



(51) International Patent Classification:
B25J 9/16 (2006.01)

(21) International Application Number:
PCT/US2020/047810

(22) International Filing Date:
25 August 2020 (25.08.2020)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/894,178 30 August 2019 (30.08.2019) US

(71) Applicant: X DEVELOPMENT LLC [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) Inventors: DUPUIS, Jean-Francois; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). GO, Keegan; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(74) Agent: DIETRICH, Allison W. et al.; Fish & Richardson P.C., PO Box 1022, Minneapolis, Minnesota 55440 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(54) Title: ROBOT PLANNING FROM PROCESS DEFINITION GRAPH




FIG. 7B

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on computer storage media, for performing robot planning using a process definition graph. The techniques can include receiving a process definition graph having a plurality of task nodes that represent respective tasks to be performed by a respective robot of a plurality of robots, wherein each task node is associated with a location at which the task will be performed; generating, from the process definition graph, an initial modified process definition graph that adds constraints for respective swept volumes occupied by each task represented by the plurality of task nodes; and generating, from the initial modified process definition graph, a refined process definition graph, wherein the refined process definition graph includes respective motion plans for robots moving between tasks, wherein the motion plans define transitions that avoid the swept volumes occupied by each task represented by the plurality of task nodes.



(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*
-

ROBOT PLANNING FROM PROCESS DEFINITION GRAPH

CROSS-REFERENCE TO RELATED APPLICATION

This application claims the benefit of US Application No. 62/894,178, filed August 30, 2019, which is incorporated by reference in its entirety.

5

BACKGROUND

This specification relates to robotics, and more particularly to planning robotic movements.

Robotics planning refers to scheduling the physical movements of robots in order to perform tasks. For example, an industrial robot that builds cars can be programmed to first pick up a car part and then weld the car part onto the frame of the car. Each of these actions can themselves include dozens or hundreds of individual movements by robot motors and actuators.

Robotics planning has traditionally required immense amounts of manual programming in order to meticulously dictate how the robotic components should move in order to accomplish a particular task. Manual programming is tedious, time-consuming, and error prone. In addition, a schedule that is manually generated for one workcell can generally not be used for other workcells. In this specification, a workcell is the physical environment in which a robot will operate. Workcells have particular physical properties, e.g., physical dimensions, that impose constraints on how robots can move within the workcell. Thus, a manually programmed schedule for one workcell may be incompatible with a workcell having different robots, a different number of robots, or different physical dimensions.

In many industrial robotics applications, the primary success or failure criteria of a schedule is the time it takes to complete a task. For example, at a welding station in a car assembly line, the time it takes for robots to complete welds on each car is a critical aspect of overall throughput of the factory. When using manual programming, it is often difficult or impossible to predict how long the resulting schedule will take to complete the task.

SUMMARY

This specification describes how a system can generate a schedule for one or more robots using an underconstrained process definition graph. In general the system can generate initial transitions by planning robot motion plans around task volumes of a

workcell. The system can add constraints for volumes occupied by tasks in order to generate a process definition graph that includes motion plans that avoid the volumes occupied by each task.

In general, one innovative aspect of the subject matter described in this specification can be embodied in methods that include receiving a process definition graph, the process definition graph having a plurality of task nodes that represent respective tasks to be performed by a respective robot of a plurality of robots, each task node being associated with a location at which the task will be performed; generating, from the process definition graph, an initial modified process definition graph that adds constraints for respective swept volumes occupied by each task represented by the plurality of task nodes; and generating, from the initial modified process definition graph, a refined process definition graph, the refined process definition graph including respective motion plans for robots moving between tasks, the motion plans defining transitions that avoid the swept volumes occupied by each task represented by the plurality of task nodes. Other embodiments of this aspect include corresponding computer systems, apparatus, computer program products, and computer programs recorded on one or more computer storage devices, each configured to perform the actions of the methods. A system of one or more computers can be configured to perform particular operations or actions by virtue of having software, firmware, hardware, or a combination of them installed on the system that in operation causes or cause the system to perform the actions. One or more computer programs can be configured to perform particular operations or actions by virtue of including instructions that, when executed by data processing apparatus, cause the apparatus to perform the actions.

This and other embodiments can each optionally include one or more of the following features. In some embodiments, generating the refined process definition graph includes generating motion plans for a plurality of combinations of the plurality of robots performing the plurality of tasks.

In some embodiments, generating the motion plans for the plurality of combinations of the plurality of robots includes selecting an initial ordering of the plurality of tasks; and iteratively generating respective motion plans for a respective robot to perform each subsequent task in the initial ordering until no additional motion plans can be generated.

In some embodiments, the method includes selecting a different ordering of the plurality of tasks; and iteratively generating respective motion plans for each subsequent

robot in the different ordering to perform a respective task until no additional motion plans can be generated.

In some embodiments, generating, from the initial modified process definition graph, a refined process definition graph occurs before any tasks have been assigned to any of the plurality of robots.

In some embodiments, generating, from the initial modified process definition graph, a refined process definition graph occurs before any ordering of execution has been assigned to the plurality of tasks.

In some embodiments, the method includes generating, from the refined process definition graph, a schedule for the plurality of robots that specifies executing motion actions that avoid the swept volumes occupied by tasks represented by the plurality of task nodes in the graph.

Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. Using transformers to iteratively manipulate an underconstrained process definition graph to generate a robotics schedule dramatically reduces the amount of manual programming required in order to program robots. The system can generate a fully constrained schedule for an arbitrary number of robots from an initial underconstrained process definition graph. Using a process definition graph makes specifying robot actions more flexible, faster, and less error-prone. In addition, the techniques described below allow for exploring a much larger search space than what could be achieved by mere manual programming. Therefore, the resulting schedules are likely to be faster and more efficient than manually programmed schedules.

Generating initial transitions by planning around tasks improves the computational efficiency of processes for generating robot motion plans. For example, planning around tasks can reduce the number of possible motion plans for each robot and reduce the number of conflicts to be resolved later in the schedule planning process. By generating initial transitions, the system does not need to evaluate every combination of robots performing tasks. The system may evaluate only the combinations of robots performing tasks that will not interfere with other task volumes. Initial transition generation can include efficiently planning around any combination of robots performing a task, and generating motion plans for multiple combinations of the robots performing the tasks. Planning robot motion plans around task volumes can reduce the total number of conflicts between robots in free motion and robots performing tasks.

