.

2.2 Discrete and Continuous States

A low-level state extends a discrete state by associating actual values (interpretations) with each Skolem symbol. Thus, every low-level state corresponds to a unique discrete state representation. Conversely, each discrete state represents a set of possible low-level states.

Since the low-level has complete information about states, it can compute (though not efficiently) the truth value of any atom in the vocabulary at any step, given an interpretation of every Skolem symbol used in the atom. E.g., the low-level can compute the truth value of *obstructs*(*cva1*, *obj10*, *obj17*) where *cval* is a numeric vector representing a robot pose. However, we wish to avoid the determining the value of such predicates as far as possible; the discretization approach discussed in the introduction is undesirable primarily due to its requirement that all such facts be precomputed.

If an action is not executable at a certain step in the lowlevel, then the low-level execution layer computes and reports a set of truth valuations of atoms showing that the action's preconditions do not hold.

Details of the low-level are presented in 3.3. In the next section, we present our overall algorithm and its properties.

3 Discrete and Continuous Planning

As noted above, the initial state for a planning problem is defined using the ground atoms which are true in the initial state. However, the truth values of ground atoms over Skolem symbols like $obstructs(gp_obj17, obj10, obj17)$ are not known initially. Domain specific default truth values are assumed for such atoms. We discuss the relationship between the choice of defaults and completeness of the approach in Sec. 3.2.

3.1 Discrete Planning and the Interface

Our overall algorithm is shown in Alg. 1. In line 2, a classical planner is called with state *s*, which is initialized with the input problem state. If no solution is found, the stored values of Skolem symbols in the lower layer are cleared and all facts using Skolem symbols are reset to their default values using *clear_cache()*. The test in lines 3 and 5 ensure that this is done only once per iteration. If the problem is unsolvable even after clearing the cache, line 6 terminates the algorithm.

In line 7, the computed plan π is passed to the low-level, which uses sampling to estimate the values of Skolem symbols and RRT-based techniques to implement movements representing each action in the plan. Because the high initial state in line 2 used default values for facts over Skolem symbols, π may not be fully executable in the low level. If this is the case, the low-level reports the step/action of the plan that failed and the reason for that failure, represented by an assignment of atoms that violate the action precondition (a minimal violating assignment is sufficient). These components of the low-level are described in detail Sec. 3.3.

As noted earlier, the execution of an operator may void the interpretation of a symbol. E.g., moving an object invalidates the cached grasping pose. If the operators that invalidate a symbol's cached value are known, they can be provided to execute() within the cache voiding structure C. The low-level then reinterprets all affected Skolem symbols after executing an action.

In line 9, the UpdateState() subroutine uses the PDDL descriptions of actions to propagate *s* forward using π until *errStep* and adds *violatedPrecons* to the resulting state. The entire iteration then repeats with the resulting state unless a preset planning time/resource limit is reached.

3.2 Completeness

We now discuss the conditions under which our solution approach is complete. In doing so we show that under certain conditions, effective default assignments for atoms can be ontained easily. We use the notion of *probabilistically-complete* (LaValle 2006) to categorize sampling based algorithms that are guaranteed to find a solution with sufficiently many samples.

The following definition uses the concept of positive and negative occurrences of atoms in formulas. Intuitively an occurrence of an atom in a formula is negative (positive) if it is negated (non-negated) in the negated-normal-form of the formula. This notion of occurrence is sufficient for our purposes as we deal only with problems with finite universes

A	Algorithm 1: Discrete-Continuous Interface						
Input: PDDL state s & domain D; cache-voiding ops C							
1	1 repeat						
2	$\pi \leftarrow \text{classicalPlanner}(s, D)$						
3	if plan not found & not cleared_cache then						
4	clear_cache(); continue						
5	else if plan not found then						
6	return "unsolvable"						
	end						
7	$\langle errStep, violatedPrecons \rangle \leftarrow execute(\pi, s, C)$						
	if errStep is null then						
8	return "success"						
	else						
9	$s = UpdateState(s, \pi, errStep, violatedPrecons)$						
	end						
I	until resource limit reached						

and all quantifiers can be compiled into conjunctions or disjunctions. The following lemma follows from the definition of positive and negative occurrences:

Lemma 1. Suppose an atom p(c) occurs only positively (negatively) in a ground formula φ . If s is an assignment under which φ is true then it must also be true under an assignment s' that makes p(c) true (false) and is the same as s for all other atoms.

