

MIT Open Access Articles

*Sensor-based reactive symbolic
planning in partially known environments*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Vasilopoulos, Vasileios, et al., "Sensor-based reactive symbolic planning in partially known environments." 2018 IEEE International Conference on Robotics and Automation (ICRA 2018), Brisbane, Australia, 2018 (Piscataway, N.J.: IEEE, 2018): p. 5683-90 doi 10.1109/ICRA.2018.8460861

As Published: 10.1109/ICRA.2018.8460861

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/125881>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Sensor-Based Reactive Symbolic Planning in Partially Known Environments

Vasileios Vasilopoulos, William Vega-Brown, Omur Arslan, Nicholas Roy, Daniel E. Koditschek

Abstract—This paper considers the problem of completing assemblies of passive objects in nonconvex environments, cluttered with convex obstacles of unknown position, shape and size that satisfy a specific separation assumption. A differential drive robot equipped with a gripper and a LIDAR sensor, capable of perceiving its environment only locally, is used to position the passive objects in a desired configuration. The method combines the virtues of a deliberative planner generating high-level, symbolic commands, with the formal guarantees of convergence and obstacle avoidance of a reactive planner that requires little onboard computation and is used online. The validity of the proposed method is verified both with formal proofs and numerical simulations.

I. INTRODUCTION

In this paper, we address a very specific instance of the Warehouseman’s Problem [1] as a challenging setting in which to advance the formal integration of deliberative and reactive modes of robot assembly planning and control. We posit a planar disk-shaped robot with velocity controlled unicycle kinematics placed in an indoor environment with known floor-plan, cluttered with convex obstacles of unknown number and placement. The robot’s task is to bring a collection of known disk-shaped objects from their initial placement to their prescribed destination by approaching, attaching and then pushing it into place, making sure to avoid any collisions with the known walls, other objects and unanticipated obstacles along the way.

A. Motivation and Related Work

The problem of using a higher-level planner to inform sub-goals of a lower-level planner has been examined previously, and we build on prior work in hybrid systems and task and motion planning. However, most work has focused on ad hoc abstractions that perform well empirically. For example, Wolfe et al. [2] use a task hierarchy to guide the search for a low-level plan by expanding high-level plans in a best-first way. This approach guarantees hierarchical optimality: it will generate the best plan which can be represented in a given task hierarchy. Ensuring optimality has always been difficult to achieve due to computational complexity. Berenson et al [3] and Konidaris et al [4] use specific




Fig. 1: A depiction of an intermediate stage of an assembly process. The robot is tasked to move two objects from their start to their final configuration using a gripper and a LIDAR. The deliberative planner outputs a reference path (purple) which the reactive planner has to follow, while avoiding the unexpected obstacles (grey) in the (potentially) nonconvex workspace. The resulting piecewise differentiable object trajectory for one object is shown in red.

formulations of hierarchy without guaranteeing optimality. Kaelbling and Lozano-Perez [5] avoid the computational cost by committing to decisions at a high level of abstraction, before a full low-level plan is available. Vega-Brown and Roy [6] provided a further step towards tractable planning that incorporated complex kinematic constraints, but no general approach exists for incorporating complex dynamics.

On the other hand, planning the rearrangement of movable objects has long been known to be algorithmically hard (e.g., PSPACE hardness was established in [1]) and a lively contemporary literature [7], [8] continues to explore conditions under which the additional complexity of planning the grasps results in a deterministically undecidable problem. While that interface has been understood to be crucial for decades [9], the literature on reactive approaches to this problem has been far more sparse. Purely reactive formulations and provably correct solutions to partial [10] and then complete [11] scalar versions of the problem motivated the empirical study [12] of a simplified proxy of the planar problem we consider here.

Vasileios Vasilopoulos is with the Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA 19104, vvasilo@seas.upenn.edu

William Vega-Brown and Nicholas Roy are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, {wrvb,nickroy}@csail.mit.edu

Omur Arslan and Daniel E. Koditschek are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, {omur,kod}@seas.upenn.edu

B. Contributions

We present a provably correct architecture (Fig. 3) for planning and executing a successful solution to this Warehouseman’s problem by decomposition into an offline “deliberative” planning module and an online “reactive” execution module. The deliberative planner, adapted from the probabilistically complete (and optimal) algorithm of [13], is assigned the job of finding an assembly plan, while the reactive planner accepts each next step of that planned sequence, and uses online (LIDAR-style) sensory measurements to avoid the unanticipated obstacles (as well as the known walls and objects) by switching between following the deliberative planner’s specified path or instead following a sensed wall. The wall following algorithm is guaranteed to maintain the robot distance from the wall within some specified bounds, while making progress along the wall boundary.

After imposing specific constraints on how tightly packed the unknown obstacles and the known objects’ initial and final configurations can be, we prove that the hybrid control scheme generated by this reactive planner must succeed in achieving any specified step of the deliberative sequence with no collisions along the way.

C. Organization of the paper

The paper is organized as follows. Section II describes the problem and summarizes our approach. Section III gives a brief outline of the high-level deliberative planner that generates the sequence of appropriate symbolic commands to accomplish the task at hand, without any information about the internal obstacles. Section IV describes the fundamental idea of reactively switching between a path following and a wall following mode, for both a holonomic and a nonholonomic robot, while Section V extends our reactive ideas to the navigation problem of a nonholonomic robot grasping a passive object and using its sensor to position it at a desired location. Section VI combines the ideas from the previous two sections and describes the low-level, online implementation of the symbolic action command set. Section VII presents illustrative numerical examples for the ideas presented, while Section VIII summarizes our observations and ideas for future work.