The details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

5 BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram that illustrates an example system.

FIG. 2 is a flowchart of an example process for iteratively applying transformers to generate a final process definition graph.

FIG. 3 illustrates an example underconstrained process definition graph.

10 FIGS. 4A-B illustrate generating transitions for a portion of a process definition graph.

FIG. 5 is a visual illustration of a schedule.

FIG. 6 is a flowchart of an example process for generating robot motion plan schedules.

15 FIGS. 7A-C are diagrams that illustrate initial transition generation.

FIG. 8 is a flowchart of an example process for generating initial transitions.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

20 FIG. 1 is a diagram that illustrates an example system 100. The system 100 is an example of a system that can implement the techniques described in this specification.

The system 100 includes a number of functional components, including a process definer 110, a planner 120, a collection of transformers 130, a user interface device 140, an onsite execution engine 150, and a robot interface subsystem 160. Each of these
25 components can be implemented as computer programs installed on one or more computers in one or more locations that are coupled to each other through any appropriate communications network, e.g., an intranet or the Internet, or combination of networks.

The system 100 also includes a workcell 170 that includes N robots 170a-n. The overall goal of the planner 120 and other components of the system 100 is to generate,
30 from the underconstrained process definition graph 115, a schedule that will be executed by the robots 170a-n to accomplish one or more tasks. The resulting schedule can be represented in a variety of ways and may be represented as a near-fully constrained or a

fully constrained process definition graph, e.g., the workcell-specific process definition graph 135. As mentioned above, a common goal metric for such schedules is elapsed time, and thus the planner 120 can aim to generate a schedule that causes the robots 170a-n to accomplish the one or more tasks in the shortest possible amount of time.

5 In this specification, a robot is a machine having a base position, one or more movable components, and a kinematic model that can be used to map desired positions, poses, or both in one coordinate system, e.g., Cartesian coordinates, into commands for physically moving the one or more movable components to the desired positions or poses. In this specification, a tool is a device that is part of and is attached at the end of the
10 kinematic chain of the one or more moveable components of the robot. Example tools include grippers, welding devices, and sanding devices. That is, a robot can include one or more tools.

 In this specification, a task is an operation to be performed by a tool. For brevity, when a robot has only one tool, a task can be described as an operation to be performed
15 by the robot as a whole. Example tasks include welding, glue dispensing, part positioning, and surfacing sanding, to name just a few examples. Tasks are generally associated with a type that indicates the tool required to perform the task, as well as a position within a workcell at which the task will be performed.

 In this specification, a motion plan is a data structure that provides information for
20 executing an action, which can be a task, a cluster of tasks, or a transition. Motion plans can be fully constrained, meaning that all values for all controllable degrees of freedom for the robot are represented explicitly or implicitly; or underconstrained, meaning that some values for controllable degrees of freedom are unspecified. In some implementations, in order to actually perform an action corresponding to a motion plan,
25 the motion plan must be fully constrained to include all necessary values for all controllable degrees of freedom for the robot. Thus, at some points in the planning processes described in this specification, some motion plans may be underconstrained, but by the time the motion plan is actually executed on a robot, the motion plan can be fully constrained. In some implementations, motion plans represent edges in a task graph
30 between two configuration states for a single robot. Thus, generally there is one task graph per robot. In some implementation, a motion plan can include instructions for the robot to “rest,” i.e., to stay in the current position.

 In this specification, a motion swept volume is a region of the space that is occupied by a least a portion of a robot or tool during the entire execution of a motion

plan. The motion swept volume can be generated by collision geometry associated with the robot-tool system.

In this specification, a transition is a motion plan that describes a movement to be performed between a start point and an end point. The start point and end point can be represented by poses, locations in a coordinate system, or tasks to be performed.

Transitions can be underconstrained by lacking one or more values of one or more respective controllable degrees of freedom (DOF) for a robot. Some transitions represent free motions. In this specification, a free motion is a transition in which none of the degrees of freedom are constrained. For example, a robot motion that simply moves from pose A to pose B without any restriction on how to move between these two poses is a free motion. During the planning process, the DOF variables for a free motion are eventually assigned values, and path planners can use any appropriate values for the motion that do not conflict with the physical constraints of the workcell.

In this specification, a schedule is data that assigns each task to at least one robot. A schedule also specifies, for each robot, a sequence of actions to be performed by the robot. A schedule also includes dependency information, which specifies which actions must not commence until another action is finished. A schedule can specify start times for actions, end times for actions, or both.

To initiate schedule generation, a user can provide a process description and workcell-specific information to a process definer. The process description can include a high-level description of the tasks to be completed. The workcell-specific information can include data that represents properties of the workcell, including physical dimensions, the locations and dimensions of obstacles or other hardware in the workcell, and the type and number of robots in the workcell.

From the process description and the workcell-specific information, the process definer can generate an underconstrained process definition graph. Alternatively or in addition, a user can provide the underconstrained process definition graph to the planner. In this specification, a process definition graph, or for brevity, a graph, is a directed acyclic graph having constraint nodes and action nodes.

Action nodes represent actions for a robot to perform, which can include nodes representing tasks or clusters of tasks, e.g., as specified in the original process description. Action nodes can also represent transitions that robots can perform, e.g., transitions between tasks or other locations in the workcell.

Constraint nodes represent particular relationships between children nodes that must be preserved in any schedule. In general, constraint nodes can represent existence constraints or time constraints. An existence constraint specifies a constraint on which children can be selected. A time constraint specifies a constraint on the timing among children. For example, a constraint node can have as children two action nodes, and the constraint node can represent a time constraint that a first action represented by a first of the two action nodes must start before a second action represented by a second of the two action nodes.

Being an underconstrained process definition graph means that the graph lacks various kinds of information required to actually drive the robots 170a-n to accomplish the tasks. The graph can be underconstrained in a variety of ways. For example, the graph can lack any sense of time, scheduling, or ordering between tasks. Being underconstrained can also mean that various properties of task nodes are partially defined or undefined.

Thus, in some implementations the planner 120 can receive a process definition graph 115 having nodes representing the tasks to be performed as specified by the process description, but without specifying any ordering between the tasks, without specifying any assignment of any of the robots 170a-n to any of the tasks, and without specifying what movements the robots should undertake in order to prepare to perform the tasks.