Definition 1. A planning domain is uniform wrt a goal g and a set of predicates R if for every $r \in R$,

- 1. Occurrences of r in action preconditions and goal are either always positive, or always negative
- 2. Actions can only add atoms using r in the form (positive or negative) used in preconditions and the goal g

Let P_S be the set of predicates whose atoms may use Skolem symbols as one of their arguments, and let $D = \langle \mathcal{R}, \mathcal{A} \rangle$ be a planning domain that is uniform wrt a goal gand P_S . In the following, consider atoms over a fixed set of constants U. Let s_{part} be an assignment of truth values to atoms over $R \setminus P_S$. Let $s_{default}$ be an assignment of truth values to atoms over P_S , assigning the atoms of each positively (negatively) occurring predicate the truth value true (false).

Proposition 1. The state $s_{part} \cup s_{default}$ has a solution plan for a goal g iff there is some assignment s_0 of atoms over P_S such that $s_{part} \cup s_0$ has a solution plan for reaching g.

Proof. Suppose there is an assignment s_0 under which $s_{part} \cup s_0$ has a solution π and which assigns an atom $p(\bar{c}_1)$ true, while all occurrences of p in action preconditions and g are negative. Consider the assignment s'_0 which assigns $p(\bar{c}_1)$ false, but is otherwise the same as s_0 .

We show that π solves $s_{part} \cup s'_0$ as well. Suppose not, and that $a(\bar{c}_2)$ is the first action in π whose preconditions are not satisfied when π is applied on $s_{part} \cup s'_0$. In this failed execution, $a(\bar{c}_2)$ must have been applied on a state s'_k that differs only on $p(\bar{c}_1)$ from the corresponding state s_k in the execution of π on s_0 . This is because all preceding action applications succeded and have the same deterministic effects in both executions. Let the ground formula representing the preconditions of this application of $a(\bar{c}_2)$ be φ . By Lemma 1, φ must be satisfied by s'_k , and we get a contradiction: $a(\bar{c}_2)$ must be applicable on s'_k . The case for positive defaults is similar.

Thus, in planning problems that are uniform wrt to the set of predicates which use Skolem symbols, it is easy to obtain a default truth assignment for atoms over these predicates, under which the problem is solvable if there is any assignment to those atoms under which it is solvable. The following result provides sufficient conditions under which our approach is complete. Let D be a planning domain and U a set of typed constants. A dead-end for $\langle D, U \rangle$ wrt g is a state over $atoms(\mathcal{R}, U)$ which has no path to g.

We say that two low-level states are physically similar if they differ only on the interpretation of Skolem symbols and atoms using Skolem symbols. We use the symbol [s] to denote the discrete representation of a low-level state s, and $[s]_{default}$ to represent the version of [s] obtained by using the default assignment for atoms using Skolem symbols, and the assignments in [s] for all other atoms.

The following result presents sufficient conditions under which our approach solves any problem which is solvable under some evaluation of Skolem symbols.

Theorem 1. Let D be a planning domain, U a set of typed constants, g a goal, and P_S the set of predicates in D that use Skolem symbols, such that $\langle D, U \rangle$ does not have deadends wrt g and D is uniform wrt g and P_S .

If the low-level sampling routines for Skolem symbols are probabilistically complete, then for any low-level state s, if there exists a physically similar state s' such that [s'] is solvable by the high-level planner then the execution of Alg. 1 with inputs $[s]_{default}$ and D reaches the goal.

Proof. The proof follows from the following two points:

1. For any given interpretation of Skolem symbols represented by a low-level state s'_{ℓ} , Alg. 1 will eventually either discover the truth values of all atoms using Skolem symbols or identify the interpretation as unsolvable.

