

Integrating Reactive Behavior and Planning: Optimizing Execution Time through Predictive Preparation of State Machine Tasks

Andreas Schierl, Alwin Hoffmann, Ludwig Nägele and Wolfgang Reif
Institute for Software & Systems Engineering
University of Augsburg, Germany
Email: {schierl,hoffmann,naegele,reif}@isse.de

Abstract—Tasks that change the physical state of a robot take a considerable amount of time to execute. However, many robot applications spend the execution time waiting, although the following tasks might require time to prepare. This paper proposes to amend tasks with a description of their expected outcomes, which allows planning successive tasks based on this information. The suggested approach allows sequential and parallel composition of tasks, as well as reactive behavior modeled as state machines. The paper describes the means of modeling and executing these tasks, details different possibilities of planning in state machine tasks, and evaluates the benefits achievable using the approach.

Keywords-planning; reactive behavior; state machine; robot programming; optimization

I. INTRODUCTION

Most classical industrial robot arms are currently used for fixed and recurring preprogrammed tasks. There, planning (at run time) and reaction to indeterministic events are of minor importance. When it comes to more modern or mobile robots, the situation is different: Due to the greater variability in the environment, these robots must be able to sense external events and react accordingly, and can thus be seen as reactive systems.

In this context, it can be observed that when executing a task, performing its physical actions often requires a considerable amount of time compared to the computations needed to prepare them. Still, it is often possible to tell details about the positive outcome of the task even before the task is fully executed. In contrast, for failures the amount of information that can be given in advance is limited (e.g. where the robot will be when the error occurs is usually unknown). While these details could considerably help to plan ahead, many software frameworks often ignore this potential and spend the execution time waiting for completion (cf. Sect. II).

To improve this situation, we propose to model tasks including their expected outcome(s) and details about the corresponding situation(s). These expected outcomes include successful execution, but also detectable errors and relevant situations that occur during execution. Using these tasks, it becomes possible to plan the following tasks even before the execution of the current task has completed. These tasks can be combined in a sequential or parallel way, can perform



Figure 1: Robots handing over a baton, adapted from [2]

case distinctions and can even be used to model reactive behavior through state machines. Composing these tasks, the programmer can configure details about the relationship between the tasks and the quality of transition: For some events, an immediate reaction may be required, so the corresponding reaction task has to be planned before the event can occur, while other events allow some time, so that more relaxed planning schemes can be used.

To evaluate the approach, the cooperation between two mobile robots was modeled and implemented (cf. Fig. 1). There, two KUKA youBots [1] drive in parallel, and hand over a baton while in motion. To realize the reactive behavior and coordination, state machines are used, while the possibility of advanced planning reduces the waiting time between the different robot actions. Additionally, a theoretical example with more complex planning is introduced and analyzed, showing the possible advantages of the approach.

This paper is structured as follows: After an overview of different ways to handle complex robot behavior (Sect. II), the approach of defining and executing tasks is outlined in Sect. III. Sect. IV describes the experiments conducted to evaluate the approach and points out the corresponding results. Finally, Sect. V gives a conclusion and outlook.

II. RELATED WORK

For modeling and composing tasks or device capabilities, different approaches exist in *ROS* and *OROCOS*.

ROS *actionlib* [3] allows to model interruptible tasks through *Actions*. *Actions* represent long-running, preemptible tasks that provide feedback and notify about their result, and that can be canceled when the goal has to be changed. They can be seen as an extensible way to model device capabilities, allowing extensions by introducing new components that provide *Action* servers. To combine multiple *Actions*, a new component can provide an *Action* that invokes the corresponding *Actions* in parallel or sequentially. However, no timing guarantees are given for this type of composition, as it relies on network communication between the different *Action* servers. Furthermore, *Actions* do not share a common interface and thus cannot directly be passed between different components (e.g. a planner cannot provide its result as a generic *Action* that can be passed to another component deciding when it should be executed).