The planner 120 can then perform an iterative process to begin solving constraints in the underconstrained process definition graph 115. The final output of this process is data representing a schedule, which can be a workcell-specific process definition graph 135, which, for brevity may also be referred to as a final schedule. The workcell-specific process definition graph 135 contains enough information to drive the robots 170a-n to complete the one or more tasks specified in the original underconstrained process definition graph 115. Thus, the workcell-specific process definition graph 135 will generally specify which robots will be performing which tasks. The workcell-specific process definition graph 135 can also specify the timing, scheduling, ordering and movement actions to be taken by each robot between tasks. Additionally, the movements specified by the workcell-specific process definition graph 135 can take into account the physical attributes and obstacles of the workcell 170.

The onsite execution engine 150 receives the workcell-specific process definition graph 135 and issues commands 155 to the robot interface subsystem 160 in order to actually drive the movements of the moveable components, e.g., the joints, of the robots

170a-n. In some implementations, the robot interface subsystem 160 provides a hardware-agnostic interface so that the commands 155 issued by onsite execution engine 150 are compatible with multiple different versions of robots. During execution the robot interface subsystem 160 can report execution data 157 back to the onsite execution engine 5 150 so that the onsite execution engine 150 can make real-time or near real-time adjustments to the robot movements, e.g., due to local faults or other unanticipated conditions.

In execution, the robots 170a-n generally continually execute the commands specified explicitly or implicitly by the motion plans to perform the various tasks or 10 transitions of the schedule. The robots can be real-time robots, which means that the robots are programmed to continually execute their commands according to a highly constrained timeline. For example, each robot can expect a command from the robot interface subsystem 160 at a particular frequency, e.g., 100 Hz or 1 kHz. If the robot does not receive a command that is expected, the robot can enter a fault mode and stop 15 operating.

In some implementations, the planner 120 and the process definer 110 are cloud-based systems that are physically removed from a facility that physically houses the workcell 170, while the onsite execution engine 150 is local to the facility that physically 20 houses the workcell 170. This arrangement allows the planner 120 to use massive cloud-based computing resources to consider many possibilities for robot schedules, while also allowing for real-time reaction to unanticipated events by the onsite execution engine 150.

As stated above, the planner 120 can generate a workcell-specific process definition graph 135 from the initial underconstrained process definition graph 115. To 25 do so, the planner 120 can repeatedly apply a number of transformers from a collection of transformers 130. Each transformer is a stateless function that takes as input an underconstrained process definition graph and resolves variables in the underconstrained process definition graph. As part of this process, a transformer can modify action nodes, constraint nodes, or both, by adding, modifying, or deleting these nodes, and generates as 30 output a modified process definition graph. This process is described in more detail below with reference to FIG. 2.

The transformers to be used can also be specified by a user as a transformer configuration 117. In other words, the user can specify which of the collection of

transformers 130 to be used when iteratively modifying the initial underconstrained process definition graph 115.

The planner 120 can also optionally invite users to make transformer selections while generating a final schedule. In other words, the planner 120 can provide a candidate process definition graph 123 to a user interface device 140. The user interface device 140 can then present a user interface that allows a user to input a user transformer selection 125, which directs the planner 120 to perform the next iteration using a transformer selection 127 specified by the user. This interactive process can allow the planner 120 to take into consideration constraints and other real-world considerations that were not or could not be specified as part of the original process description.

FIG. 2 is a flowchart of an example process for iteratively applying transformers to generate a final process definition graph. The process can be implemented by one or more computer programs installed on one or more computers in one or more locations and programmed in accordance with this specification. For example, the process can be performed by the planner 120 shown in FIG. 1. For convenience, the process will be described as being performed by a system of one or more computers.

The system receives an underconstrained process definition graph (210). As described above, a process definition graph is a directed acyclic graph that can include at least two different types of nodes, constraint nodes and action nodes.

FIG. 3 illustrates an example underconstrained process definition graph 300. The graph includes a root node 310, two constraint nodes 322 and 324, and six action nodes 332, 334, 342, 344, 346, and 348.

Each of the action nodes represents one or more physical actions to be performed by a particular robot. For example, an action node can represent an action to move a robot arm, apply a weld, open a gripper, close a gripper, or any other appropriate action that can be performed by a robot. Action nodes can also be composite actions that include multiple other actions. For example, a PickAndPlace composite action can include a separate pick action and a place action.

Notably, the action nodes need not specify which robot in a workcell will perform each action. The action nodes can be partially defined either by specifying some values, e.g., a duration or a location for each action. The action nodes can also be partially defined by specifying options, e.g., a list or a range of acceptable durations.

The AllofInOrder constraint node 322 specifies that its children nodes have to be performed in a particular order. In other words, the constraint node 322 introduces a

constraint into the plan that Action 1, represented by the action node 332, has to occur before any of the other actions, and that Action 7, represented by the action node 334, has to occur after any of the other actions.

5 The AllOfInOrder constraint node 322 specifies that its children nodes have to be performed in a particular order. In other words, the constraint node 322 introduces a constraint that Action 1, represented by the action node 332, has to occur before any of the other actions, and that Action 7, represented by the action node 334, has to occur after any of the other actions.

10 The AllOf constraint node 322 specifies simply that all of its children have to be performed, but does not specify an ordering. Thus, through the iterative planning process, the planner is free to schedule Actions 3-6 in any appropriate order that does not violate any other constraints in the graph.

15 One concrete example of an application that might use a graph similar to the graph 300 is welding a part in a factory. For example, the initial action node 332 can represent an action that moves the part into place, and the final action node 334 can represent an action to move the part to its destination. Those are examples of actions that have to happen first and last, as specified by the AllOfInOrder constraint node 322.

20 Actions 3-6 can be welding actions for four separate weld points on the part. For most welding applications, the order does not matter. Thus, while the AllOf constraint node 324 imposes a constraint that all welds have to be performed, the AllOf constraint node 324 does not impose any constraints on the ordering in which the welds must be performed.

25 Instead, the planner will use various transformers in order to search the space of schedules that satisfy the constraints imposed by the final graph and choose a schedule that is best according to a particular objective function. For example, the planner can simply search the space of schedules that satisfy the constraints and identify which schedule executes all the actions the fastest. TABLE 1 lists some common constraint node types.