This is because in every non-final iteration of Alg. 1, line 7 must be executed and will return a non-empty assignment of atoms constituting a violated precondition. The entire set of atoms that violated preconditions are generated from is finite. Once they have been returned by the low level, they will never be generated again for this interpretation. This is because the atom is incorporated into the new discrete state accurately. Subsequent actions can only make it true (false) when it occurs positively (negatively) in preconditions. As a result, the discrete planning layer will never consider an atom that occurs positively to be true if its truth value was returned as false (true) by the low-level—unless an action set it to true in which case the low-level also considers it to be true.

If at any stage a discrete state s with default values for some atoms using Skolem symbols is unsolvable, then replacing them with any other truth values will not make the problem solvable. Thus, if at any stage the classical planner call in line 2 fails for a discrete state s with default values that are inconsistent with the low-level state for some atoms using Skolem symbols, the problem is indeed unsolvable under that interpretation.

2. As a result of probabilistic completeness, the low-level will discover all possible interpretations of Skolem symbols.

The uniform property used in this result plays a role similar to *simple domains* defined by Bonet and Geffner (Bonet and Geffner 2011). Neither categorization is more general than the other. In contrast to simple domains, uniformity is less restrictive in not requiring invariants over initially unknown facts, while simple domains are less restrictive in not enforcing all occurrences of unknown predicates to be of the same polarity.

Extension of this result to more general situations is left for future work.

3.3 Low-Level Execution and Interpretation of Skolem Symbols

Arbitrary sampling techniques could be used for interpreting Skolem functions. For efficiency, we used methods that searched for interpretations satisfying required conditions which cannot be achieved by highlevel planning alone. For instance, the symbol $gp(obj_1)$ needs to satisfy graspable($gp(obj_1), obj_1$) and $\land \forall obj_2 \neg$ obstructs($gp(obj_1), obj_2, obj_1$). Since obstructs facts can be cleared by high-level actions but graspable is not achieved by any, we used a sampling process that achieves graspable.

We drew upon the wealth of approaches developed in the robotics literature for solving navigation, motion planning and obstacle avoidance problems. In this work we assumed that the continuous planner has complete knowledge of the state of the world. This can obtained using, for example, 3D sensors (Hsiao et al. 2010) and map-making techniques (Marder-Eppstein et al. 2010). Both arm motion planning and base movements between poses (i.e. the *moveto* action) are performed using Rapidly Exploring Random Trees (RRTs) (LaValle and Kuffner Jr 2000).

Interpreting Grasping Pose Symbols A robot pose from which an object can be grasped (i.e, a grasping pose or one that satisfies *graspable*) must satisfy two conditions: the robot base should be at a collision free configuration from where the object is reachable, and there must exist a collision-free trajectory from this configuration to one that brings the manipulator into the grasping configuration. Grasping an object requires finding a manipulator configuration that allows the robot to robustly hold it. We precalculated and stored the grasping configurations for each object in OpenRave (these configurations can be computed on-the-fly if necessary).

Our approach for finding a grasping pose is inspired by Probabilistic Roadmaps (Kavraki et al. 1996), and summarized in Algorithm 2. The first step in lines 2-3 is to generate a robot's base configuration that is within a reachability area around the object. The presence of obstacles, the nonlinearity of the robot actuators and the high-dimensionality of the search space render this problem very challenging (Stulp, Fedrizzi, and Beetz 2009; Berenson, Kuffner, and Choset 2008). While pre-calculating poses using inverse

Algorithm 2: Find a Grasping Pose

_	Input: Object o						
1	$ikfound \leftarrow False$						
2	repeat						
3	$p \leftarrow \text{sampleAround}(o)$						
4	$t \leftarrow \text{sampleTorsoHeight}()$						
5	$gps \leftarrow generateGrasps(p, t, o)$						
6	for $gp_i \in gps$ do						
7	if <i>IKsolution</i> (p, t, gp_i) then						
8	$ikfound \leftarrow True; solnPose \leftarrow p$						
9	if $obstaclesFreeArmPath(p, t, gp_i)$ then						
10	return (p, t, gp_i)						
	end						
	end						
	end						
	until max number of attempts reached						
	/* No soln found, return obstructions	*					
11	if <i>ikfound</i> then						
12	2 $objs \leftarrow FindObstructingObjects(o, solnPose)$						
13	return CreateObstructionAtoms(<i>objs</i>)						
	end						
		-					

kinematic (IK) caching could have been used, we found that sampling from a sphere centered around the object yields the same qualitative results. During sampling, poses outside the environment's boundaries or in collision with the environment are discarded.