II. PROBLEM FORMULATION

In this work, we consider a first-order, nonholonomically-constrained, disk-shaped robot, centered at $\mathbf{x} \in \mathbb{R}^2$ with radius $r \in \mathbb{R}_{>0}$ and orientation $\psi \in S^1$, using a gripper to move circular objects in a closed, compact, not necessarily convex workspace $\mathcal{W} \subset \mathbb{R}^2$ as shown in Fig. 1, whose boundary $\partial\mathcal{W}$ is assumed to be known. The robot dynamics are described by

$$(\dot{\mathbf{x}}, \dot{\psi}) = \mathbf{B}(\psi)\mathbf{u}_{ku} \quad (1)$$

with $\mathbf{B}(\psi) = \begin{bmatrix} \cos \psi & \sin \psi & 0 \\ 0 & 0 & 1 \end{bmatrix}^T$ the differential constraint matrix and $\mathbf{u}_{ku} = (v, \omega)$ the input vector¹ consisting of

¹Throughout this paper, we will use the ordered set notation $(*, *, \dots)$ and the matrix notation $[* \quad * \quad \dots]^T$ for vectors interchangeably.

a linear and an angular command. The robot is assumed to possess a LIDAR, positioned at \mathbf{x} , with a 360° angular scanning range and a fixed sensing range $R \in \mathbb{R}_{>0}$ and is tasked with moving each of the $n \in \mathbb{N}$ movable disk-shaped objects, centered at $\mathbf{p} := (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n) \in \mathcal{W}^n$ with a vector of radii $(\rho_1, \rho_2, \dots, \rho_n) \in (\mathbb{R}_{>0})^n$, from its initial configuration to a user-specified goal configuration $\mathbf{p}^* := (\mathbf{p}_1^*, \mathbf{p}_2^*, \dots, \mathbf{p}_n^*) \in \mathcal{W}^n$. We assume that both the initial configuration and the target configuration of the objects are known. In addition to the known boundary of the workspace $\partial\mathcal{W}$, the workspace is cluttered by an unknown number of fixed, disjoint, convex obstacles of unknown position and size, denoted by $\mathcal{O} := (O_1, O_2, \dots)$. To simplify the notation, also define $\mathcal{O}_w := \mathcal{O} \cup \partial\mathcal{W}$.

We adopt the following assumptions to guarantee that any robot-object pair can go around any obstacle in the workspace along any possible direction, introduced only to facilitate the proofs of our formal results, without being necessary for the existence of some solution to the problem.

Assumption 1 (Obstacle separation) *The obstacles \mathcal{O} in the workspace are separated from each other by clearance² of at least $d(O_i, O_j) > 2(r + \max_k \rho_k)$, $i \neq j$, with k an index spanning the set of movable objects. They are also separated from the boundary of the (potentially nonconvex) workspace \mathcal{W} by at least $d(O_i, \partial\mathcal{W}) > 2(r + \max_k \rho_k)$ for all i .*

Assumption 1 means that there exists $\eta \in \mathbb{R}_{>0}$ such that

$$\eta = \min \left\{ \min_{\substack{i,j \\ i \neq j}} d(O_i, O_j), \min_i d(O_i, \partial\mathcal{W}) \right\} \quad (2)$$

and $\eta > 2(r + \max_k \rho_k)$.

Also, in order to ensure successful positioning of all the objects to their target configuration using reactive control schemes, it is convenient to impose a further constraint on how tightly packed the desired goal configuration can be.

Assumption 2 (Admissible object goals) *For any object $i \in \{1, \dots, n\}$, $d(\mathbf{p}_i^*, \mathcal{O}_w) > \rho_i + 2r$.*

The robot’s gripper can either be engaged or disengaged; we will write $g = 1$ when the gripper is engaged and $g = 0$ when it is disengaged.

In order to accomplish the task of bringing every object to its designated goal position, we endow the deliberative planner with a set of three symbolic output action commands:

- **MOVETOOBJECT**(i, \mathcal{P}) instructing the robot to move and grasp the object i along the piecewise continuously differentiable path $\mathcal{P} : [0, 1] \rightarrow \mathcal{W}$ such that $\mathcal{P}(0) = \mathbf{x}$ and $\mathcal{P}(1) = \mathbf{p}_i$.
- **POSITIONOBJECT**(i, \mathcal{P}) instructing the robot to push the (assumed already grasped) object i toward its designated goal position, \mathbf{p}_i^* , along the piecewise continuously differentiable path $\mathcal{P} : [0, 1] \rightarrow \mathcal{W}$ such that $\mathcal{P}(0) = \mathbf{p}_i$ and $\mathcal{P}(1) = \mathbf{p}_i^*$.

²Here the clearance between two sets A and B is defined as $d(A, B) := \inf\{\|\mathbf{a} - \mathbf{b}\| \mid \mathbf{a} \in A, \mathbf{b} \in B\}$




Fig. 2: A depiction of a disk-shaped robot with radius r (grey) moving a disk-shaped object with radius ρ_i (yellow).




Fig. 3: An outline of the control approach followed in order to position the objects. A high-level, deliberative planner outputs a sequence of symbolic actions that are realized and executed sequentially in low-level using a reactive controller.

- $\text{MOVE}(\mathcal{P})$ instructing the robot to move along the piecewise continuously differentiable path $\mathcal{P} : [0, 1] \rightarrow \mathcal{W}$ such that $\mathcal{P}(0) = \mathbf{x}$.

This symbolic command set, comprising the interface between the deliberative and reactive components of our planner enforces the following problem decomposition into the complementary pair:

- 1) Find a *symbolic plan*, i.e a sequence of symbolic actions whose successful implementation is guaranteed to complete the task.
- 2) Implement each of the symbolic actions using the appropriate commands \mathbf{u}_{ku} according to the robot's equations of motion shown in (1), while avoiding the perceived unanticipated by the deliberative planner obstacles.

Fig. 3 depicts this problem decomposition and the associated interface between the deliberative and reactive components of our architecture.

III. DELIBERATIVE PLANNER

In order to obtain plans suitable for the reactive planner to track, we use a high-level planner that combines the factored orbital random geometric graph (FORGG) construction [6] with the approximate angelic A* (AAA*) search algorithm [14]. FORGG extends the asymptotic optimality guarantees of the PRM* algorithm to problems involving discontinuous differential constraints like contact and object manipulation. Searching this planning graph using conventional methods like A* is computationally expensive, due to the size of

search space. To facilitate efficient search, we employ the angelic semantics developed by Marthi *et al.* [13] to encode bounds on the possible cost of sets of possible plans. AAA* uses these bounds to guide the search, allowing large parts of the search space to be pruned away and accelerating the search for a near-optimal high-level plan.