Constraint Node Name	Constraint
AllOf	All children must be performed
AnyOf	Any children can be performed
OneOf	Exactly one child must be performed
InOrder	Children must be performed in an order

	specified by the representation in the graph
DirectlyInOrder	Children must be performed in an order specified by the representation in the graph, without any intervening actions
MustOccurDuring	Child must be performed during a specified time period
MustNotOccurDuring	Child must not be performed during a specific time period
Conflict	Children cannot occur at overlapping times
Cluster	All children must be performed by the same robot
If	If one child occurs, the subsequent child must happen after it

TABLE 1

A constraint node can combine multiple constraint node types. For example, in this notation, the AllOfInOrder node had a name that specified two constraint types, both the AllOf constraint type as well as the InOrder constraint type.

5 As shown in FIG. 2, the system applies local transformers (220). As described above, each transformer is a function that further constrains the process definition by assigning values to variables or manipulating or expanding nodes in the graph.

In general a system can apply a transformer by matching certain patterns associated with nodes in the graph. As one example, applying a transformer can include
 10 looking for underconstrained properties in the graph and assigning values to those properties. As another example, a transformer can have one or more defined node types as input and the system can determine that the transformer is a match for the node types when a matching arrangement of nodes found in the graph.

The system can distinguish between local transformers and global transformers
 15 and apply local transformers first. A local transformer is a transformer that does not require reconsideration of all nodes in the graph and which affects only a sufficiently small subset of the graph. The system can use any appropriate definition for what is a sufficiently small subset of the graph, e.g., only transformers that affect their direct

children or nodes not more than N links away. As a particular example, a local transformer can be applied to a single node in the graph.

An example of a global transformer is a “clustering” transformer, which takes consideration all tasks in the graph that change the position of one or more robots, and proposes ordering constraints that ensure that the robots move efficiently between tasks (for example, avoiding doubling back by robots where possible).

The system can apply local transformers first in order to quickly generate many additional constraints in the graph. For example, if an action node has a particular pose, the system can apply an inverse kinematics (IK) transformer that will generate the kinematic parameters for achieving the pose. Thus, for a given pose, the system can consider an IK transformer to be a local transformer because the transformer affects only a single action node.

The system generates transitions (230). Transitions are actions taken by robots to move from one configuration to another. Thus, if two actions are to be performed by a robot in sequence, the system can generate a transition between the actions by determining how the robot can move from a pose for one action to a pose for another. In some implementations, the system implements transition generators as transformers that seek to match on two action nodes in sequence that so far have no assigned intermediate transition.

Some transformers are designed to generate many alternative options that can all be considered when performing scheduling. For example, when generating transitions, the system can generate multiple different ways of moving from one action to another. The system can represent each generated alternative transition as a separate node in the graph. Since only one transition is needed, the system can constrain the alternatives with an appropriate OneOf constraint node.

FIGS. 4A-B illustrate generating transitions for a portion 400a of a process definition graph. In FIG. 4A, an AllOfInOrder constraint node 410 represents that a welding action represented by an action node 420 should be performed before a welding action represented by another action node 430. In some implementations, during transition generation the system searches the graph for gaps between sequenced actions that lack transitions. Thus, the system can identify the portion 400a as a portion of the graph having two sequenced actions but lacking a transition between them.

FIG. 4B illustrates a modified portion 400b of the process definition graph after transition generation. As shown, the system generated four alternative transitions for

transitioning between the action node 420 and the action node 430. Each alternative transition is represented by its own action node, e.g., the transition nodes 450, 460, 470, and 480. The transition nodes 450-480 are constrained by a OneOf constraint node 440 that represents that the system should execute one and only one of the generated
5 alternative transitions.

This example illustrates how transformers can add additional constraints to the graph in the form of alternatives rather than selections. In other words, the transformers seek to increase the space of possible alternatives schedules rather than attempting to solve the constraints of a schedule in one pass.

10 Therefore, after the system performs transition generation, the process definition graph typically has many more nodes that it previously did, with each node representing a possible transition between actions.

As shown in FIG. 2, the system performs conflict identification (240). In other words, the system identifies actions in the graph that cannot occur at the same time, the same space, or both. For example, the system can generate a swept volume for each
15 action node in the graph and identify which actions are potentially conflicting. Notably, the system can perform the conflict identification process before scheduling occurs. The system can then perform a deconfliction process after scheduling occurs. In other words, the system identifies conflicts before scheduling, but need not strive to avoid conflicts
20 during scheduling.

The system performs scheduling (250). In general, the scheduling process solves for underconstrained values in the graph until either no more values need to be solved or the system determines that no solution can be found. As described above, a schedule specifies one robot to perform each task, and for each robot, a sequence of actions to be
25 performed. A schedule can also specify a start time for each action and dependencies between actions.

For example, to perform scheduling, the system can receive as input a set of possible graphs, e.g., a set of every possible alternative graph that can be generated from the graph and that accomplishes the required tasks. The system can process the set of
30 possible graphs to generate an output that is a selection of action nodes to execute as well as dependencies (i.e., a selection of constraint nodes) that specify requirements of the schedule. That is, the output of the scheduling process can be an updated graph, generated from the original graph, with extra constraints that remove the uncertainties of the original graph.

As a particular example, referring to the example depicted in FIG. 4B, the output of the scheduling process might be a graph that includes an AllOfInOrder node with three child nodes 420, 460, and 430. That is, the system selected the second transition 460 from the set of possible transitions 450, 460, 470, and 480, removing the uncertainty of the OneOf node 440.

As another particular example, there may be a situation where two robots need to cross each other, the system can identify (e.g., during conflict identification) the constraint that the two robots cannot collide when crossing each other. The output of the scheduling process can therefore include a scheduling constraint that ensures that the two robots do not perform the movement at the same time, for example, by generating a dependency between the completion of the first movement and the beginning of the second movement (e.g., using an InOrder node or a MustNotOccurDuring node).

In general, the system can use a variety of solvers to determine that no more constraints need to be solved. For example, the system can use a circuit solver to determine when sequenced starting and ending points have no gaps in between. If gaps remain, the system can use transition generation to fill the gaps or raise an error if no solution exists.

The system can also insert rests into the actions and adjust the values of the rests in order to find a scheduling solution. For example, one perhaps not ideal schedule is that a robot can perform an action and then wait for all other robots to finish their tasks before doing anything else. But by using rests, the system increases the chances that a scheduling solution can be found.