Collisions can be efficiently calculated by approximating the robot's shape to be either a cube or a cylinder. At line 3 we generate candidate values for the PR2's torso height and retrieve the manipulator grasping configurations at line 4.

The conditions for satisfying a grasping pose are checked in lines 6 & 7. Previously found grasping solutions are cached and retained until an object is moved or when the discrete planner calls *clear_cache()*. If an IK solution exists but no obstacle-free solution is found within a maximum number of iterations then the simulator is used to calculate all the objects in collision with the robot's arm for a computed IK solution. Logical atoms are created using these objects and returned in line 13. This last step requires a few additional inputs like the predicate name and ordering of arguments, but are omitted for readability.

Interpreting Symbols for putDown By substituting the grasping manipulator poses at line 5 with a set of manipulator poses that lay above a tabletop, Algorithm 2 can be generalized to find base poses to put down objects on free areas of a surface. Our current implementation does not take into account collisions between objects held by the robot and the environment. Therefore only obstacle free manipulator configurations are considered when looking for an empty area above a tabletop. Although our simulator provides a realistic implementation of robots' kinematics and collision detection, to speed up the computations we did not simulate physics. Therefore once the robot opens the gripper holding an object we forced the object to fall down and rest on the target surface in a horizontal position.

4 Empirical Evaluation

We tested our proposed approach in two different scenarios, simulated using OpenRave¹. The first scenario involves grasping objects from a cluttered tabletop (see Figure 2). This problem requires the robot to identify objects that prevent a collision free trajectory to the target object, and moving them to a separate side table. The second scenario sees the robot setting up a dining table with three pairs of kitchen objects (see Figure 3). The robot can make use of a tray to help carrying more than one object, but there are constraints on how the objects can be stacked on the tray.

Our approach can be used with any classical planner. For the first experiment we used a well-established satisficing planner, FastForward (FF) (Hoffmann and Nebel 2001). The second experiment required a cost-sensitive planner. We used FastDownward (FD) (Helmert 2006) and always chose the second plan that it produced in its anytime mode if it was produced within 350s, and the first plan otherwise. When reporting the execution time we did not calculate the time required to actually perform an action, e.g. move the robot's base or the arms. In other words the robot instantaneously moves all of its joints.

All the experiments discussed below were carried out on an AMD Opteron dual-core 2GHz machine with 4GB RAM. **Pick and Place Scenario** An instance of the pick-andplace problem is illustrated in Figure 2. To guarantee that objects in the middle of the table are not easily reachable we approximated their shape with over-sized bounding boxes. We modeled this problem as a realistic simulation of a situation the robot might face, where the robot can make free use of the available empty space between objects to reach its goal. Thus a highly cluttered table with 80 objects may have obstructions from nearly every possible graspable pose, but the number of obstructions from each may be small. If grasping fails, the continuous planner returns a list of obstructions.

We performed 30 tests over 3 randomly generated environments. Each environment has the same configuration of static objects (tables and walls) but different configuration of movable objects. Each random environment has between 50 and 80 objects randomly placed on the tabletop. In these experiments we used 200 samples of base configurations generated while searching for grasping poses (outer loop of Alg. 2). Table 1 summarizes the results. Each row in the table is an average of 10 runs. For each run we randomly selected a target object with at least one obstruction to be the grasping goal. As a baseline for comparison we used the time a classical planner would take to solve a version of the same problem with all obstructions precomputed and a discretized set of sampled configurations. For each object in the environment, we sampled 200 poses and kept only those from where objects are graspable. This resulted in 80 to 130 configurations for our problems. The last column of Table 1 shows the time required to pre-calculate all these obstructions, which is significantly higher than the time required for the computation and execution of the whole plan in our

¹Source code is available at https://github.com/ lorenzoriano/OpenRaving



Figure 2: An example execution of a pick and place experiment. From left to right: 1) The initial disposition of the objects on the table. The blue object is the target, while the green objects are the ones identified as obstructing the target. 2) The robot selects the first object to remove. 3) The robot puts the object on a free area of an additional table. 4) The robot removes the second obstruction. 5) The robot removes the third obstruction. 6) The target object is finally reachable.