With this construction, the deliberative planner is supplied with the initial position and size of the robot and the objects to be placed, along with any information assumed to be known (boundary of the workspace, walls, interior obstacles etc.) and outputs a series of symbolic action commands (MOVEOBJECT, POSITIONOBJECT, MOVE) each associated with a collision-free path \mathcal{P} in order to accomplish the task at hand.

IV. REACTIVE PLANNING FOR SINGLE ROBOTS

In this section we describe the (low-level) reactive algorithms which guarantee collision avoidance and (almost) global convergence³ to the plan provided by the (high-level) deliberative planner, described in Section III. First, we focus on the navigation problem of a single (fully actuated or nonholonomically-constrained) robot, using tools from [16] and [17], and we will show in Section V how to extend these principles for the case of gripping contact.

A. Doubly-reactive planner for holonomic robots

First we consider a fully actuated disk-shaped robot centered at $\mathbf{x} \in \mathbb{R}^2$ with radius $r > 0$, moving in a closed-convex environment (denoted by $\mathcal{W} \subset \mathbb{R}^2$) towards a goal location $\mathbf{x}^* \in \mathbb{R}^2$. Although we use a differential drive robot for our assembly problem here, we find it useful to present the basic algorithm for fully actuated robots, especially since it will be used in Section V. The robot dynamics are described by

$$\dot{\mathbf{x}} = \mathbf{u}(\mathbf{x}) \quad (3)$$

with $\mathbf{u} \in \mathbb{R}^2$ the input. The sensory measurement of the LIDAR at $\mathbf{x} \in \mathcal{W}$ is modeled as in [17] by a polar curve $\rho_{\mathbf{x}} : (-\pi, \pi] \rightarrow [0, R]$ as follows⁴

$$\rho_{\mathbf{x}}(\theta) := \min \left(\begin{array}{c} R \\ \min \{ \|\mathbf{p} - \mathbf{x}\| \mid \mathbf{p} \in \partial\mathcal{W}, \text{atan2}(\mathbf{p} - \mathbf{x}) = \theta \} \\ \min_i \{ \|\mathbf{p} - \mathbf{x}\| \mid \mathbf{p} \in O_i, \text{atan2}(\mathbf{p} - \mathbf{x}) = \theta \} \end{array} \right) \quad (4)$$

We will also use the definitions of free space \mathcal{F} , line-of-sight local workspace $\mathcal{LW}_{\mathcal{L}}(\mathbf{x})$ and line-of-sight local free space $\mathcal{LF}_{\mathcal{L}}(\mathbf{x})$ at \mathbf{x} from [17].

Under the preceding definitions, it is shown in [17] that the control law

$$\mathbf{u}(\mathbf{x}) = -k(\mathbf{x} - \Pi_{\mathcal{LF}_{\mathcal{L}}(\mathbf{x})}(\mathbf{x}^*)), k \in \mathbb{R} \quad (5)$$

with $\Pi_A : \mathbb{R}^2 \rightarrow A$ denoting the projection function onto a convex subset $A \subseteq \mathbb{R}^2$, i.e

$$\Pi_A(\mathbf{q}) := \arg \min_{\mathbf{a} \in A} \|\mathbf{a} - \mathbf{q}\| \quad (6)$$

³It is well-known that the basin of a point attractor in a non-contractible space must exclude a set of measure zero [15].

⁴See [17] for a discussion on the choice of LIDAR range R to avoid obstacle occlusions.

asymptotically drives almost all configurations in \mathcal{F} to the goal \mathbf{x}^* while avoiding obstacles and not increasing the Euclidean distance to the goal along the way.

B. Reactive path following

For a fixed goal \mathbf{x}^* , the reactive control law in (5) guarantees convergence only for convex workspaces (punctured by obstacles).

Therefore, inspired by [18], we apply the idea from Section IV-A to the problem of a robot following a navigation path $\mathcal{P} : [0, 1] \rightarrow \mathring{\mathcal{F}}$, that joins a pair of initial and final configurations $\mathbf{x}^0, \mathbf{x}^1 \in \mathring{\mathcal{F}}$ in a potentially nonconvex workspace and lies in the interior of the free space, i.e. $\mathcal{P}(0) = \mathbf{x}^0, \mathcal{P}(1) = \mathbf{x}^1$ and $\mathcal{P}(\alpha) \in \mathring{\mathcal{F}}, \forall \alpha \in [0, 1]$.

As demonstrated in [18], the *projected-path goal* $\mathcal{P}(\alpha^*)$ with α^* determined as⁵

$$\alpha^* = \max\{\alpha \in [0, 1] \mid \mathcal{P}(\alpha) \in B(\mathbf{x}, d(\mathbf{x}, \partial\mathcal{F}))\} \quad (7)$$

replaces \mathbf{x}^* in (5) as the target goal position and is constantly updated as the agent moves along the path. Note that in the LIDAR-based setting presented here, the distance of the agent from the boundary of the free space $d(\mathbf{x}, \partial\mathcal{F})$ can easily be determined as

$$d(\mathbf{x}, \partial\mathcal{F}) = \min_{\theta} \rho_{\mathbf{x}}(\theta) - r \quad (8)$$

C. Reactive wall following

As described in Section II and shown in Fig. 1, the path \mathcal{P} might not lie in the free space since the deliberative planner is only aware of the boundary of the workspace and not of the position or size of the internal obstacles. For this reason, we present here a novel control law for reactive wall following, inspired from the ‘‘bug algorithm’’ [19], that exhibits desired formal guarantees.

The wall following law is triggered by saving the current index α_s^* of the path \mathcal{P} when the distance of the agent from the boundary of its free space, given in (8), drops below a small critical value ϵ , i.e. when $d(\mathbf{x}, \partial\mathcal{F}) < \epsilon$. This would imply that the robot enters a ‘‘danger zone’’ within the vicinity of an unexpected obstacle. The goal now would be to follow the boundary of that obstacle without losing it, in order to find the path again.