For reactive behavior on the task-level with *ROS*, Bohren et al. [4] introduced *SMACH* state machines. In *SMACH*, *ROS* *Actions*, *Services* or Python code can be defined as *States*, and *States* can be composed into new *States* in a concurrent, sequential or state-machine form. Therefore, each *State* defines different outcomes that can be used in composite *States* to define reactions or handle error conditions. Furthermore, states in *SMACH* (and thus state machines) can be made available to other components as *actionlib* *Actions*. This composition mechanism is similar to the one presented in this paper, however without timing guarantees, so that guaranteed reactions to events cannot be specified this way. Furthermore, *SMACH* (as well as *actionlib*) provides no explicit semantics or metadata for *Action* or *State* post conditions, making further planning during execution harder to achieve. *SMACH* state machines are executed in a blocking way and thus do not offer the possibility of using the execution time of robot actions for further (motion) planning steps.

For *OROCOS*, restricted Finite State Machines (*rFSM*) [5] allow modeling reactive behavior. An *rFSM* describes a hierarchical state machine without parallelism, aimed at the coordination of robot applications. According to the *pure-coordination* pattern [5], these state machines only process and raise Boolean events which have to be provided by monitor components or handled by configurator components, which in turn manipulate or reconfigure the active components in order to achieve the goal. *rFSM* state machines are implemented on top of the programming language *LUA* with specialized memory allocation and garbage collection, and can thus be executed with real-time guarantees.

As an extension to *rFSM* state machines, Scioni et al. [6] describe how to achieve *preview coordination*: In this approach, the execution environment takes hints about ex-

ecution probabilities based on likelihood labels on some transitions. Using these labels, likely successor states can be prepared (performing some of their work) while the previous state is still active, as long as there is no conflict between the preparation steps and the actions performed in the current state. This allows reducing the execution time, while keeping the action definitions coherent (instead of moving the preparation step into the previous state).

The preview coordination mechanism introduces a form of decoupling between workflow and capability execution and makes use of meta data about conflicts between states, but the *rFSM* mechanism still does not include further semantic descriptions for the results of states. This way, the following tasks cannot analyze and prepare for the expected results of the previous *State*, a powerful and important feature offered by the approach introduced in this paper.

Another related approach has been introduced by Angerer [7] in the form of Robotics API *Activities*. It can be seen as the basis of the approach suggested in this paper, but had a stronger focus on real-time and specific ways of modeling meta data, and does not support reactive behavior in the form of state machines. This approach has been further extended by Schierl [2] and along with further research led to the results presented in this paper.

III. APPROACH

To facilitate planning during execution, we propose to model individual robot tasks along with a description of their expected outcome(s) (as described in Sect. III-A). Based on these outcome description(s), it becomes possible to plan the successor task (or at least start planning) before the current task has been fully executed – maybe even before execution starts. This feature can be exploited during execution, as described in Sect. III-B. When it comes to more complex tasks with different viable outcomes, an implementation that is based on simple control flow tends to become confusing, so it becomes helpful to define the desired behavior in a model-based way. Here, the formalism of state machines is often used. Our approach allows modeling complex tasks based on state machines, as described in Sect. III-C. For these state machines, different choices concerning execution exist that may be configured and affect the resulting performance (Sect. III-D).

A. Defining Tasks and Outcomes

Initially, a task to be executed by a specific actuator (e.g. move to a given position, pick up an item, or even bring me a beer) is modeled as an *Activity* (cf. Fig. 2). As the same task can be performed more than once (if you are with friends), it becomes helpful to keep track of different task executions. So, when a task is to be executed in a given situation, its *Activity* is instantiated through *createHandle()* yielding an *ActivityHandle*.




Figure 2: Relevant classes for task definition and execution

Being responsible for a single task execution, the *ActivityHandle* offers state tracking that notifies about execution progress, and is also responsible to decide how the desired task is to be executed in a given situation. It receives a situation description (see below) the task is to be executed in, and has to provide instructions on how to act, and what outcome to expect after the action has been performed. Thus, the *ActivitySchedule* takes a situation description called *ActivityResult*, and provides an *ActivitySchedule* that describes the task to execute, as well as a set of possible *ActivityResult* that can occur while/after executing the task. This way, it provides possible post conditions for the execution of the task, while the preconditions are derived from the results of the previous task. If the previous task has multiple possible outcomes, the *ActivitySchedule* is provided with the different acceptable outcomes, and can plan for each of the envisioned situations.