As part of the scheduling process, the system can assign tasks to individual robots in the workcell. Thus, at the end of the scheduling process, the system can generate a visual representation of which tasks are performed by which robots and when.

FIG. 5 is a visual illustration of a schedule. The information illustrated in FIG. 5 can still be represented as a process definition graph. And in fact, the deconfliction process that occurs after scheduling can operate on the process definition graph generated by the scheduling process. However, the visualization shown in FIG. 5 is useful for understanding the constraints encoded into the graph.

In FIG. 5, three robotic components in a workcell have been assigned to perform a variety of tasks. Time generally moves from left to right in FIG. 5, and the arrows represent time ordered constraints between actions by different robotic components. For

actions performed by the same robotic component, the ordering of the actions implies a time ordered constraint.

Thus, the example schedule illustrated in FIG. 5 first has the fixturing 510 of the workcell perform a part transfer 512 and then a close clamps action 514. For example,
5 the close clamps action 514 can represent clamps closing down on the part to secure it for welding.

At some time after the transfer 512 starts, the first robot 520 performs a joint move 522 to get into position for welding. However, the action weld action 524 cannot happen until the actual clamps are closed, as illustrated by the arrow from the closed
10 clamps action 514 to the weld action 524.

At some time after the clamps are closed, a second robot 530 performs a joint move 532 to get into position for a weld 534. The joint move 532 also has a constraint that it has to happen after the clamps are closed by the close clamps action 514.

The first robot 520 can then perform a second joint move 526 to get into position
15 for the second weld 528, but not until the second robot has finished its joint move 532, as represented by the arrow from the joint move 532 to the joint move 526. The first robot can then perform the second weld 528.

Finally, the fixturing 510 can open the clamps with an open clamps action 516,
20 but not until both welds 528 and 534 are completed. The fixturing 510 can then perform a transfer action 518 to move the part along in an assembly line.

As shown in FIG. 2, the system performs deconfliction (260). In general, the deconfliction process also uses transformers that aim to solve possible conflicts in the schedule. As described above, although conflicts may have been identified before
25 scheduling, the scheduler was not necessarily bound by such conflicts during scheduling. In general, generating an initial solution using an underconstrained graph that ignores some conflicts provides computational advantages over trying to solve all constraints and all possible conflicts in one pass.

The system optionally runs additional user-selected transformers (270). As described above, the system can provide a user interface that seamlessly allows some
30 human design in the process. In particular, the system can present a user interface that allows a user to manually specify a next transformer to be applied in the graph or some other manipulation of the graph. As part of this user interface, the system can also present a graphical simulation of the robots executing a particular candidate schedule. For example, if a particular transition between actions seems too awkward or otherwise

not ideal, a human can manually select a different transition. As another example, humans tend to be better at performing geometric or spatial groupings. Thus, a human may want to impose a Cluster constraint node for a particular group of actions that are close together in space, time, or both. This process is described in more detail below with
5 reference to FIG. 6.

The system determines whether a goal has been reached (280). The system can use a goal solver to determine whether a process definition graph meets various goal criteria. As mentioned above, total time is often a critical goal criterion. Each action in the schedule can be associated with a respective duration, and then the system can use the
10 respective durations to determine whether a schedule exists that meets the goal criteria. The system can also use other goal criteria, for example, power used, or some combined score of multiple goal criteria.

If the goal is reached (280), the system outputs the final schedule (branch to 285). For example, the final schedule can then be provided to an onsite execution engine for
15 execution by robots in a workcell.

If the goal is not reached, the system determines whether to perform more iterations (290). In some implementations, the system automatically performs another iteration as long as the system has not determined that a solution is not possible. The system can alternatively or in addition perform an additional iteration only if a maximum
20 number of iterations has not been reached.

If no more iterations are to be performed and the goal has not been reached, the system can raise an error (branch to 295). At that point, the users can determine whether to modify the original process definition in order to try to find a valid solution with
25 different inputs.

In some implementations, as depicted in FIG. 2, if more iterations are to be performed, the system returns to step 220 and reapplies local transformers. In some other implementations, the system can return to step 240 to perform another round of conflict identification. That is, the system might not need to reapply the local transformers or to
generate transitions again.

One of the advantages to representing all phases of the planning process as a
30 process definition graph is that the history of the process is represented very explicitly. The state of the schedule as it evolved over the iterations becomes quite clear. This allows actions to be easily undone and the schedule's modifications rewound in time and tried again with different transformers or parameters. This also makes debugging the

schedules much easier because it becomes clear how the schedule got to where it is and by which transformers.

In some implementations, the system can perform the process depicted in FIG. 2 multiple times to generate multiple different candidate schedules. For example, the system might perform the process multiple times in parallel, e.g., by selecting different sequences of transformations to generate the different candidate schedules. The system can then select the final schedule from the set of candidate schedules according to one or more criteria, e.g., a time to complete the one or more required tasks, a measure of efficiency, a measure of safety, etc. That is, the system can search the space of possible schedules by evaluating different transformation sequences. As a particular example, the system might evaluate a particular sequence of transformation sequences and determine to “backtrack” to a particular transformation in the sequence and being a new sub-sequence from that point6 is a flowchart of an example process 600 for generating schedules. The example process 600 illustrates an example technique for iteratively applying various transformers to a process definition graph. The process can be implemented by one or more computer programs installed on one or more computers in one or more locations and programmed in accordance with this specification. For example, the process can be performed by the planner 120 shown in FIG. 1. For convenience, the process will be described as being performed by a system of one or more computers.

The system plans for reachable tasks (610). Starting with the task requirements of the workcell, the system finds for each robot, a motion plan representing a transition to all of the tasks that the robot can reach. As described above, a motion plan can provide some or all information for executing a transition.

The system generates initial transitions by planning around tasks (612). The system plans transitions that avoid motion swept volumes for each task. Thus, the transitions avoid conflicts with any a robot that might be performing each task. To do so, the system can plan the transitions iteratively by considering every executable combination of robots performing each task. Initial transition generation thus provides an initial set of possible transitions that avoid volumes that are occupied by tasks. Initial transition generation is described in more detail below with reference to FIGS. 7A-C and FIG. 8.