Therefore, the robot first needs to select a specific direction to consistently follow the boundary of the obstacle along that direction. Since our problem is planar, there are only two possible direction choices: clockwise (CW) or counterclockwise (CCW). Also, since the robot has only local information about the obstacle based on the current LIDAR readings, a greedy selection of the wall following direction is necessary.

Let $\theta_m \in (-\pi, \pi]$ be the LIDAR angle such that $\rho_{\mathbf{x}}(\theta_m) = \min_{\theta} \rho_{\mathbf{x}}(\theta)$ corresponds to the minimum distance from the blocking obstacle. Let $\mathbf{n}_w(\mathbf{x}) := -(\cos \theta_m, \sin \theta_m)$ denote the normal vector to the boundary of the obstacle at the

⁵Here $B(\mathbf{q}, t) := \{\mathbf{p} \in \mathcal{W} \mid \|\mathbf{p} - \mathbf{q}\| \leq t\}$, i.e. the ball of radius t centered at \mathbf{q} .




Fig. 4: An example of computing the wall following local free space $\mathcal{L}\mathcal{F}_w(\mathbf{x})$ (cyan) as the intersection of the local free space $\mathcal{L}\mathcal{F}_L(\mathbf{x})$ (green) and the offset disk \mathcal{D}_w (magenta) for a robot with radius r positioned at \mathbf{x} , encountering an obstacle within its LIDAR footprint $L_{ft}(\mathbf{x})$ (red).

point of minimum distance and $\mathbf{t}_w(\mathbf{x}) = \mathbf{J} \mathbf{n}_w(\mathbf{x})$, with $\mathbf{J} := \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$, the corresponding tangent vector.

Our proposed method uses the inner product $\mathbf{t}_{w,0} \cdot \mathbf{t}_{\mathcal{P}}(\alpha_s^*)$, with $\mathbf{t}_{w,0}$ denoting the tangent vector to the boundary of the obstacle at the beginning of the wall following phase and $\mathbf{t}_{\mathcal{P}}(\alpha_s^*)$ the tangent vector of the path \mathcal{P} at α_s^* . Then, the value of a variable a is set to 1 for CCW motion and to -1 for CW motion (fixed for all future time) according to

$$a = \begin{cases} 1, & \text{if } \mathbf{t}_{w,0} \cdot \mathbf{t}_{\mathcal{P}}(\alpha_s^*) \geq 0 \\ -1, & \text{if } \mathbf{t}_{w,0} \cdot \mathbf{t}_{\mathcal{P}}(\alpha_s^*) < 0 \end{cases} \quad (9)$$

since $\mathbf{t}_w(\mathbf{x})$ has counterclockwise direction around the obstacle by construction.

Define the *offset disk at x*

$$\mathcal{D}_w(\mathbf{x}) := \{\mathbf{p} \in \mathcal{W} \mid \|\mathbf{p} - \mathbf{x}_{\text{offset}}(\mathbf{x})\| \leq \epsilon\} \quad (10)$$

with ϵ selected according to Assumption 1 to satisfy

$$0 < \epsilon < \frac{1}{2} \left[\eta - 2(r + \max_j \rho_j) \right] \quad (11)$$

with η given in (2) and

$$\mathbf{x}_{\text{offset}}(\mathbf{x}) := \mathbf{x} - (\rho_{\mathbf{x}}(\theta_m) - r) \mathbf{n}_w(\mathbf{x}) \quad (12)$$

Then define the *wall following local free space* $\mathcal{L}\mathcal{F}_w(\mathbf{x})$ at \mathbf{x}

$$\mathcal{L}\mathcal{F}_w(\mathbf{x}) := \mathcal{L}\mathcal{F}_L(\mathbf{x}) \cap \mathcal{D}_w(\mathbf{x}) \quad (13)$$

Since $\mathcal{L}\mathcal{F}_L(\mathbf{x})$ is convex [17], $\mathcal{L}\mathcal{F}_w(\mathbf{x})$ is convex as the intersection of convex sets.

The *wall following law* is then given as

$$\mathbf{u}(\mathbf{x}) = -k(\mathbf{x} - \mathbf{x}_p(\mathbf{x})) \quad (14)$$

with

$$\mathbf{x}_p(\mathbf{x}) := \mathbf{x}_{\text{offset}}(\mathbf{x}) + \frac{\epsilon}{2} \mathbf{n}_w(\mathbf{x}) + a \frac{\epsilon\sqrt{3}}{2} \mathbf{t}_w(\mathbf{x}) \quad (15)$$

Lemma 1 If $\mathbf{x} \in \mathcal{F}$ and $d(\mathbf{x}, \partial\mathcal{F}) < \epsilon$ with ϵ chosen according to (11):

- (i) The wall following free space $\mathcal{L}\mathcal{F}_w(\mathbf{x})$ contains \mathbf{x} in its interior.
- (ii) $\mathcal{L}\mathcal{F}_w(\mathbf{x}) = \mathcal{D}_w(\mathbf{x}) \cap H_{\mathbf{n}_w}(\mathbf{x})$ with $H_{\mathbf{n}_w}(\mathbf{x})$ the half space

$$H_{\mathbf{n}_w}(\mathbf{x}) = \{\mathbf{p} \in \mathcal{W} \mid (\mathbf{p} - \mathbf{x}_h(\mathbf{x})) \cdot \mathbf{n}_w \geq 0\}$$

$$\text{and } \mathbf{x}_h(\mathbf{x}) = \mathbf{x} - \frac{1}{2}(\rho_{\mathbf{x}}(\theta_m) - r).$$

- (iii) The point $\mathbf{x}_p(\mathbf{x})$ lies on the boundary of $\mathcal{L}\mathcal{F}_w(\mathbf{x})$.

Proof. Included in [20]. ■

Proposition 1 With the choice of ϵ in (11), the wall following law in (14) has the following properties:

- (i) It is piecewise continuously differentiable.
- (ii) It generates a unique continuously differentiable flow, defined for all future time.
- (iii) It has no stationary points.
- (iv) The free space \mathcal{F} is positively invariant under its flow.
- (v) Moreover, the set $\left\{ \mathbf{p} \in \mathcal{W} \mid \frac{\epsilon}{2} < d(\mathbf{p}, \partial\mathcal{F}) < \epsilon \right\}$ is positively invariant under its flow.

