The Task Motion Kit

Neil T. Dantam[*†]   Swarat Chaudhuri[*]   Lydia E. Kavraki[*]

I. Introduction

Robots require novel reasoning systems to achieve complex objectives in new environments. Everyday activities in the physical world couple discrete and continuous reasoning. For example, to set the table in Fig. 1, the robot must make discrete decisions about which objects to pick and the order in which to do so, and execute these decisions by computing continuous motions to reach objects or desired locations. Robotics has traditionally treated these issues in isolation. Reasoning about discrete events is referred to as task planning while reasoning about and computing continuous motions is the realm of motion planning. However, several recent works have shown that separating task planning from motion planning—that is finding first a series of actions that will later be executed through continuous motion—is problematic; for example, the next discrete action may specify picking an object, but there may be no continuous motion for the robot to bring its hand to a configuration that can actually grasp the object to pick it up. Instead, Task–Motion Planning (TMP) tightly couples task planning and motion planning, producing a sequence of steps that can actually be executed by a real robot to bring the world from an initial to a final state. This article provides an introduction to TMP and discusses the implementation and use of an open-source TMP framework that is adaptable to new robots, scenarios, and algorithms.

TMP presents challenges both in algorithmic design and software engineering. Interaction between the discrete, task component and the continuous, motion component imposes requirements not faced by stand-alone task planners or motion planners. The planner may need to consider alternate task plans in an efficient way until finding one that can actually be executed by the robot at hand, whereas typical task planners generate only a single plan. In addition, actions where the robot grasps and rearranges objects will change the kinematics and configuration space in which the robot can move, whereas typical motion planners assume a fixed configuration space. Thus, we cannot expect to combine existing tools for isolated task planning and motion planning and produce frameworks that can consistently use high-level specifications of behavior to produce motion. Instead, we must handle the possible interactions of discrete and continuous components to identify task plans and executable motions.

The Task–Motion Kit (TMKit)[1] is an end-to-end system for probabilistically complete task–motion planning and real-time execution. TMKit follows the high-level design shown in Fig. 2 to implement the algorithm of [1] and at the same time provides a general framework to integrate multiple methods for task planning, motion planning, and task–motion interaction. Shared abstractions and data structures are fundamental aspects of TMKit that enable coupling of task planning, motion planning, and real-time estimation and control. TMKit is modular, extensible, and we are adapting it to additional methods for TMP [2], [3]. Whenever appropriate, we employ widely-used formats and protocols to promote compatibility. The resulting system generates real-time, collision-free robot motion from high-level specifications. To our knowledge, this is the first publicly-available, general-purpose TMP framework. Sharing this project with the community will encourage the implementation of more TMP approaches and provide a valuable tool for the development and comparison of related

Figure 1: An example Task–Motion Planning problem: setting a table. The input for TMP includes **(a)** the start state, **(b)** the goal state, and a set of allowable actions (e.g., pick, place, etc.). TMP finds the output **(c)** which consists of a sequence of discrete actions (the task plan) and their corresponding continuous paths (or motion plans).

techniques.

II. Background

A. Task Planning

Task planning identifies a sequence of discrete actions that change an initial state into a desired goal state or condition, given a *task domain* that defines the available actions and their preconditions and effects (see Sec. IV-A). This field evolved largely from pioneering work on the Stanford Research Institute Planning System (STRIPS) [4]. The leading approaches for efficient task planning are heuristic search [5] and constraint satisfaction [6].

Off-the-shelf task planners typically focus on efficiently finding a single plan. In contrast, TMP often requires search through multiple, alternate task plans as discussed in the introduction. This raises an inherent challenge: motion planners used to compute paths are at best probabilistically complete for high-dimensional systems. Consequently, we cannot generally prove the non-existence of corresponding motion plans. To address this challenge, our system does not use an off-the-shelf task planner but rather employs a newly-introduced task planner capable of efficiently generating alternate plans.

B. Motion Planning

Motion planning identifies a continuous path of valid configurations—i.e., joint positions—from an initial state to a desired goal state (see Sec. IV-B). Sampling-based motion planners are widely used for high-dimensional systems [7]. Such sampling-based planners offer *probabilistic completeness*, guaranteeing that the planner will eventually find a solution if one exists. However, if solution does not exist, a sampling-based planner cannot prove this negative; in such case, the planner would not terminate or would run until a timeout. Motion planners based on gradient descent or optimization are also common and highly efficient, but they do not offer the same probabilistic completeness guarantees as the sampling-based motion planners. This work uses sampling-based planners that offer probabilistic completeness guarantees because probabilistic completeness of the overall framework is a desired property. Conveniently, high-quality, open source implementations of such planners are available [8]. Future integration of alternate motion planning approaches is possible, with their accompanying set of trade-offs, but the integration of motion planners in TMP needs special attention to address the coupling of task planning and motion planning.

Off-the-shelf motion planning frameworks often abstract the details of robot kinematics or assume the kinematic equations are fixed or change infrequently, with only configurations changing during planning [8]. In contrast, TMP requires rapid updates to kinematic equations; as the robot grasps and transfers objects, these objects' poses change between fixed values and functions of robot configuration. Moreover, these changes may involve more than just the individual grasped object, such as in the case of moving a tray or pushing a cart containing other objects. Consequently, kinematic representations capable of efficient updates are required for TMP (see Sec. IV-B).

C. Task–Motion Planning

Task–Motion Planning (TMP) takes an initial state to a desired goal state through the concurrent or interleaved production of high-level, discrete action sequences via task planning and continuous, collision-free paths via motion planning.