This differs from typical planning mechanisms that are based on pre/post conditions and try to combine fitting actions, but instead allows applying a given task in many different situations.

To describe the expected situation, the *ActivityResult* contains status information about each individual actuator contributing to the task, as well as a description about the envisioned geometric and logical situation of the world. The actuator-specific information may contain details about a gripper, e.g. whether it is open or closed, as well as joint positions, velocities and accelerations for robot arms.

To talk about the world state, a model as described by Schierl et al. [8] is used. It is based on geometric features

(modeled as frames), along with relations that are defined between the frames that form an undirected multigraph. These relations describe the logical situation, e.g. that an object is placed on the ground or grasped by a gripper (or, rather their respective frames are), as well as geometric information to give the exact position where the frames are relative to each others. To give a description about the future world state, a view of the world model with both position changes and topology changes is required. For tasks that include robot motions, a position change that gives the (expected) new place of the robot or all of its components is recorded. This geometric information is used whenever performing geometric calculations based on the world view. Additionally, some tasks change the connectivity between different frames: When a robot picks up an object from the ground, the relation between ground and object has to be removed, and a new relation between gripper and object has to be established. These topology changes are also relevant for further planning, because a grasped object changes the shape of a robot: some motions that were previously possible now result in collisions. Thus, the world view provided by the task includes sets of removed and added relations. After executing the task, the described world state has to be in place, either implicitly because the robot moved to the promised position, or explicitly by enacting the topology changes.

B. Executing and Composing Tasks

Once a task has been defined, i. e. created as an *Activity*, it can be executed in a blocking or non-blocking way. For a specific execution instance, the *Activity* creates an *ActivityHandle* that takes the *ActivityResult* of the previous *Activity* and prepares *ActivitySchedules* for each of the situations. The *ActivitySchedules* are activated so that they trigger the execution of the respective task once their *ActivityResult* is reached. Additionally, the *ActivitySchedules* provide their possible *ActivityResult* for use with following *Activities*.

In the blocking case, the method *execute()* blocks until a completion *ActivityResult* of a corresponding *ActivitySchedule* has been reached. Fig. 3a shows a depiction of this behavior. The *Activities* are shown as vertical life lines with time running from top to bottom, with the different boxes denoting the phases planning *P* and execution *E*. However, in this case the situation description contained in the *ActivityResult* is of little use, because once the result has been reached, the described situation is already reached, and the execution time has already been wasted waiting.

To use the execution time, non-blocking execution with *beginExecute()* allows the control flow to continue once any of the *ActivitySchedules* has been triggered and the execution of the task has thus started. The following task can then be prepared for different possible start situations (based on the *ActivityResult* of the currently running task), and immediately execute one of the planned solutions if the




Figure 3: Blocking and non-blocking execution of simple and composed tasks with planning (P) and execution (E)

corresponding situation occurs. This situation is shown in Fig. 3b – here, the dotted lines denote times where *B* is already planned and waiting to be executed. The preparation of *B* happens while *A* is running, and *B* starts running once *A* is completed. In contrast, if *B* is independent of *A*, e.g. if it controls different devices, it starts immediately when prepared, as shown in Fig. 3c.

Additionally, these tasks can easily be combined: For sequential execution, the sequential task provides the second task with all the *ActivityResults* of the first task's *ActivitySchedules*, so that the second task can react to all expected outcomes of the first task, while the results of the second task are provided as results of the sequence. In cases where determinism is required, this kind of planning can even allow real-time guarantees (given a corresponding software framework that supports scheduling of real-time tasks): planning the second task for all possible outcomes of the first task before the first task's execution starts guarantees that the second task can immediately take over no matter how long the first task takes. In Fig. 3d, the sequence *S* is executed after *A*, and the planning of *S1* and *S2* completely happens before *S1* is started, guaranteeing that *S2* can be started immediately when *S1* ends. However, this determinism is only possible if the outcomes are fully known in advance: if details of an outcome description depend on sensor data, these *ActivityResults* can only be provided when the sensor data becomes available, and planning of the next task is delayed to this moment (no longer guaranteeing determinism).