The system solves, ignoring replannable motion plans and fixturing conflicts (614). In this context, solving refers to generating a schedule that resolves as many

underconstrained values as possible in the process definition graph. A replannable motion plan is a motion plan that is underconstrained and thus is flexible to be modified later on in the process, e.g., during conflict resolution. Some example criteria for determining if a motion plan is replannable are length, duration, and moving volume of the motion plan. In some implementations, the volume of intersection between a swept volume pair can be used to define the severity of a conflict. For example, a replannable motion plan can be a motion plan in which the volume of the intersection between the swept volume pair is less than a particular threshold value, e.g., 10, 100, or 1000 cubic centimeters.

10 The system repairs small conflicts (616). In other words, the system modifies the process definition graph to fix conflicts that are under the particular threshold value. To do so, the planner can run an algorithm that replans around the movement that is known to be happening during the execution of the motion plan in conflict. Both motion plans involved in a conflict are then replanned. Given a motion plan involved in one or many
15 conflicts, all of the motion plans with an overlapping execution schedule to the motion plan in concern are collected. From that collection, two sets are formed: the set of conflicting plans and the set of non-conflicting plans. For some conflicts, the system reschedules the motion plan so that the conflict is avoided. In this way, the algorithm for repairing small conflicts replans a motion plan against all of the possible combinations of
20 conflicting plans.

 The system solves, ignoring fixturing conflicts (618). In other words, the system resolves for the schedule using the repaired small conflicts but still ignoring conflicts with the fixturing.

 The system replans conflicting ends (620). The system can resolve end conflicts separately from mid-motion conflicts. For example, when a robot approaches the fixture, the end of the motion may be in conflict with the fixture. The system finds a waiting point so that the robot can be as close as possible to the task without being in collision with the movement of the fixture. The system then solves again (622) with the replanned ends.

30 The system splits (624) motion plans. Splitting creates two or more paths from an existing path. The system can split motion plans in order to resolve conflicts by allowing the conflicting parts of paths to be scheduled independently. The system can introduce a stopping point to separate two paths in time.

The system performs additional sequences of repair and solving of the motion plan schedules (626, 628, 630). The solution includes the motion plan execution sequence, or schedule, for each robot of the workcell. The schedule can specify each robot's motion actions that avoid conflict, including conflict with the volumes occupied by the task nodes of the process definition graph. The end of the process can be a fully constrained process definition graph, which can be executed by an onsite execution engine, e.g., the onsite execution engine 150 of FIG. 1.

FIGS. 7A-C are diagrams that illustrate initial transition generation. Initial transition generation includes planning around tasks, as described above with reference to step 612 of the process 600. Initial transition generation provides an initial set of possible transitions that avoid volumes of space that will be occupied by robots performing tasks.

The system can perform initial transition generation before generating any schedule and thus before assigning any tasks to individual robots. Instead, the system can use the initial transition generation process to generate multiple different possible assignments of robots to tasks, with each robot transitioning to each possible task in a way that avoids motion swept volumes of tasks in the workcell. Thus, initial transition generation may include planning around every combination of robots performing a task, and generating motion plans for multiple combinations of the robots performing the tasks. During initial transition generation, the system can plan for all possible free motions, avoiding the volumes of space that will be occupied by robots performing the tasks. To plan a free motion, the system can use any non-conflicting space to move a robot to a destination. By planning around tasks in initial transition generation, the system can reduce the number of conflicts that need to be resolved later in the schedule generating process.

FIG. 7A illustrates a workcell 700 that includes three welding robots 170a-c. Each task of the workcell 700 is a welding action, e.g., Weld1 to Weld6, to be performed by one of the robots 170a-c. The tasks are not assigned to any particular robot, and thus it is possible at this stage of the planning process that any robot can be assigned to perform any of the tasks. In addition, at this stage in the process, the tasks have no ordering, and thus, the robots 170a-c can perform the tasks in any order.

Each task occupies a motion swept volume, or for brevity, a task volume. The tasks Weld1 to Weld6 thus each occupy one of the task volumes 701, 702, 703, 704, 705, and 706, respectively. Each task volume represents a volume of space that may be occupied by a robot as it performs the respective task. Each task is required to be

performed by one of the robots. Therefore, each of the task volumes will be occupied during a period of time of the final resulting schedule. Thus, though the task volumes may not be occupied continuously, task volumes have a higher probability of robot occupancy at any given time, compared to surrounding volumes of space.

5 For example, an example task Weld3 occupies an example task volume 703. As a robot performs the task Weld3, the task volume 703 is occupied, and a conflict occurs if another robot crosses through task volume 703 at the same time. Initial transition generation can avoid transitions that cross through task volume 703, thus reducing the likelihood of conflict.

10 FIG. 7B illustrates an example assignment of initial transitions that avoid task volumes. A transition path can be a path that a robot takes while transitioning between tasks. Transition paths can connect any of two task volumes. For example, a transition path 720 connects task volume 701 and task volume 705. Transition paths can also connect a robot start or end position with a task volume. For example, transition path 713
15 connects a robot start position 172c for the robot 170c with task volume 703.

The planner can perform initial transition generation by generating transition paths that do not cross through task volumes. For example, the planner can generate transition path 720. Transition path 720 is a free motion path between task volume 701 and task
20 volume 705. The shortest straight-line distance between the task volume 701 and the task volume 705 would cross into the task volume 706. Thus, the planner can generate the transition path 720 that causes a robot to completely navigate around the task volume 706. Thus, a robot scheduled to move along transition path 720 will not interfere with robots performing Weld6. Notably, the system can generate the conflict-free transition path 720 before determining whether or not any robot will actually occupy the task
25 volume 706 at the time that the transition path 720 is traversed.

In some examples, it may not be possible for the planner to assign all tasks to robots while planning around all task volumes. For example, planning around all task volumes may result in too many constraints such that the schedule becomes unsolvable. This can result due to particular constraints on entry and exit points of tasks in the process
30 definition graph. For example, some tasks may not have any suitable exit points, and thus, candidate schedules that specify executing those tasks early on may be unsolvable. The planner can thus apply an algorithm for planning around tasks that improves computational efficiency by exploring many candidate assignments while avoiding overconstraining the solution.

FIG. 7B illustrates an example of planning around tasks. The technique illustrated by FIG. 7B first establishes an order of tasks. At this phase of the process, before a schedule has been generated, the order represents merely the order in which robots will be assigned to perform the tasks, rather than an actual order in time that robots will actually perform the tasks. In other words, the order exists only for the purposes of initial assignments and transition generations, rather than for specifying an ordering in time during execution.