Most prior work on TMP focuses on computational performance rather than completeness or generality, which is emphasized in this work. [9] applies geometric constraints to limit the motion planning space or prove motion infeasibility in special cases. Hierarchical Planning in the Now (HPN) [10] interleaves planning and execution, reducing search depth but requiring reversible actions—e.g., rearranging objects but not pouring a cup down a drain—when backtracking. [11] extends a hierarchical task planner with geometric primitives, using shared literals that relate task-level symbols with motion-level geometric entities. [12] interfaces an off-the-shelf task planner and motion planning using a heuristic to remove objects that potentially block the robot's path. [13] formulates the *motion* side of TMP as a constraint satisfaction problem over a discretized, preprocessed subset of the configuration space. The Robosynth framework [14] uses a Satisfiability Modulo Theories (SMT) solver to generate task and motion plans from a static roadmap, employing plan outlines to guide the planning process. FFRob [15] develops an FF-like [5] task-layer heuristic based on a lazily-expanded roadmap. Overall, these methods set aside the broad challenge of ensuring probabilistic completeness that arises from interactions between the task and motion layers. In contrast, the framework we implement focuses on probabilistically complete TMP.

A smaller number of task and motion planners do achieve probabilistic completeness. The aSyMov planner [16] combines a heuristic-search task planner with lazily-expanded roadmaps. Our implementation of [1] in TMKit operates differently at the task level, motion level, and interface level, yielding different performance characteristics than aSyMov. For example, aSyMov's composed roadmaps could be amortized over multiple runs but composing roadmaps for object interactions may be expensive. In contrast, [1] finds a new motion plan each run, but efficiently updates scene data structures to handle object interaction (see Sec. IV-B). Furthermore, TMKit is extensible to both forward-search [5] and constraint-based [6] task planners.

While source code is available for some specific methods such as [12], we believe that TMKit is the first publicly available framework that is extensible to multiple methods and domains. A key to this extensibility is our abstraction of the interaction between task and motion layers via the *domain semantics* (see Sec. IV-C and Fig. 6) that enable the introduction of new actions and domains without any necessary changes to the framework itself.

D. Plan Execution

Motion planners make certain assumptions to achieve sufficient performance, and the execution step must correct those assumptions in real-time. Specifically, motion planners typically assume (1) a given model for the kinematics and geometry of the robot and environment and (2) that motion between nearby joint configurations is possible. In reality, geometric

models contain numerous errors due to imprecise lengths, encoder calibration error, flexing of assumedly rigid bodies, inaccurate object detection, inaccurate camera calibration, etc. Thus, despite the *precision* or repeatability of many robots, *accurate* motion to correct poses still presents challenges. In addition, robot motion is subject to dynamic constraints on velocity, force, current, etc. The execution step must track the planned path in a way that is physically feasible, and it must correct for the inevitable and sundry errors.

III. Overview of Task–Motion Planning

We first discuss the high-level approach to task–motion planning and execution shown in Fig. 2.

A. Input to TMP

The input to TMP includes the discrete Task Domain, the continuous Motion Domain, and the coupling of these two sides.

1) Task Domain: The task domain defines the discrete actions the robot can take, including their preconditions and effects. For example, the action `pick-up` may have a precondition that the object is on the table and the effect that the object is in the robot's hand. Sec. IV-A describes our implementation of task domains, and Fig. 6 provides a complete example of the `pick-up` action with preconditions and effects.

2) Motion Domain: The motion domain defines the 3-dimensional (3D) positions objects in the environment, the kinematic structure of robot joints, and the geometry—i.e., meshes—of objects and robot links. Collectively, we call the robot and environment the *scene* and the tree or graph of the local coordinate frames of environment objects and robot links the *scene graph*. A scene graph defines the *configuration space* of a robot. For a given configuration, computing the forward kinematics of each frame in the scene graph provides the mesh positions; then, specialized collision checkers [17] determine whether those positions are in collision.

We specify both an initial scene—consisting of the robot and environment—and a goal scene for the planner. Then, we map from these scenes to task states using the domain semantics.

3) Domain Semantics: The domain semantics defines the coupling between the discrete task domain and the continuous motion domain. Two types of functions are necessary. First, we need a function to map from a scene graph to a discrete state for the task planner. For example, if the scene graph defines some object `a`'s position relative to the robot's hand, the domain semantics would set a discrete variable `holding-a`, indicating that the robot is holding object `a`, to true. Second, we need a function to map from a discrete task action to a motion planning problem (start and goal states) for the motion planner. For example, the `pick-up` action would start at the robot's current configuration and move to a goal that is a grasping configuration for the object to be picked up. Sec. IV-C discusses our implementation of the domain semantics.

B. Task–Motion Planner

The Task–Motion Planner finds the sequence of discrete task actions from the task domain and their corresponding motion plans, based on the domain semantics, that will take the system from some initial state or scene to a desired state or scene. This planning process is structured as an alternation between task planning to identify the discrete actions and motion planning to identify the paths for each action. Some task plans may include infeasible actions, e.g., picking up an object that is blocked by something in front of it. In this case, motion planning would fail—that is, exceed a timeout—and we would go back to the task planner to find a different task plan, e.g., first moving the blocking object out of the way. Sec. IV-C discusses our implementation of a Task–Motion Planner.

C. Task–Motion Control

The Task–Motion Control phase executes the plan in real-time. Each path produced by the motion planner is a sequence of waypoints that the robot must move through. To execute this motion plan, we compute a reference position, velocity, etc. for the robot at each timestep by interpolating between the waypoints. In addition, we must correct positioning error in following the motion plan through feedback control. Finally, we operate the gripper to grasp and release objects as specified by the actions in the plan. Sec. IV-D discusses our control and execution implementation.