To achieve parallel execution, the parallel task provides both subtasks with the same initial situation (as *ActivityResults*), and combines resulting *ActivitySchedules* (and optionally *ActivityResults*) to create the *ActivitySchedule* and *ActivityResults* of the parallel task. This situation is shown

in Fig. 3e, where the parallel tasks *P1* and *P2* are planned and thus allow *P* to start them at the same time after *A* is completed. Of course, the parallel task has to make sure that the subtasks do not conflict, e.g. by checking that they control different devices.

Additionally, simple case distinctions can be made based on the possible *ActivityResults*. A task with case distinction classifies the *ActivityResults* of the previous task and forwards them to the corresponding subtask depending on the decision condition. Fig. 3f shows this situation, with *C* as a case distinction for the result of *A*, deciding to execute *C2* and unload *C1*.

C. Using State Machines for Complex Tasks

For longer or repeating task sequences, these composition mechanisms that fully plan ahead may reach their limit, because the number of tasks that has to be prepared grows exponentially with the number of successive tasks (given that each has more than one possible result). Thus, it becomes helpful to limit the amount of pre-planning. One way to achieve this while supporting complex compositions of tasks is the use of state machines – a step often taken when specifying reactive behavior.

The proposed approach uses *Activities* to define states of the state machine. For each state (or *Activity*), transitions can be given that specify a switch to another *Activity* if certain *ActivityResults* occur. Additionally, one distinct *Activity* is chosen as start state, and *ActivityResults* of some *Activities* can be defined as transitions to a final state.

As each state can be entered more than once during execution, for its *Activity* multiple *ActivityHandles*, *ActivitySchedules* and also *ActivityResults* are prepared. The resulting set of *ActivityResults* for an *Activity* is dynamic, so the transition cannot give a complete list. Instead, it works




Figure 4: State machine definition and execution with planning (P) and execution (E)

as a predicate to *ActivityResult*s that chooses whether a given *ActivityResult* qualifies for the transition, which can be as generic as *any error* or as specific as *any successful execution of the given grasp*.

As an example, Fig. 4a shows a state machine of the three tasks *A*, *B*, and *C*, switching from *A* to *C* if *A* succeeds, and to *B* if it fails. State *B* is always followed by *C*, leading to two different paths reaching *C*.

In addition to the predicate and following state, the transition can also be annotated with qualifiers that direct the order or timing of preparation. For transitions that may not be missed (such as error recovery strategies that bring the robot into a safe state), a transition can be marked as reliable, requesting it to be prepared before the originating *Activity* starts. These reliable transitions are handled similar to sequential tasks described before, making sure that the following state can be executed no matter how long or short the first state takes. While having the same advantages of sequential tasks, reliable transitions also inherit their drawbacks, especially the amount of preparation necessary if more than one *ActivityResult* is handled by a reliable transition, as well as the fact that cycles in the reliable transition graph are forbidden.

Fig. 4c amends the transition from *A* to *B* with a stereotype «reliable», defining that the execution environment has to guarantee that the transition will be taken if the failure occurs.

D. Executing State Machines

Regarding execution, state machine tasks differ from the composed tasks mentioned previously: While for sequential, parallel or conditional tasks planning can completely be performed up front, state machine tasks require some planning while the task is running. The *Activity* for the initial state is prepared traditionally, based on the *ActivityResult*s of the previous task. Then, the *ActivityResult*s of the created

ActivitySchedule are compared against the transitions originating from the start state to decide whether they have to be handled by switching to another state (i.e. *Activity*). However, for found transitions and *Activities*, different options exist when to prepare them for the corresponding results, as a trade-off between preparation of unnecessary traces and not having prepared a required task. For the scope of this paper, the preparation is usually delayed until their originating state is entered (i.e. its *Activity* started running). Then, the target states for outgoing transitions are prepared using the reported *ActivityResult*s of the current state, so that they can react to the occurrence of the corresponding result. Additionally, the outgoing transitions of the target state are analyzed, so that they can be handled once the target state is entered. This way, the states are prepared with a look-ahead of one transition, which works fine as long as the typical execution time of a state is longer than the preparation times for all following states.