Proof. Included in [20]. ■

We find it useful to include the following definition

Definition 1 The rate of progress along the boundary of the observed obstacle at \mathbf{x} is defined as

$$\sigma(\mathbf{x}) := \frac{\mathbf{u}(\mathbf{x}) \cdot \mathbf{t}_w(\mathbf{x})}{\|\mathbf{u}(\mathbf{x})\|} \quad (16)$$

By combining all these results, we arrive at the following Theorem:

Theorem 1 With a selection of ϵ as in (11), the wall following law in (14) has no stationary points, leaves the robot's free space \mathcal{F} positively invariant under its unique continuously differentiable flow, and steers the robot along the boundary of a unique obstacle in \mathcal{O} in a clockwise or counterclockwise fashion (according to the selection of a in (9)) with a nonzero rate of progress σ , while maintaining a distance of at most $(r + \epsilon)$ and no less than $(r + \frac{\epsilon}{2})$ from it.

Proof. Included in [20]. ■

The robot exits the wall following mode and returns to the path following mode once it encounters the path again, i.e when $\alpha^* = \max\{\alpha \in [0, 1] \mid \mathcal{P}(\alpha) \in B(\mathbf{x}, d(\mathbf{x}, \partial\mathcal{F}))\} > \alpha_s^*$. An immediate Corollary of Theorem 1, along with path continuity of \mathcal{P} and Assumptions 1 and 2 is the following:

Corollary 1 If the robot enters the wall following mode, it will exit it in finite time and return to the path following mode.

Finally, since both the path following law [18] and the wall following law generate continuously differentiable flows, we find it useful to explicitly state the following result, in the sense of sequential composition [21].

Theorem 2 In a workspace where Assumption 1 is satisfied, any composition of path following and wall following phases generates a unique piecewise continuously differentiable flow for \mathbf{x} , defined for all future time.

D. Extension to nonholonomic robots

As shown in [17], the preceding results can easily be extended for the case of a differential-drive robot driving towards a goal \mathbf{x}^* , whose dynamics are given in (1). Here, we will use a slightly different than [17] control law since the robot possesses a gripper and must only move in the forward direction to grasp objects. The following inputs are used

$$v = \max \left\{ -k \begin{bmatrix} \cos \psi \\ \sin \psi \end{bmatrix}^T (\mathbf{x} - \Pi_{\mathcal{L}\mathcal{F}_v(\mathbf{x})}(\mathbf{x}^*)), 0 \right\} \quad (17)$$

$$\omega = -k \operatorname{atan2}(\beta_2, \beta_1) \quad (18)$$

with

$$\beta_1 = \begin{bmatrix} \cos \psi \\ \sin \psi \end{bmatrix}^T \left(\mathbf{x} - \frac{\Pi_{\mathcal{L}\mathcal{F}_\omega(\mathbf{x})}(\mathbf{x}^*) + \Pi_{\mathcal{L}\mathcal{F}_\epsilon(\mathbf{x})}(\mathbf{x}^*)}{2} \right) \quad (19)$$

$$\beta_2 = \begin{bmatrix} -\sin \psi \\ \cos \psi \end{bmatrix}^T \left(\mathbf{x} - \frac{\Pi_{\mathcal{L}\mathcal{F}_\omega(\mathbf{x})}(\mathbf{x}^*) + \Pi_{\mathcal{L}\mathcal{F}_\epsilon(\mathbf{x})}(\mathbf{x}^*)}{2} \right) \quad (20)$$

in order to constrain the robot motion to forward only and align with the desired target respectively. Here $\mathcal{L}\mathcal{F}_v(\mathbf{x}), \mathcal{L}\mathcal{F}_\omega(\mathbf{x})$ are used as in [17].

Based on the preceding analysis, for a differential drive robot, we will use $\mathbf{x}^* = \mathcal{P}(\alpha^*)$ (with α^* shown in (7)) in the path following mode and $\mathbf{x}^* = \mathbf{x}_p(\mathbf{x})$ in the wall following mode. The following Theorem summarizes the qualitative properties of the wall following law for differential drive robots.

Theorem 3 With a selection of ϵ as in (11), the unicycle wall following law in (17), (18) with $\mathbf{x}^* = \mathbf{x}_p(\mathbf{x})$ as in (15) leaves the robot's free space \mathcal{F} positively invariant under its unique continuously differentiable flow, aligns the robot with a $\mathbf{t}_w(\mathbf{x})$ (according to the selection of a in (9)) in finite time and steers the robot along the boundary of a unique obstacle in \mathcal{O} in a clockwise or counterclockwise fashion (depending on a) with a nonzero rate of progress σ afterwards, while maintaining a distance of at most $(r + \epsilon)$ from it.

Proof sketch. Included in [20]. ■

We summarize the proposed method for switching between a path following and a wall following phase and generating velocity commands for a differential drive robot following a reference path \mathcal{P} in [20, Algorithm 1], with the definition of an auxiliary symbolic action $\text{NAVIGATEROBOT}(\mathcal{P}, r, \epsilon, \delta)$.

V. REACTIVE PLANNING FOR GRIPPING CONTACT

In this Section, we describe a method for generating suitable motion commands online for two objects in contact,

of which one is a differential drive robot and uses a gripper to push the other, passive object on the plane. Our method consists of generating “virtual” commands for different points of interest in the robot-object pair and translating them to “actual” commands for the robot using simple kinematic maps.