IV. TMKit Implementation

We now discuss our task–motion system TMKit. This section targets the researcher interested in TMKit for the algorithm of [1] or for implementing new TMP approaches. Fig. 3 outlines the major software components in our system implementation. TMP involves many different software modules, and our design choices were also influenced by the need to support real-time execution. The key to integrate these components in our system was identifying the appropriate abstractions for task and motion domains and relating these abstractions through the *domain semantics*. Using these suitable abstractions not only eases development but also increases flexibility by providing a uniform interface to domain information such as task state or scene geometry.

A. Task Domain

We represent the task domain by the Task Language (a) in Fig. 3. Generally, task domains are specified using a variety of notations and logics, but at a fundamental level, all these representations define some type of transition system, automaton, or formal language. The de facto standard syntax for task planning is the Planning Domain Definition Language (PDDL) [18], which our framework also takes as input. PDDL (see Fig. 6) defines parameterized actions with preconditions and effects based on first-order logic. Our task planning algorithm [1], however, is not specific to PDDL and assumes only that the state space is finite and compactly represented with a set of variables. Thus, new task domains can be created in PDDL, and the underlying algorithm is adaptable to other notions as well.
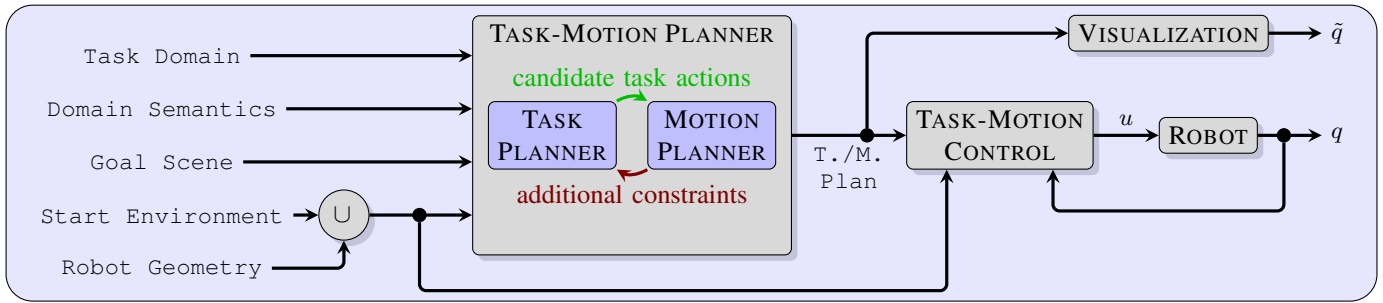
Figure 2: High-Level Planning and Execution Block Diagram. The inputs are the task domain definition (see Sec. IV-A Fig. 6), the environment and robot geometries, combined to produce the scene graph (see Sec. IV-B and Fig. 4), and the domain semantics that relate the task and motion layers (see Sec. IV-C and Fig. 6). The Task–Motion Planner generates a plan based on these inputs. The Task–Motion Control layer executes the plan, sharing a geometric representation—the scene graph—with the planning layer. The control output u drives the robot, resulting in configuration q. In a parallel layer, we visualize the system at simulated configuration $\tilde{q}$. Sec. III provides an overview of the system, and Sec. IV discusses its implementation.
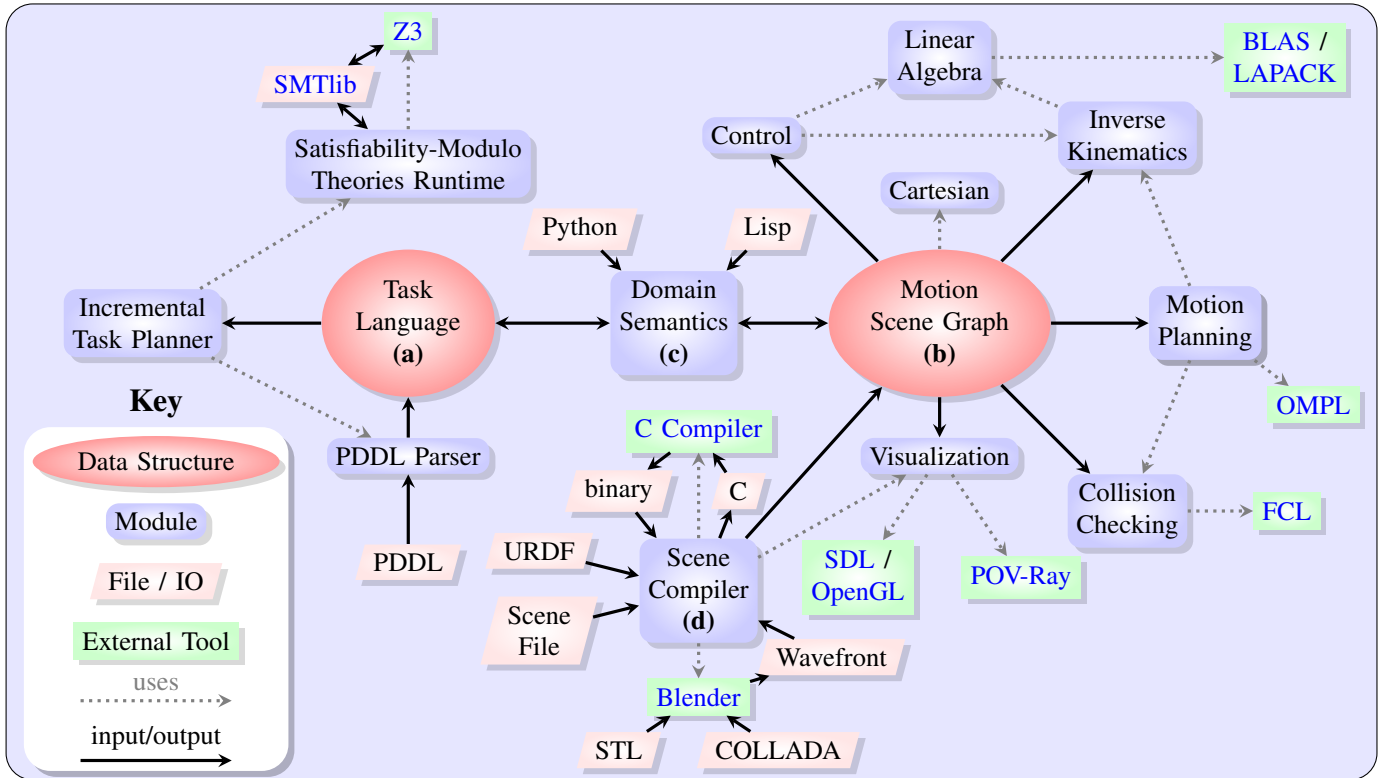