For example, the tasks in FIG. 7B can be ordered Weld1, Weld2, Weld3, Weld4, Weld5, and Weld6. The planner can then iteratively generate motion plans for robots to perform respective tasks until no additional motion plans can be generated, e.g., due to no scheduling solutions existing. In this process, the planner can assign any one of the robots 170a-c to the first task in the order, and generate a motion plan for the robot to perform the first task. The planner can then assign any one of the robots 170a-c to the second task in the order, and generate a corresponding motion plan for the robot performing the second task. The planner can continue through the tasks in order, generating motion plans for each robot in sequence. During this process, the system can assign one robot to perform two or more tasks.

For example, the planner can start with Weld1 by assigning a robot, e.g., the robot 170a, to perform Weld1 in task volume 701. The system can apply one or more heuristics for determining which robot to perform Weld1. In this example, the robot 170a was simply the closest robot to the task volume 701. For subsequent tasks, the system can consider other factors besides distance, e.g., the likelihood that a robot will be available to perform a task at a particular time.

After assigning the robot 170a to Weld1, the system can generate a motion plan 711 for the robot 170a to transition to performing Weld1 in task volume 701. In FIG. 7B, transition paths performed by the robot 170a are represented by solid lines.

The system can turn to the next task in the ordering, Weld2, by assigning a robot, e.g., the robot 170b, to perform Weld2. The system can generate a motion plan 712 for the robot 170b to perform Weld2 in task volume 702. In FIG. 7B, transition paths performed by the robot 170b are represented by dashed lines.

The system can turn to the next task in the ordering, Weld3, by assigning a robot, e.g., the robot 170c, to perform Weld3. The system can generate a transition path 713 for the robot 170c to perform Weld3 in task volume 703. In FIG. 7B, transition paths performed by the robot 170c are represented by dashed and dotted lines.

Notably, the transition path 713 does not take a straight-line path between the start position 172c and the task volume 703. Instead, the system generates the transition path 713 to avoid the task volume 704, even though it is unclear at this point in the process whether any robot will actually be occupying the task volume 704 at the time that the robot 170c follows the transition path 713.

The system can continue generating initial transitions for each task in the initial ordering of task, including generating motion plans for respective robots to perform each task. For example, the system assigns the robot 170b to perform Weld4 by following transition path 714.

The system also assigns the robot 170a to perform Weld5 by following the transition path 720 that explicitly avoids the task volume 706. In other words, the system generates the initial transition between the task volume 701 and the task volume 706 to avoid the task volume 706 even though it is not clear that any robot will be in the task volume 706 at the time that the robot 170a follows the transition path 720.

Finally, the system assigns the robot 170c to perform Weld6 by following the transition path 721.

In some situations, the system may not be able to plan for a certain task. For example, after planning for the tasks Weld1, Weld2, and Weld3, the system may not be able to plan for the task Weld4. The system may not be able to generate a motion plan for a robot to perform Weld4 while avoiding the task volumes 701, 702, 703, 705, and 706. In some examples, the system can determine not to plan for the task Weld4, and continue the sequence to plan for Weld5, followed by Weld6. In some examples, the planner can determine not to plan for the task Weld4, and to stop the sequence. These situations are not necessarily failure states because the scheduling process can always simply assign each task to be performed in serial by any robot that can reach it.

FIG. 7C illustrates initial planning around tasks due to a different task ordering. After evaluating the tasks in a particular order, the planner can then reorder the tasks. For example, the system can reorder the tasks, e.g. randomly or pseudorandomly, as Weld2, Weld4, Weld3, Weld6, Weld5, and Weld1. The planner can then again iteratively perform initial transition generation for each task in the modified task order.

The planner first assigns the first task in the modified ordering, Weld2, to the robot 170b, which in this example simply happens to be the closest robot to the task volume 702. The system generates a transition path 731 that transitions the robot from the start position 172b to the task volume 702.

The planner can then turn to the next task in the modified ordering, Weld4, by assigning a robot, e.g., the robot 170c, to perform Weld4. The system can generate a transition path 732 for the robot 170c to perform Weld4 in the task volume 704.

5 The planner can then turn to the next task in the modified ordering, Weld3, by assigning a robot, e.g., the robot 170a, to perform Weld3. The system can generate an initial transition path 733 that transitions the robot 170a from the start position 172a to the task volume 703 and which also avoids any other task volumes. In this example, the task volume 701 is situated between the start position 172a and the target task volume 703. Thus, the planner generates the transition path 733 so that it avoids the intervening task
10 volume 701.

Similarly, for Weld6 and Weld5 in that order, the system assigns Weld6 to the robot 170c and generates a transition path 734 that avoids the task volume 703 and then assigns the Weld5 to the robot 170b and generates a transition path 735 that avoids the intervening task volume 703 as well.

15 And as stated above, each of these transition paths avoid the task volumes before it is decided which robots will be performing the tasks in those task volumes and before it is decided whether a robot will be in the task volume at all at the time that the transition path is traversed. This means, for example, that system may generate initial transitions that cause a robot to avoid a task volume for a task that that exact same robot may
20 perform at another time. Although this seems counterintuitive, the initial transition generation process that avoids task volumes allows the system to more efficiently explore a large number of plausible schedules by filtering out schedules that are likely to result in conflicts that are expensive to resolve in terms of computational resources, time, human input, or some combination of these. This process also tends to filter out implausible
25 schedules that have no solution earlier in the process. As shown in FIG. 6, the initial transition generation might be the second step of the process. Thereafter, as the initial schedules are further elaborated with many different alternatives, the initial transition generation process ensures that the computational resources to explore these many different alternatives are expended only on schedules that are plausibly solvable in the
30 first instance.

In other words, by generating initial transitions that avoid task volumes early in the process, the system can explore a large search space without actually evaluating every conceivable combination of robots performing tasks. Rather, the system can evaluate

only the combinations of robots performing tasks that will not interfere with other task volumes.

The system can continue to evaluate multiple randomly generated orderings of tasks as described above. By evaluating multiple orderings of tasks, the system can efficiently assess many combinations of assigning robots to tasks, while avoiding conflicts between tasks and robot free motion. After reordering the tasks multiple times, the planner can select an ordering of tasks that optimizes the set of initial transitions. For example, the planner can select an ordering of tasks that optimizes any or all of the number of initial transitions, the number of assigned tasks, the number of avoided task volumes, the number of conflicts between transitions and tasks, or the transition time.