A. Gripping contact kinematics

Consider the robot gripping an object i , as shown in Fig. 2. We can find the position of the object center of mass $\mathbf{x}_i \in \mathcal{W}$ from the position of the robot center of mass \mathbf{x} as

$$\mathbf{x}_i := \mathbf{x} + (\rho_i + r) \mathbf{e}_{\parallel} \quad (21)$$

where $\phi_i = \text{atan2}(\mathbf{x}_i - \mathbf{x})$ and $\mathbf{e}_{\parallel} = (\cos \phi_i, \sin \phi_i) \in \mathbb{R}^2$ is the unit vector along the line connecting the two bodies. Since, the velocity of the object center of mass will be $\dot{\mathbf{x}}_i = \dot{\mathbf{x}} + (\rho_i + r) \dot{\phi}_i \mathbf{e}_{\perp}$ with $\mathbf{e}_{\perp} = (-\sin \phi_i, \cos \phi_i) \perp \mathbf{e}_{\parallel}$, and since the robot has a grip on the object along its line of motion, so that the orientation of the robot ψ is always equal to the robot-object bearing angle ϕ_i , we can use (1) to write

$$\dot{\mathbf{x}}_i = \mathbf{T}_i \mathbf{u}_{ku} \quad (22)$$

with the Jacobian \mathbf{T}_i given by

$$\mathbf{T}_i = \begin{bmatrix} \cos \psi & -(\rho_i + r) \sin \psi \\ \sin \psi & (\rho_i + r) \cos \psi \end{bmatrix} \quad (23)$$

and $\mathbf{u}_{ku} = (v, \omega)$ the input vector as defined above.

Similarly, consider the circumscribed circle enclosing the robot and the object with radius $(\rho_i + r)$, as shown in Fig. 2. Its *center point* is located at

$$\mathbf{x}_{i,c} = \mathbf{x} + \rho_i \mathbf{e}_{\parallel} \quad (24)$$

Following a similar procedure as above, we can show that

$$\dot{\mathbf{x}}_{i,c} = \mathbf{T}_{i,c} \mathbf{u}_{ku} \quad (25)$$

with the Jacobian $\mathbf{T}_{i,c}$ given by

$$\mathbf{T}_{i,c} = \begin{bmatrix} \cos \psi & -\rho_i \sin \psi \\ \sin \psi & \rho_i \cos \psi \end{bmatrix} \quad (26)$$

B. Generating virtual commands

For the planning process, the fact that both \mathbf{T}_i and $\mathbf{T}_{i,c}$ are always non-singular implies that we can describe the robot-object pair as either a dynamical system of the form

$$\dot{\mathbf{x}}_i = \mathbf{u}_i(\mathbf{x}_i) \quad (27)$$

or a dynamical system of the form

$$\dot{\mathbf{x}}_{i,c} = \mathbf{u}_{i,c}(\mathbf{x}_{i,c}) \quad (28)$$

since we can always prescribe (virtual) arbitrary velocity commands \mathbf{u}_i or $\mathbf{u}_{i,c}$ for either the object itself or for the center point and then translate them to (actual) inputs \mathbf{u}_{ku} through (22) or (25) respectively ($\mathbf{u}_{ku} = \mathbf{T}_i^{-1} \mathbf{u}_i$ or $\mathbf{u}_{ku} = \mathbf{T}_{i,c}^{-1} \mathbf{u}_{i,c}$).

Since the circumscribed circle centered at $\mathbf{x}_{i,c}$ is the smallest circle enclosing both the robot and the object and since Assumption 1 guarantees only that $\eta > 2(r + \max_k \rho_k)$,

we conclude that it is beneficial to consider the dynamical system (28) (and generate virtual commands for the center point $\mathbf{x}_{i,c}$) when following the path \mathcal{P} that the high-level planner provides. However, this will eventually position $\mathbf{x}_{i,c}$ to \mathbf{p}_i^* , instead of the object \mathbf{x}_i (which is desired). Therefore, once the center point is placed to \mathbf{p}_i^* , we have to switch to the system (27) and generate virtual commands for the object \mathbf{x}_i to carefully position it to \mathbf{p}_i^* . Assumption 2 guarantees that this is always possible. We can think of generating commands \mathbf{u}_i and $\mathbf{u}_{i,c}$ as a trade-off between careful object positioning and agility in avoiding obstacles respectively.

C. LIDAR Range transformation

As described above, the robot-object pair is treated as a single holonomic agent with radius $\rho_i + r$ centered at $\mathbf{x}_{i,c}$ when following the reference path \mathcal{P} . However, we know that the LIDAR is positioned on the robot and its range measurements are given with respect to \mathbf{x} . Therefore, we need a mechanism for translating these measurements from \mathbf{x} to $\mathbf{x}_{i,c}$. To this end, we describe the observed points from the LIDAR using the function $\mathbf{x}_{\text{LIDAR}} : (-\pi, \pi] \rightarrow \mathcal{W}$

$$\mathbf{x}_{\text{LIDAR}}(\theta) = \mathbf{x} + \rho_{\mathbf{x}}(\theta) (\cos \theta, \sin \theta) \quad (29)$$

and find the equivalent ranges from $\mathbf{x}_{i,c}$ as

$$\rho_{\mathbf{x}_{i,c}}(\theta) = \min\{R - \rho_i, \|\mathbf{x}_{\text{LIDAR}}(\theta) - \mathbf{x}_{i,c}\|\} \quad (30)$$

since $R - \rho_i$ is the minimum distance that can be observed from $\mathbf{x}_{i,c}$ when no obstacles are present and corresponds to the ray along the orientation ψ of the robot towards the object.

We summarize the proposed algorithm for switching between a path following and a wall following phase and generating velocity commands for a robot-object pair following a reference path \mathcal{P} in [20, Algorithm 2], with the definition of an auxiliary symbolic action $\text{NAVIGATEROBOTOBJECT}(\mathcal{P}, r, \rho_i, \epsilon, \delta)$.

VI. LOW-LEVEL IMPLEMENTATION OF SYMBOLIC LANGUAGE

In this section, we describe the low-level implementation and realization of the three symbolic actions introduced in Section II, used to solve our assembly problem.

A. Action MOVETOBJECT

The low-level implementation of this symbolic action is quite straightforward, since the robot just needs to follow the plan provided by the high-level planner and navigate to a specific object using the auxiliary action NAVIGATEROBOT [20, Algorithm 1]. The only caveat is that the robot needs to be aligned with the object it needs to pick up in order to engage the gripper. Since, no continuous law can guarantee both position and orientation convergence for a nonholonomically-constrained, differential drive robot [22] and a discontinuous law needs to be introduced, we compute

$$\tilde{\alpha} := \min\{\alpha \in [0, 1] \mid \mathcal{P}(\alpha) \in B(\mathbf{p}_i, \rho_i + r)\} \quad (31)$$

and “truncate” the path to $\mathcal{P}([0, \tilde{\alpha}])$. In this way, the robot will navigate to $\mathcal{P}(\tilde{\alpha})$ (within a δ tolerance) which satisfies $\|\mathcal{P}(\tilde{\alpha}) - \mathbf{p}_i\| = \rho_i + r$ as desired. Then, in order to align the robot with the object, the linear command v is set to zero and the angular command is set to

$$\omega = -k(\phi_i - \psi) \quad (32)$$

until $\phi_i = \psi$. The low-level implementation is shown in [20, Algorithm 3].