Figure 3: Map of software components. The key data structures are **(a)** the task language (see Sec. IV-A) and **(b)** the scene graph (see Sec. IV-B). These data structures are connected by **(c)** the domain semantics definitions (see Sec. IV-C). **(d)** the scene compiler is also an important component (see Fig. 5). This system integrates the following external tools and formats: *BLAS/LAPACK* (Basic Linear Algebra Subprograms / Linear Algebra PACKage): High-performance linear algebra routines with many vendor-optimized implementations. *COLLADA* (COLLAborative Design Activity): An interchange file format for 3D applications. *FCL* (Flexible Collision Library): A popular software library for collision checking. *GCC* (GNU Compiler Collection): A compiler suite from the GNU project. *OMPL* (Open Motion Planning Library): A popular software library for sampling-based motion planning. *POV-Ray* (Persistence Of Vision RAYtracer): An open source ray-tracing program. *PDDL* (Planning Domain Definition Language): Cross-platform library to access graphics, audio, mouse, keyboard, etc. *SDL* (Simple DirectMedia Layer): Cross-platform library to access graphics, audio, mouse, keyboard, etc. *SMT* (Satisfiability Modulo Theories): A decision problem combining logic and additional *theories*, e.g. integer constraints, lists, arrays. *STL* (STereoLithography): A file format for CAD software. *URDF* (Universal Robot Definition Format): XML file type for robot kinematics. *XML* (eXtensible Markup Language): A tree-structured, general-purpose file format. *Z3*: A high-performance theorem prover/SMT solver.

B. Motion Domain

The motion domain is represented by the scene graph (b) in Fig. 3. Motion planning algorithms are typically defined in terms of abstract *configuration spaces* [8], while robot manipulators are modeled as kinematic trees or scene graphs of joints and links in packages such as OpenRave[2], Orocos KDL[3], and MoveIt![4]. Existing implementations, however, focus on only a subset of the TMP pipeline shown in Fig. 2. Consequently, TMKit, uses a new, streamlined scene graph representation that enables direct task–motion translation, efficient updates, and real-time kinematics.

The scene graph is a tree representing relative Cartesian poses, with data attached at each node for geometry (e.g., meshes), inertial parameters, joint limits, etc. Fig. 4 shows how the scene graph edges correspond to symbolic multiplication or chaining of transformations in the Cartesian space. Starting from the global root "0" (see Fig. 4) and multiplying the relative pose of each local coordinate frame along a chain yields the global pose of the frame at the end of the chain. Picking and placing objects, common operations in TMP, are represented by *reparenting* a frame in the scene graph, i.e., changing the object's parent between the hand and a support surface such as a table.

We now discuss the details of implementing the scene graph data structure, which provides essential infrastructure for TMP systems. Our scene graph implementation offers a unique set of features that make it suitable for both TMP and real-time execution. We use two variations on the structure: a *mutable* version and a *persistent*, i.e., purely functional version. Both variations can be efficiently updated at runtime, e.g., when the robot picks up a tray of objects, and they share underlying data for geometric objects via reference counting so that data for large meshes is not copied. The mutable version is based on indexed arrays and avoids heap allocations—which may impose unacceptable pauses—after construction, making it suitable for real-time operation. The persistent version is based on weight-balanced binary trees that efficiently create partial copies on updates, useful during planning when we backtrack to a previous point in the search and a previous version of the scene graph. To enable efficient multi-threaded access, e.g., when performing inverse kinematics, motion planning, and visualization in separate threads, we separate the scene graph object from the data for states and configurations.

We also provide a compiler (see Fig. 5) enabling scene graphs to be specified in domain specific languages. Our compiler supports the widely-used ROS Universal Robot Definition Format (URDF). Additionally, due to the difficulty of writing URDF by hand, we also introduce a new, compact scene file syntax (see Fig. 6). Compiling scenes offers several performance and administrative advantages:

1) Compiled scenes are fast to load because the operating system directly maps into memory (via `mmap`) the included mesh data, eliminating runtime parsing and processing.

[2]http://openrave.org/
[3]http://www.orocos.org/kdl
[4]http://moveit.ros.org

2) Compiled scenes reduce memory use compared to runtime parsing when multiple operating system processes operate on the same scene, because the memory mapped scene graphs in different processes share physical memory.