Figure 3: An example scenario of a dining set-up experiment. From left to right: 1) The initial configuration of the environment. 2) The robot places a plate on the tray. 3) Another plate and a bowl are placed on the tray. 4) The robot carries the tray with four objects to the dining table. 5) The first objects are removed from the top of the stack on the tray and placed on the appropriate locations on the table. 6) The final goal configuration is achieved.

approach.

Dining Table Set-Up Scenario The second experiment illustrates the general applicability of our approach to heterogeneous problems. The robot's task is to move six objects from an initial location to a dining table, as shown in Figure 3. The robot can only move objects between the "kitchen" and "dining" areas by placing them on a tray, moving the tray to the target table and then placing the objects from the tray to their desired configuration. Multiple round-trips are allowed. For this experiment we considered two sets of plates, bowls and mugs.

Placing objects on the tray is an action that can fail, thus generating violated preconditions. This simulates dynamical instability between pairs of objects which is difficult to pre-compute without execution. The lower level enforced this by not allowing a plate to be stacked over other objects, and not allowing anything to be stacked over a mug. These conditions are initially unknown to the high-level planner. We also made the cost of moving across rooms 100 times the cost of other actions.

As picking up the tray requires two-arm manipulation we pre-calculated both manipulators poses for grasping it. Grasping objects, placing them on the tray or on the table uses the same primitives described in Sec. 3.3.

Solving the dining table scenario required 230s, including execution. Given the violations induced by stacking objects on the tray, the classical planner had been called 3 times. The solution plans typically attempt to stack a few objects on the tray and transport them in batches.

5 Related Work

This work is related to, and builds upon a vast history of research in planning and robotics. The field of discrete

#Objs	#Obstrns	Time(s)	#Replan	Precomp(s)
80	2.6	602	2.3	4245
65	2.0	228	2.6	3667
50	1.8	139	2.1	1777

Table 1: Results for the cluttered table scenario. For comparison, times for precomputing obstruction facts are shown in the last column. The central 3 columns are averages of 10 independent runs, showing the number of obstructions, total time for planning and execution and number of calls made to the classical planner.

planning has made several advances in scope and scalability, mainly through the development of efficient methods for computing heuristic functions automatically from a domain definition (Bonet and Geffner 2001; Hoffmann and Nebel 2001; Helmert, Haslum, and Hoffmann 2007). These advances are among the main motivations in developing our approach. An alternate representation of our high-level problem would be a partially observable problem with nondeterministic actions (Bonet and Geffner 2000). However, this is a much harder problem in general (Rintanen 2004) and an offline contingent planning process would be unfeasible in our setting.

Various researchers have investigated the problem of combining discrete and continuous planning. Alami et al. (1998) describe a system architecture that uses a translation module for converting high-level actions into continuous trajectories. Volpe et al. (2001) also use a similar module, referred to as "time-line" interaction. However, these approaches do not discuss specific methods for compiling away continuous parameters of operators in order to facilitate discrete planning. Plaku et al. (2010) propose a sampling based approach for guiding the low-level search process with information from a high-level planner. They also allow non-deterministic actions. However, communication is only one-way: only the first action in the high-level plan is used, and only to bias low-level search process. In contrast, the low-level search process in our approach provides information back to the high-level planner which uses it to generate more relevant plans. Hauser (2010) observes almost the same problems in tasks like pick-and-place. His approach uses a STRIPS like language for specifying actions, but does not use a classical planner for the high-level search. Another approach that combines discrete and continuous planning is described in (Choi and Amir 2009). These authors propose including kinematics constraints among the information provided to the discrete planner. Their approach automatically generates logical actions from the output of a geometric motion planning algorithm. However this approach also has only one-way communication between the discrete and continuous planner. Moreover their approach is difficult to adapt to non-motion planning problems. Wolfe et al. (2010) use angelic hierarchical planning to define a hierarchy of high-level actions over primitive actions. However the high-level actions in their approach have continuous action arguments and need to be sampled. In contrast, we define a whole discrete problem, the size of which is independent of the discretization. They also do not model problems with preconditions like obstructions which require geometric reasoning. Our approach also includes an angelic interpretation: the skolemized functions in our discrete, highlevel operators are assumed to have a value that satisfies the preconditions, and is selected by the agent. Kaelbling et al. (2011) combine discrete and continuous planning by relying upon a very specific regression-based planner with task-specific components. Levihn et al. (2012) present an approach for the problem of navigation among movable obstacles, but unlike our approach, they assume that the effects of discrete, high level actions on obstructions and geometric regions are known a-priori. They also do not address the problem of utilizing general purpose classical planners in conjunction with motion planning.