B. Action POSITIONOBJECT

From the preceding analysis in Section V and using the auxiliary action NAVIGATEROBOTOBJECT, we can construct the POSITIONOBJECT algorithm as shown in [20, Algorithm 4]. Since the task of NAVIGATEROBOTOBJECT is to bring the object close enough to the destination in order to allow careful positioning (allowed by Assumption 2), a final positioning step is required. To this end, instead of generating virtual commands for the center of the circumscribed circle of the robot-object pair as shown in (28), we generate commands for the center of the object itself, as shown in (27), according to the following law

$$\mathbf{u}_i = -k(\mathbf{x}_i - \mathbf{p}_i^*) \quad (33)$$

These virtual commands are then translated to actual robot commands according to (22).

C. Action MOVE

This action is exactly like MOVETOBJECT, but there is no final orientation requirement. Its low-level implementation is shown in [20, Algorithm 5].

Note here that the formal results accompanying both the path following phase [18] and the wall following phase (Theorems 1 and 3) along with Theorem 2 guarantee that every symbolic action command will be successfully executed.

VII. NUMERICAL EXAMPLES

In this Section, we provide numerical examples⁶ of assembly processes in various environments using the symbolic action commands described above.

A. Environment packed circular obstacles

First, we test our algorithm in a rectangular, 20x20m workspace, packed with circular obstacles, whose position and size are unknown to the deliberative planner. The minimum separation η between the obstacles is chosen to be only slightly above (e.g 5cm) the minimum allowed value prescribed by Assumption 1, in order to demonstrate the validity of our approach, deriving from the formal guarantees of Theorem 1. The goal is to place an object to a desired position, shown in 5. The deliberative planner outputs a plan comprising of two actions: MOVETOBJECT(1, \mathcal{P}_1) \rightarrow POSITIONOBJECT(1, \mathcal{P}_2), whose sequential execution and the corresponding reference paths $\mathcal{P}_1, \mathcal{P}_2$ are depicted in Fig. 5.




Fig. 5: A depiction of a packed two stage assembly process with a fixed timestep, with the separation value just above the minimum allowed value.

B. Cluttered environment with walls

In this paragraph, we demonstrate the execution of a more challenging task. The robot should position the two obstacles depicted in Fig. 6 to their predefined positions within a polygonal workspace with walls, whose locations are provided a-priori to the deliberative planner, and then return to a “nest” location. The workspace is packed with several convex, not-necessarily circular obstacles. As shown in Fig. 6, the deliberative planner outputs a high-level plan comprising of five actions: MOVETOBJECT(1, \mathcal{P}_1) \rightarrow POSITIONOBJECT(1, \mathcal{P}_2) \rightarrow MOVETOBJECT(2, \mathcal{P}_3) \rightarrow POSITIONOBJECT(2, \mathcal{P}_4) \rightarrow MOVE(\mathcal{P}_5), which is successfully executed by the reactive planner. An example for an object trajectory during this execution is shown in Fig. 1. Notice that, in contrast with several reactive wall following schemes that require an estimate of the wall curvature, our scheme can easily handle obstacles with corners. It is also worth noting that the deliberative planner hit the maximum number of expansions allowed and had difficulties extracting a feasible plan when it was provided the exact position and size of every obstacle, due to the highly packed construction. This highlights another benefit of our approach; we can significantly reduce the computational load of high-level planners by tasking them only with the extraction of the action sequence required, and using the reactive planner for local obstacle avoidance and convergence online. This happens because the computational load of the reactive planner remains the same regardless of the number of obstacles.

Finally, it is worth noting that the proposed scheme is capable of executing a sequence of symbolic commands provided by the deliberative planner, even when Assumptions 1 or 2 or the obstacle convexity are not satisfied. In the accompanying video, we provide examples of successful assemblies even in the absence of both obstacle convexity and enough separation.

VIII. CONCLUSION AND FUTURE WORK

This paper presents a provably correct architecture for planning and executing a successful solution to an instance of the Warehouseman’s problem by decomposition into an offline “deliberative” planning module and an online “reac-

⁶All simulations were run in MATLAB using ode45 and a gain $k = 2$.

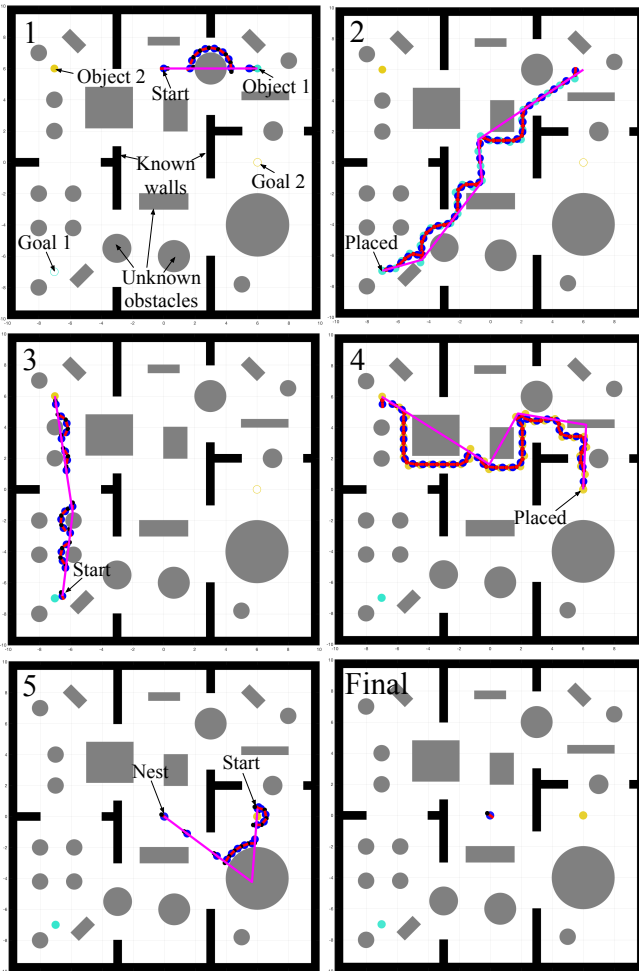


Fig. 6: An illustration of the assembly process described in Section VII-B, with a fixed timestep. The walls and boundaries of the workspace, known to the deliberative planner, are shown in black and the unexpected obstacles handled by the reactive planner are shown in grey.