3) Compiled scenes are easy to distribute to other machines, e.g., a cluster, which may lack scene sources, utilities, or support libraries. Only the executable or shared library is required to load the compiled scene, reducing potential runtime dependencies.

4) Compiled scenes avoid the need to include large parsing libraries, e.g., an XML parser, in real-time processes and reduce the dynamic allocations necessary to load scene.

5) Multiple compiled scenes can be flexibly composed both statically ahead-of-time and dynamically at run-time, improving overall scene construction efficiency when some portions of the scene are fixed and others changing. For example, we can pre-compile the scene for the robot, which remains constant based on the robot's mechanical design. However, the locations of objects on a table may change frequently, so we can compose the pre-compiled robot scene with the separate or dynamically-generated scene for these objects.

The scene graph data structure and scene compiler provide the necessary geometric support for TMP and plan execution.

C. Task–Motion Planning

Our TMP implementation follows the overall structure of Fig. 2, based on the algorithm of [1]. In the task layer, we use an incremental, constraint-based task planner. In the motion layer, we include a variety of sampling-based motion planners through the Open Motion Planning Library (OMPL) [8].

The key to achieving generality in our planner is the selection of abstractions. Our task languages can model arbitrary finite state task domains, and our scene graphs can model arbitrary rigid body robots and environments.

We relate the task and motion domains by defining a *Domain Semantics* (see Fig. 3.c). The domain semantics defines the conversion of the scene graph to a task state and defines functions to refine high-level task actions by computing the corresponding motion plans. Concretely, the domain semantics in TMKit are functions written in Python or Common Lisp. Fig. 4 includes an example of a task state computed from a scene graph, and Fig. 6 contains an example of a refinement function in the domain semantics for pick-and-place manipulation. The same semantics definition may be used across multiple scenes or problem instances. Changes to the task domain, e.g., new actions, do require updating the domain semantics but require no changes to the planner itself. By abstracting task-motion interaction to these separate domain semantics definitions, our planning system generalizes across domains.

D. Output and Execution

The immediate output of our system is a Task–Motion plan describing the sequence of task actions and corresponding motion plans. We benchmark the performance and scalability of the overall approach in [1]. Fig. 6 shows a fragment of such

Figure 4: The scene graph abstraction for a simplified version of the Baxter robot and the corresponding task state: **(a)** kinematic equation to compute the right wrist pose. **(b)** kinematic equation to compute the left wrist pose. **(c)** the local coordinate frames of the corresponding scene graph overlaid on the Baxter. **(d)** the task state corresponding to the scene graph for a pick-and-place task domain, partially shown in Fig. 6. This state abstracts the object relationships in the scene graph. Note that the task state is computed automatically via the *Domain Semantics* and that additional or different task predicates may be used by modifying the Domain Semantics (see Sec. IV-C). The transform $^a\mathcal{S}_b$ is the relative Cartesian pose between parent a and child b. The coordinate frame labels for Baxter consist of the left (l) or right (r) arm; the shoulder (s), elbow (e) or wrist (w); and the zeroth (0), first (1), and second (2) joint. The other frame labels represent the global root (0), table (T), or blocks (A, B, C).
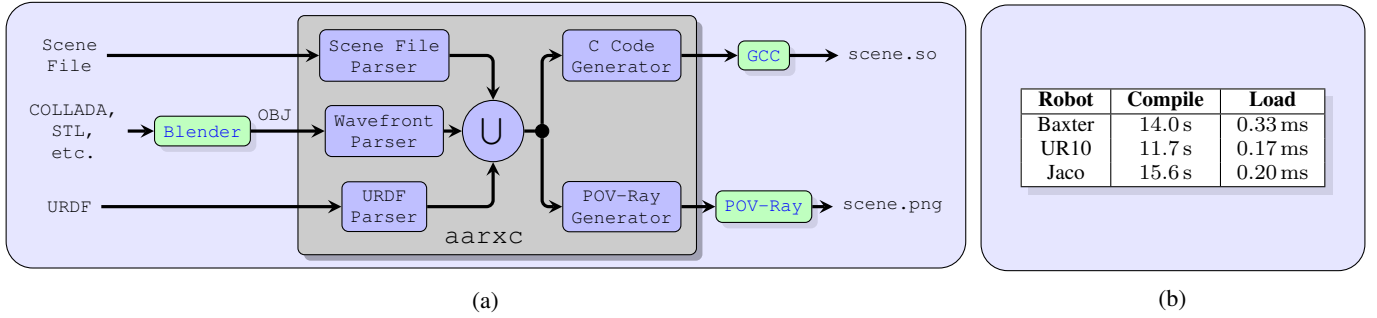

Figure 5: The scene graph compiler `aarxc`. **(a)** Compiler block diagram. The compiler includes parsers for scene files, Wavefront OBJ meshes, and ROS URDF files. It uses the Blender 3D modeling program to convert a variety of meshes to the conventional Wavefront OBJ format. The compiler translates the loaded scene graph to optimized C code for later fast loading and real-time execution. It can also translate scene graphs to input for the POV-Ray raytracer for high-quality visualization. **(b)** Compile times—including mesh processing, code generation, and C compilation—and load times for common robots using Blender 2.77 and GCC 4.9.2 on an Intel® Core™ i7-4790. Example plans and planning times are presented in Fig. 6 and Fig. 8. [1].

a plan for the table-setting example, represented using a plain-text, line-based file format, which is both human-readable and can be efficiently parsed. Each line indicates either a task action, the joints moved during a motion plan, a waypoint in a motion plan, or a reparent operation, e.g., picking or placing an object by changing an edge in the scene graph from the table to hand or vice versa. The resulting file defines the interleaving of task actions and motion plans.