Approaches for planning modulo theories (Gregory et al. 2012) (PMT) and planning with semantic attachments (Hertle et al. 2012) address similar high level problems of planning in hybrid domains. However, these approaches use an extended planning language and require appropriately extended planners. Both of these approaches require deterministic functions for computing "external" predicates during search, and do not address the representational problem of using a classical planner to plan with actions and fluents that use continuous arguments. In contrast, our approach only requires information from the lower level when a high level plan fails and works with arbitrary classical planners. Our approach can also be used in situations where a model of the effects of low level actions is not available, and the action execution failures cannot be predicted accurately (as in the stacking problem) and are only known when they occur. As a consequence of this generality, the completeness of our approach depends on the problem domain (the presented analysis provides a sufficient but not necessary condition).

Grasping objects in a cluttered environment is still an open problem in robotics. In (Dogar and Srinivasa 2011)

an approach is proposed that replaces pick-and-place with pushing objects away from the desired trajectory. This represents an interesting solution to the reachability problem that we aim at exploring as an application of our framework.

6 Conclusions

In this paper we presented an approach for planning in robotics that relies upon the synergy between a classical and a continuous planner. One of the advantages of our proposed approach is that not all the facts need to be pre-computed before planning, but only the ones that are required by the execution. The classical planner is also able to efficiently represent these facts without having to deal with representations that grow with the sampling process used in the lower layer. As our experiments show, this is significantly more efficient than solving precomputed discrete problems with sampling-based discretizations. Our solution also addressed a broad underlying problem in robotics: that the truth value of many predicates can only be found when the robot tries to perform an action. Both of our test scenarios featured this problem.

Although we simulated dynamical instability by using fixed rules, a different implementation could utilize a physics simulator for predicting if an object is unstable when placed over a stack. It could thus prune away actions that a planner alone would not know were wrong. While we conducted our main experiments in a pick-and-place scenario, the second experiment shows that our proposed approach is not limited to this particular application. Introducing costs, handling constraints and plans with heterogeneous actions are supported by our algorithm through its modularity.

Future work will include a more thorough analysis of optimality and termination guarantees on replanning. We also plan to study situations where action costs, in addition to facts, are made available during execution. Although our simulation is extremely faithful in terms of the capabilities and limitations of the robot and in terms of construction of the problem scenarios, the most natural next step in this work will be to test the plans with a real robot.

Acknowledgments

We thank Malte Helmert and Joerg Hoffmann for providing versions of their planners with verbose outputs that were very helpful in our implementation. We also thank John Schulman for helpful comments. This work was supported by the NSF under Grant IIS-0904672.

References

Alami, R.; Chatila, R.; Fleury, S.; Ghallab, M.; and Ingrand, F. 1998. An architecture for autonomy. *International Journal Of Robotics Research* 17:315–337.

Berenson, D.; Kuffner, J.; and Choset, H. 2008. An optimization approach to planning for mobile manipulation. 2008 IEEE International Conference on Robotics and Automation 1187–1192.

Bonet, B., and Geffner, H. 2000. Planning with incomplete information as heuristic search in belief space. In *Proc. of*

the 6th International Conference on Artificial Intelligence Planning and Scheduling, 52–61.

Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artif. Intell.* 129(1-2):5–33.

Bonet, B., and Geffner, H. 2011. Planning under partial observability by classical replanning: Theory and experiments. In *IJCAI*, 1936–1941.