Fig. 4b shows the preparation timing resulting for the state machine defined in Fig. 4a. *B* and *C* (denoted *C1*) are prepared once *A* is started, and when the failure occurs, *B* is started and *C* is unloaded. A new instance of *C* (denoted *C2*) has to be prepared for the new situation resulting of *B*, which is then executed once *B* finishes. Using this method yields a behavior that is similar to the one of non-blocking execution (cf. Fig. 3b), where the following task is prepared once the previous is started.

However, this mechanism of delaying planning outgoing transitions once the state is entered is not sufficient to guarantee that a certain transition can be taken without further delay, especially if the result happens quickly after entering the state. Thus, for reliable transitions, planning may not be delayed until this time. Instead, reliable transitions and their target state have to be prepared along with their originating state, before the originating state is entered.

This behavior is shown in Fig. 4d. Here, *A* and *B* are




Figure 5: State machine: Robot interaction, adapted from [2]

planned before A is started, making sure that the occurrence of the failure cannot happen while B is still being planned.

IV. EXPERIMENTAL RESULTS

The proposed approach has been evaluated in two examples. First as a real-world example, the interaction between two mobile robots has been modeled. There, two mobile robots drive parallel to each others and hand over a baton while in motion. The implementation uses two KUKA youBots and is based on the Robotics API [9]. There, both robots are controlled using their on-board computers using a C++ control core, while the high-level coordination and task execution is implemented in Java and performed from a laptop computer connected to the youBots through a wireless network.

Because the handover example (apart from error handling) is purely sequential, the first implementation was based on separate tasks for moving the arm and gripper and the execution model described in Sect. III-B. For all tasks that do not depend on the second robot, non-blocking execution was used, while the correct order of the gripping and releasing tasks was performed through blocking execution of the corresponding tasks in a common control flow. While working fine in simulation, this implementation led to unintended delays between the grasp and release operations when executed on real robots, and thus required more space for parallel driving than necessary. This was mainly caused by the unreliable network connection and inefficient network code that required multiple communication roundtrips to transmit and start the tasks.

A second implementation modeled the expected behavior through two linked state machines (cf. Fig. 5). Here, the following tasks could already be prepared while the second youBot was waiting for the first youBot, so the delay between the grasp and release operation as well as the required space for the interaction were significantly reduced.

The second, theoretical scenario was performed without connection to a robot framework, but based on a prototypical implementation with support for reliable transitions. Here,




Figure 6: State machine: Second scenario

the scenario was to drive a mobile robot to a pick-up area, take a large object, and return to its start position. In addition, during all steps a sensor had to be observed, and the robot had to stop once the sensor detected a dangerous situation. To complicate things, in the environment the shortest path to the pick up zone could not be taken while carrying the large object, but instead a detour was required.

Fig. 6 shows a state machine model of the task. The top row of states gives the main success flow, while reliable transitions handle danger occurring during this task. Additionally, a *Back off* state has been added, which is executed after the robot has stopped in case of emergency. As after stopping the robot is already in a safe state, this transition is not time-critical and thus does not need to be reliable.

The motion planning tasks, *Go to pick up zone* and *Return home* were implemented as tasks that take considerable time to plan (3 s for collision free motion planning) and to execute (5 s to move the robot along the planned path), while *Pick up* only takes time to execute (3 s to move the arm and gripper), but plans quickly (1 s). As a comparative reference simulating a system that does not support planning ahead based on result meta data, the three main success task were executed in a blocking way as a sequence.

When executed, the state machine of the second scenario implementation completed after about an average of 16.05 s, while the comparative reference took an average of 20.08 seconds. Fig. 7 shows the resulting life lines for those different execution models, clearly showing where time can be saved by planning during execution.