Finally, we execute the task–motion plan by interpolating the given motion plans and performing the indicated reparentings to grasp and release objects. There are numerous methods to interpolate the waypoint sequence of a motion plan so as to satisfy the physical limits of a robot, e.g., maximum acceleration and velocity. Given any such interpolation, we use a feedback control law to compute the command for robot:

$$\dot{q}_u(t) = \dot{q}_r(t) - k\left(q_a(t) - q_r(t)\right) \qquad (1)$$

where $\dot{q}_u$ is the velocity command, $\dot{q}_r$ is the reference velocity from the interpolated waypoints, k is a feedback gain, q_a is the actual position, and q_r is the interpolated reference position.

Finally we must communicate with the robot hardware at each time step to retrieve the actual state q_a and $\dot{q}_a$ and to

send the velocity command $\dot{q}_u$. For example, we would use Controller Area Network (CAN) bus message for the Schunk LWA4, TCP for the Universal Robot, or ROS communication for the Rethink Baxter.

V. Example Use Case

Fig. 6 illustrates an example use case of TMKit for a domain such as the table setting in Fig. 1, including the planner's specific input and output.

1) Task Domain: The task domain block shows the `pick-up` action, with its preconditions and effects. This action picks up an object from a table, so the precondition requires the object to be on the table and uncovered. The effect is that robot holds the object and the object is not on the table. The full task domain includes similar definitions for other actions to put-down objects, unstack objects, etc.

2) Motion Domain: The motion domain block shows the definition for a single object—a glass. The definition includes the object's relative position to its parent (the shelf) and the object's geometric mesh. The full task domain includes similar definitions for the other glasses and bowls as well as the links and joints of the robot.

3) Domain Semantics : The domain semantics block shows the function to find a motion plan for the `pick-up` action. This function computes the current position of the object, then it attempts to find a motion plan to bring the robot's hand to a grasping pose for that object. If motion planning fails (exceeds a timeout), the motion planning function generates an exception that the task–motion planner will catch and handle by finding a different task plan based on the feedback from the motion planner.

4) Task–Motion Planner : The task–motion planner block illustrates the alternation and feedback between task planning and motion planning. The task planner identifies a high-level plan. The motion planner attempts to find corresponding paths. Failing to do so, the motion planner provides additional constraints to the task planner which then finds a different task plan. This process iterates until finding a task plan where all actions have corresponding motion plans.

5) Task–Motion Plan: The task–motion plan block shows the first two actions of the plan: picking and placing an object. The first `pick-up` action includes the joint waypoints to move the robot's hand to the grasping position for an object, then changes the object's parent in the scene graph to the robot's hand. The second `put-down` action includes the waypoints to move the robot's hand and the grasped object to the desired location, then (unshown) will change the object's parent in the scene graph to the table, placing the object. The full task-motion plan contains the rest of the actions necessary to achieve the desired goal.

6) Plan Execution: Fig. 7 and Fig. 8 show two example task–motion plans and planning times, one for the Rethink Robotics Baxter and one for the Universal Robots UR5. The same overall framework produces the plan for each system.

We apply the framework in each case by using the URDF model of the robot for the specific system.

These examples demonstrate the modularity and extensibility of TMKit. TMKit works on multiple robots, it supports multiple types of actions—e.g., picking, placing, stacking, pushing—and it handles coupling between objects—i.e., moving cans in a bin in Fig. 8. Additional benchmark results are presented in [1].

VI. Conclusion

We have presented a new software framework for Task–Motion Planning (TMP) and execution, the Task–Motion Kit (TMKit). TMKit is available under an open source, permissive license[5]. We believe that TMKit is the first open source TMP framework that is extensible to multiple domains, different planning methods, and supports end-to-end planning and execution. Its modular design enables TMKit to generalize across hardware platforms, task domains, and TMP algorithms.

Our objective in this work was to produce a general-purpose, easy-to-use, and extensible framework for TMP. There are numerous avenues to improve and build upon this framework. Going beyond our implemented method of [1], we will extend the feedback between the task and motion layers, improve plan re-use, and incorporate additional rich constraint capabilities. Already, we are adapting and integrating the additional TMP methods of [2], [3] with TMKit. We hope the community will find this end-to-end system both easy to use and a helpful platform to demonstrate other methods for TMP.

Currently, TMKit focuses on the geometric case of motion planning, which is often sufficient for manipulation. Planning with dynamics, e.g., considering torques in the planning layer, may be necessary for other cases such as bipedal walking. However, including dynamics during motion planning may impact completeness [19], so careful analysis is necessary. Improved considerations for planning with dynamics remains an area of future work for TMKit.

An ongoing need in TMP is support to compare and benchmark different TMP algorithms and implementations. We believe that TMKit, as an extensible framework supporting common formats such as PDDL and URDF, can help meet that need. Furthermore, modular components such as the scene graph compiler (see Fig. 5) could aid the development of alternate TMP methods and implementations. We hope that TMKit will be a useful tool for other researchers to evaluate existing algorithms and extend to new approaches.