Choi, J., and Amir, E. 2009. Combining planning and motion planning. In *Robotics and Automation*, 2009. *ICRA'09*. *IEEE International Conference on*, 238–244. IEEE.

Diankov, R. 2010. *Automated Construction of Robotic Manipulation Programs*. Ph.D. Dissertation, Carnegie Mellon University, Robotics Institute.

Dogar, M., and Srinivasa, S. 2011. A framework for pushgrasping in clutter. *Robotics: Science and Systems VII*.

Fikes, R., and Nilsson, N. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. Technical report, AI Center, SRI International. SRI Project 8259.

Fox, M., and Long, D. 2003. PDDL2.1: an extension to pddl for expressing temporal planning domains. *J. Artif. Int. Res.* 20(1):61–124.

Gregory, P.; Long, D.; Fox, M.; and Beck, J. C. 2012. Planning modulo theories: Extending the planning paradigm. In *ICAPS*.

Hauser, K. 2010. Task planning with continuous actions and nondeterministic motion planning queries. In *Proc. of AAAI Workshop on Bridging the Gap between Task and Motion Planning*.

Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible abstraction heuristics for optimal sequential planning. In *Proc. of the 17th International Conference on Automated Planning and Scheduling*, 176–183.

Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Hertle, A.; Dornhege, C.; Keller, T.; and Nebel, B. 2012. Planning with semantic attachments: An object-oriented view. In *ECAI*, 402–407.

Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.

Hsiao, K.; Chitta, S.; Ciocarlie, M.; and Jones, E. 2010. Contact-reactive grasping of objects with partial shape information. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, 1228–1235. IEEE.

Kaelbling, L. P., and Lozano-Pérez, T. 2011. Hierarchical task and motion planning in the now. In *Proc. IEEE International Conference on Robotics and Automation*, 1470–1477.

Kavraki, L.; Svestka, P.; Latombe, J.; and Overmars, M. 1996. Probabilistic roadmaps for path planning in highdimensional configuration spaces. *Robotics and Automation, IEEE Transactions on* 12(4):566–580.

LaValle, S., and Kuffner Jr, J. 2000. Rapidly-exploring random trees: Progress and prospects. LaValle, S. M. 2006. *Planning algorithms*. Cambridge University Press.

Levesque, H. J.; Pirri, F.; and Reiter, R. 1998. Foundations for the situation calculus. *Electronic Transactions on Artificial Intelligence* 2:159–178.

Levihn, M.; Scholz, J.; and Stilman, M. 2012. Hierarchical decision theoretic planning for navigation among movable obstacles. In *Workshop on the Algorithmic Foundations of Robotics*, 19–35.

Marder-Eppstein, E.; Berger, E.; Foote, T.; Gerkey, B. P.; and Konolige, K. 2010. The Office Marathon: Robust Navigation in an Indoor Office Environment. In *International Conference on Robotics and Automation*.

Plaku, E., and Hager, G. D. 2010. Sampling-based motion and symbolic action planning with geometric and differential constraints. In *IEEE International Conference on Robotics and Automation*, 5002–5008.

Rintanen, J. 2004. Complexity of planning with partial observability. In *Proc. of the 14th International Conference on Automated Planning and Scheduling*, 345–354.

Stulp, F.; Fedrizzi, A.; and Beetz, M. 2009. Action-related place-based mobile manipulation. 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems 3115–3120.

Talamadupula, K.; Benton, J.; Schermerhorn, P. W.; Kambhampati, S.; and Scheutz, M. 2010. Integrating a closed world planner with an open world robot: A case study. In *Proc. of AAAI*.

Volpe, R.; Nesnas, I.; Estlin, T.; Mutz, D.; Petras, R.; and Das, H. 2001. The CLARAty architecture for robotic autonomy. In *Proc. of IEEE Aerospace Conference*, 121–132.

Wolfe, J.; Marthi, B.; and Russell, S. J. 2010. Combined task and motion planning for mobile manipulation. In *Proc. of the International Conference on Automated Planning and Scheduling*, 254–258.

Yoon, S. W.; Fern, A.; and Givan, R. 2007. FF-replan: A baseline for probabilistic planning. In *Proc. of the International Conference on Automated Planning and Scheduling*, 352–.