V. CONCLUSION

Working with robots, the performance of applications is not only limited by the available processing power of the computer, but also by the physical limitations of controlled devices. Usually, the time needed to execute a task is significantly longer than the mere computation time, but still some planning also takes considerable time. In this paper, we proposed to amend robot tasks with descriptions of their expected outcomes. This allows planning the following




Figure 7: Execution times for the second scenario with planning (P) and execution (E)

task while the current task is still running, thus reducing delays between the execution of successive robot tasks and avoiding the situation that execution time goes unused although planning steps may be pending.

These amended tasks can be executed in a blocking or non-blocking way, or combined into complex tasks. Besides sequential and parallel composition, state machines promise to allow the specification of more complex reactive behavior including recurring subtasks, while still behaving like regular tasks, so that further composition or planning remains possible.

A prototypical implementation of this mechanism has been created based on the Robotics API [9] and shown to work for cooperating robots. Additionally, the approach promises to accelerate the execution of tasks where considerable time is spent planning the next steps, such as collision free motion planning, and also allows specifying events for which a timely reaction has to be guaranteed (given a capable execution environment, such as a Realtime Robot Control Core [10] used with the Robotics API).

Still, this paper is limited to simple preparation strategies of state machines (transitions can be guaranteed, or planned when the state is entered). As a part of further research, more complex strategies might benefit from likelihood annotations [6] or estimated planning times to decide which transitions in the state machine will likely happen and should be prepared first (to reduce the risk of missed transitions). Additionally, longer sequences of transitions could be prepared in corresponding situations, e.g. to skip over states that take a very short time to plan and execute. In the given theoretical example, this could help if the grasping task took shorter, because then the execution time of the *Go to* task could be used to plan the *Return* task.

REFERENCES

- [1] R. Bischoff, U. Huggenberger, and E. Prassler, "KUKA youBot - a mobile manipulator for research and education," in *Proc. 2011 IEEE Intl. Conf. on Robot. & Autom., Shanghai, China*. IEEE, May, pp. 1–4.
- [2] A. Schierl, "Object-oriented modeling and coordination of mobile robots," Ph.D. dissertation, Universität Augsburg, 2017.
- [3] E. Fernandez, E. Marder-Eppstein, and V. Pradeep. *actionlib*. Online, accessed Feb 2016. [Online]. Available: <http://wiki.ros.org/actionlib>
- [4] J. Bohren and S. Cousins, "The SMACH high-level executive," *IEEE Robot. & Autom. Mag.*, vol. 17, no. 4, pp. 18–20, 2010.
- [5] M. Klotzbücher and H. Bruyninckx, "Coordinating robotic tasks and systems with rFSM statecharts," *J. of Software Engineering for Robotics*, vol. 3, no. 1, pp. 28–56, 2012. [Online]. Available: <http://joser.unibg.it/index.php?journal=joser&page=article&op=view&path%5B%5D=52>
- [6] E. Scioni, M. Klotzbuecher, T. De Laet, H. Bruyninckx, and M. Bonfe, "Preview coordination: An enhanced execution model for online scheduling of mobile manipulation tasks," in *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2013)*, Nov 2013, pp. 5779–5786.
- [7] A. Angerer, "Object-oriented Software for Industrial Robots," Ph.D. dissertation, University of Augsburg, Mar 2014.
- [8] A. Schierl, A. Angerer, A. Hoffmann, and W. Reif, "Consistent world models for cooperating robots: Separating logical relationships, sensor interpretation and estimation," in *2017 First IEEE International Conference on Robotic Computing (IRC)*, April 2017, pp. 101–108.
- [9] A. Angerer, A. Hoffmann, A. Schierl, M. Vistein, and W. Reif, "Robotics API: Object-Oriented Software Development for Industrial Robots," *J. of Software Engineering for Robotics*, vol. 4, no. 1, pp. 1–22, 2013. [Online]. Available: <http://joser.unibg.it/index.php?journal=joser&page=article&op=view&path=53>
- [10] M. Vistein, A. Angerer, A. Hoffmann, A. Schierl, and W. Reif, "Flexible and continuous execution of real-time critical robotic tasks," *International Journal of Mechatronics and Automation*, vol. 4, no. 1, pp. 27–38, 1 2014.