Acknowledgments

References

[1] N. T. Dantam, Z. Kingston, S. Chaudhuri, and L. E. Kavraki, "Incremental task and motion planning: A constraint-based approach," in *Robotics:*
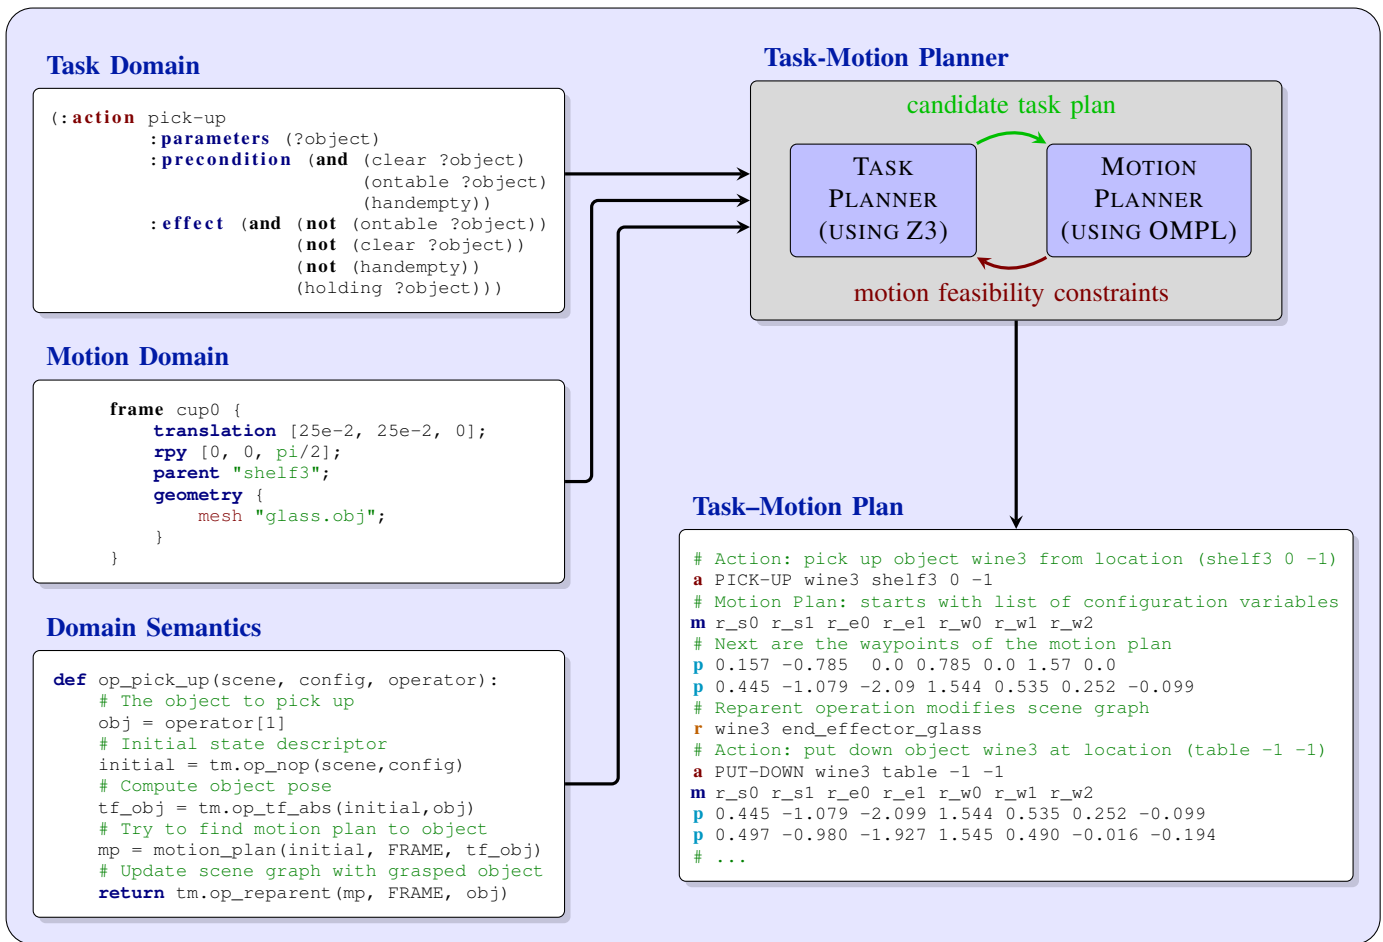
[5]http://tmkit.kavrakilab.org

Figure 6: Task–Motion Planner Implementation Diagram, showing fragments of the planner's input—i.e., the task domain, domain semantics, and motion domain—and output—i.e., the task–motion plan. Sec. V discusses this example.

Science and Systems, 2016.

[2] K. He, M. Lahijanian, L. E. Kavraki, and M. Y. Vardi, "Towards manipulation planning with temporal logic specifications," in *Intl. Conf. on Robotics and Automation*. Seattle, WA: IEEE, 2015, pp. 346–352.

[3] Y. Wang, N. T. Dantam, S. Chaudhuri, and L. E. Kavraki, "Task and motion policy synthesis as liveness games," in *International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI, 2016.

[4] R. E. Fikes and N. J. Nilsson, "STRIPS: A new approach to the application of theorem proving to problem solving," *Artificial intelligence*, vol. 2, no. 3, pp. 189–208, 1972.

[5] J. Hoffmann and B. Nebel, "The FF planning system: Fast plan generation through heuristic search," *J. of Artificial Intelligence Research*, pp. 253–302, 2001.

[6] H. Kautz and B. Selman, "Unifying SAT-based and graph-based planning," in *Intl. Joint. Conf. on Artifial Intelligence*, vol. 99, 1999, pp. 318–325.