“itive” execution module. A differential drive robot equipped with a gripper and a LIDAR sensor, capable of perceiving its environment only locally, is used to successfully position the passive objects in a desired configuration with no collisions along the way. Formal proofs and numerical results demonstrate the validity of our method.

Future work will address the problem of closely integrating the reactive and deliberative planners and shifting more of the planning burden to the reactive component, allowing the deliberative planner to consider more complex dynamics and generate paths faster and more efficiently. Moreover, we plan to experimentally verify the reactive planning ideas for pushing a grasped object, as presented in Section V.

ACKNOWLEDGMENT

This work was supported in part by the ARL/GDRS RCTA project, Coop. Agreement #W911NF-1020016 and in part by AFRL grant FA865015D1845 (subcontract 6697371).

REFERENCES

- [1] J. E. Hopcroft, J. T. Schwartz, and M. Sharir, “On the complexity of motion planning for multiple independent objects; pspace-hardness of the” warehousemans problem”, *The International Journal of Robotics Research*, vol. 3, no. 4, p. 7688, 1984.
- [2] J. Wolfe, B. Marthi, and S. Russell, “Combined task and motion planning for mobile manipulation.” in *Proceedings of the International Conference on Automated Planning and Scheduling*, 2010.
- [3] D. Berenson, S. Srinivasa, D. Ferguson, and J. Kuffner, “Manipulation planning on constraint manifolds,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009.
- [4] G. Konidaris, L. Kaelbling, and T. Lozano-Perez, “Constructing symbolic representations for high-level planning,” in *proc. AAAI*, 2014.
- [5] L. P. Kaelbling and T. Lozano-Perez, “Hierarchical task and motion planning in the now,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2011.
- [6] W. Vega-Brown and N. Roy, “Asymptotically optimal planning under piecewise-analytic constraints,” in *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, San Francisco, CA, 2016. [Online]. Available: <https://groups.csail.mit.edu/rrg/papers/vegabrown2016asymptotically.pdf>
- [7] M. Vendittelli, J.-P. Laumond, and B. Mishra, *Decidability of Robot Manipulation Planning: Three Disks in the Plane*, ser. Springer Tracts in Advanced Robotics. Springer, Cham, 2015, p. 641657. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-16595-0_37
- [8] A. Deshpande, L. P. Kaelbling, and T. Lozano-Perez, *Decidability of semi-holonomic prehensile task and motion planning*, 2016. [Online]. Available: <http://lis.csail.mit.edu/pubs/deshpande-WAFR16.pdf>
- [9] R. Chatila, “Deliberation and reactivity in autonomous mobile robots,” *Robotics and Autonomous Systems*, vol. 16, no. 2, p. 197211, Dec 1995.
- [10] D. E. Koditschek, “An approach to autonomous robot assembly,” *Robotica (Cambridge. Print)*, vol. 12, p. 137155, 1994.
- [11] H. Isil Bozma and D. E. Koditschek, “Assembly as a noncooperative game of its pieces: analysis of 1d sphere assemblies,” *Robotica*, vol. 19, no. 01, p. 93108, 2001.
- [12] C. S. Karagoz, H. I. Bozma, and D. E. Koditschek, “Feedback-based event-driven parts moving,” *Robotics, IEEE Transactions on [see also Robotics and Automation, IEEE Transactions on]*, vol. 20, no. 6, p. 10121018, 2004.
- [13] B. Marthi, S. Russell, and J. Wolfe, “Angelic hierarchical planning: Optimal and online algorithms,” in *ICAPS*, 2008.
- [14] W. Vega-Brown and N. Roy, “Optimal and near-optimal task and motion planning with abstraction,” 2017. [Online]. Available: <https://groups.csail.mit.edu/rrg/papers/vegabrown2017optimal.pdf>
- [15] D. Koditschek and E. Rimon, “Robot navigation functions on manifolds with boundary,” *Advances in Applied Mathematics*, vol. 11, no. 4, pp. 412–442, 1990.
- [16] O. Arslan and D. E. Koditschek, “Exact robot navigation using power diagrams,” in *Robotics and Automation, 2016 IEEE International Conference on*, 2016, pp. 1–8.
- [17] —, “Sensor-based reactive navigation in unknown convex sphere worlds,” in *The 12th International Workshop on the Algorithmic Foundations of Robotics*, 2016.
- [18] —, “Smooth extensions of feedback motion planners via reference governors,” in *Robotics and Automation, 2017 IEEE International Conference on*, 2017(accepted).
- [19] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. Kavraki, and S. Thrun, *Principles of robot motion: theory, algorithms, and implementations*. MIT Press, Cambridge, MA, 2005.
- [20] V. Vasilopoulos, W. Vega-Brown, O. Arslan, N. Roy, and D. E. Koditschek, “Technical Report: Sensor-Based Reactive Symbolic Planning in Partially Known Environments,” *Technical Report*, September 2017. [Online]. Available: <https://arxiv.org/abs/1709.05474>
- [21] R. Burridge, A. Rizzi, and D. Koditschek, “Sequential composition of dynamically dexterous robot behaviors,” *The International Journal of Robotics Research*, vol. 18, pp. 534–555, 1999.
- [22] R. W. Brockett, “Asymptotic stability and feedback stabilization,” in *Differential Geometric Control Theory*. Birkhauser, 1983, pp. 181–191.