[7] L. E. Kavraki and S. M. LaValle, "Springer handbook of robotics," B. Siciliano and O. Khatib, Eds. Springer, 2016, ch. 5, pp. 109–128.

[8] I. Şucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.

[9] F. Lagriffoul and B. Andres, "Combining task and motion planning: A culprit detection problem," *Intl. J. of Robotics Research*, vol. 35, pp. 890–927, July 2016.

[10] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *Intl. J. of Robotics Research*, vol. 32, no. 9-10, pp. 1194–1227, August/September 2013.

[11] M. Gharbi, R. Lallement, and R. Alami, "Combining symbolic and geometric planning to synthesize human-aware plans: toward more efficient combined search." in *Intl. Conf. on Intelligent Robots and Systems*. IEEE, 2015, pp. 6360–6365.

[12] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel, "Combined task and motion planning through an extensible planner-independent interface layer," in *Intl. Conf. on Robotics and Automation*. IEEE, 2014, pp. 639–646.

[13] T. Lozano-Pérez and L. P. Kaelbling, "A constraint-based method for solving sequential manipulation planning problems," in *Intl. Conf. on Intelligent Robots and Systems*. IEEE, 2014, pp. 3684–3691.

[14] S. Nedunuri, S. Prabhu, M. Moll, S. Chaudhuri, and L. E. Kavraki, "SMT-based synthesis of integrated task and motion plans from plan outlines," in *Intl. Conf. on Robotics and Automation*. IEEE, 2014, pp. 655–662.

[15] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "FFRob: An efficient heuristic for task and motion planning," in *Algorithmic Foundations of Robotics XI*. Springer, 2015, pp. 179–195.

[16] S. Cambon, R. Alami, and F. Gravot, "A hybrid approach to intricate motion, manipulation and task planning," *Intl. J. of Robotics Research*, vol. 28, no. 1, pp. 104–126, 2009.

[17] J. Pan, S. Chitta, and D. Manocha, "FCL: A general purpose library for collision and proximity queries," in *Intl. Conf. on Robotics and Automation*. IEEE, 2012, pp. 3859–3866.

[18] M. Ghallab, D. McDermott, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins, "PDDL—the planning domain definition language," Yale Center for Computational Vision and Control, Tech. Rep. CVC TR-98-003/DCS TR-1165, October 1998.

[19] T. Kunz and M. Stilman, "Kinodynamic RRTs with fixed time step and best-input extension are not probabilistically complete," in *Algorithmic Foundations of Robotics XI*. Springer, 2015, pp. 233–244.

(a) (b) (c)

(d) (e) (f)

Figure 7: Example task–motion plan to set a table using the Rethink Robotics Baxter. Average planning time for 10 trials was $64.8\,$s on an Intel® Core™ i7-4790. **(a)** the initial state. **(b)** picking the first glass. **(c)** placing the first glass. **(d)** placing the second glass. **(e)** placing the first bowl. **(f)** placing the second bowl.

(a) (b) (c) (d) (e) (f)
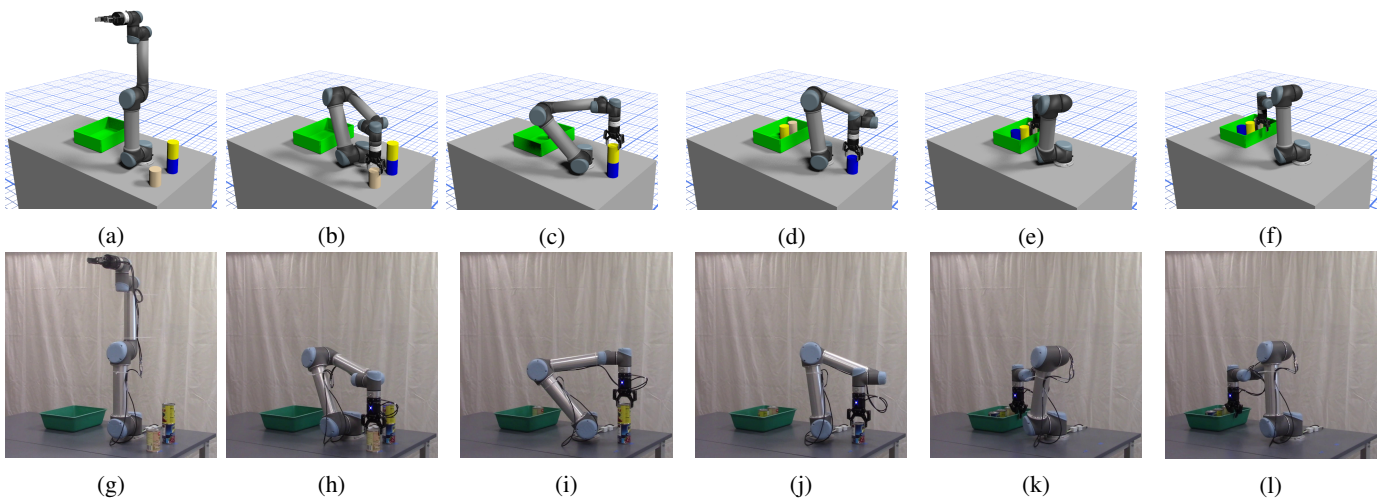
(g) (h) (i) (j) (k) (l)

Figure 8: Example task–motion plan to load and move a bin using the Universal Robots UR5. The same overall framework with a different URDF for the robot produces this plan for a different system. Average planning time for 10 trials was $8.78\,$s on an Intel® Core™ i7-4790. **(a)-(i)** simulated execution. **(j)-(r)** physical execution.