After generating initial transitions, the planner can continue the process 600, e.g., by solving, ignoring replannable motion plans and fixturing conflicts (614). Initial transition generation may not resolve all conflicts between motion plans and task volumes. Remaining conflicts between motion plans and task volumes may be solved at a later step in the planning process.

FIG. 8 is a flowchart of an example process for generating initial transitions. The process can be implemented by one or more computer programs installed on one or more computers and programmed in accordance with this specification. For example, the process can be performed by the planner 120 shown in FIG. 1. For convenience, the process will be described as being performed by a system of one or more computers.

The process includes receiving a process definition graph having a plurality of task nodes to be performed by a robot (810), generating, from the process definition graph, an initial modified process definition graph that adds constraints for respective volumes occupied by each task (820), and generating, from the initial modified process definition graph, a refined process definition graph including respective motion plans for robots moving between tasks, the motion plans avoiding the volumes occupied by each task (830).

In more detail, the system receives a process definition graph having a plurality of task nodes that represent respective tasks to be performed by a robot of a plurality of robots, each task node being associated with a location at which the task will be performed (810). The process definition graph can include for example, the portion 400a of a process definition graph. The task nodes can be, for example, the action nodes 420, 430, to be performed by one of the robots 170a-c. Each of the task nodes 420, 430, is associated with a location at which the tasks Weld1, Weld2, respectively, will be

performed. The location can be, for example, a volume of space of a workcell, e.g., the workcell 700.

The system generates, from the process definition graph, an initial modified process definition graph that adds constraints for respective volumes occupied by each task represented by the plurality of task nodes (820). For example, the initial modified process definition graph can add a constraint for avoiding the volume occupied by the task Weld1, and for the volume occupied by the task Weld2. The initial modified process definition graph can include constraints for avoiding volumes occupied by any additional tasks of the process definition graph.

The system generates, from the initial modified process definition graph, a refined process definition graph including respective motion plans for robots moving between tasks, the motion plans avoiding the volumes occupied by each task (830). The refined process definition graph can include, for example, the portion 400b. The system can represent each generated alternative transition as a separate node in the refined process definition graph. The motion plans can include the transition nodes 450, 460, 470, and 480, where each of the transition nodes 450 to 480 avoids the volumes occupied by the task nodes 420, 430. Each of the transition nodes 450 to 480 can also avoid the volumes occupied by any additional tasks of the process definition graph. The system may generate the refined process definition graph before any tasks have been assigned to any of the robots. The system can use the refined process definition graph to generate a schedule for the robots. The schedule may specify each robot's motion actions that avoid the volumes occupied by the task nodes of the process definition graph.

The robot functionalities described in this specification can be implemented by a hardware-agnostic software stack, or, for brevity just a software stack, that is at least partially hardware-agnostic. In other words, the software stack can accept as input commands generated by the planning processes described above without requiring the commands to relate specifically to a particular model of robot or to a particular robotic component. For example, the software stack can be implemented at least partially by the onsite execution engine 150 and the robot interface subsystem 160 of FIG. 1.

The software stack can include multiple levels of increasing hardware specificity in one direction and increasing software abstraction in the other direction. At the lowest level of the software stack are robot components that include devices that carry out low-level actions and sensors that report low-level statuses. For example, robots can include a variety of low-level components including motors, encoders, cameras, drivers, grippers,

application-specific sensors, linear or rotary position sensors, and other peripheral devices. As one example, a motor can receive a command indicating an amount of torque that should be applied. In response to receiving the command, the motor can report a current position of a joint of the robot, e.g., using an encoder, to a higher level of the software stack.

Each next highest level in the software stack can implement an interface that supports multiple different underlying implementations. In general, each interface between levels provides status messages from the lower level to the upper level and provides commands from the upper level to the lower level.

Typically, the commands and status messages are generated cyclically during each control cycle, e.g., one status message and one command per control cycle. Lower levels of the software stack generally have tighter real-time requirements than higher levels of the software stack. At the lowest levels of the software stack, for example, the control cycle can have actual real-time requirements. In this specification, real-time means that a command received at one level of the software stack must be executed and optionally, that a status message be provided back to an upper level of the software stack, within a particular control cycle time. If this real-time requirement is not met, the robot can be configured to enter a fault state, e.g., by freezing all operation.

At a next-highest level, the software stack can include software abstractions of particular components, which will be referred to motor feedback controllers. A motor feedback controller can be a software abstraction of any appropriate lower-level components and not just a literal motor. A motor feedback controller thus receives state through an interface into a lower-level hardware component and sends commands back down through the interface to the lower-level hardware component based on upper-level commands received from higher levels in the stack. A motor feedback controller can have any appropriate control rules that determine how the upper-level commands should be interpreted and transformed into lower-level commands. For example, a motor feedback controller can use anything from simple logical rules to more advanced machine learning techniques to transform upper-level commands into lower-level commands. Similarly, a motor feedback controller can use any appropriate fault rules to determine when a fault state has been reached. For example, if the motor feedback controller receives an upper-level command but does not receive a lower-level status within a particular portion of the control cycle, the motor feedback controller can cause the robot to enter a fault state that ceases all operations.

At a next-highest level, the software stack can include actuator feedback controllers. An actuator feedback controller can include control logic for controlling multiple robot components through their respective motor feedback controllers. For example, some robot components, e.g., a joint arm, can actually be controlled by multiple
5 motors. Thus, the actuator feedback controller can provide a software abstraction of the joint arm by using its control logic to send commands to the motor feedback controllers of the multiple motors.

At a next-highest level, the software stack can include joint feedback controllers. A joint feedback controller can represent a joint that maps to a logical degree of freedom
10 in a robot. Thus, for example, while a wrist of a robot might be controlled by a complicated network of actuators, a joint feedback controller can abstract away that complexity and expose that degree of freedom as a single joint. Thus, each joint feedback controller can control an arbitrarily complex network of actuator feedback controllers. As an example, a six degree-of-freedom robot can be controlled by six different joint
15 feedback controllers that each controls a separate network of actual feedback controllers.

Each level of the software stack can also perform enforcement of level-specific constraints. For example, if a particular torque value received by an actuator feedback controller is outside of an acceptable range, the actuator feedback controller can either modify it to be within range or enter a fault state.

To drive the input to the joint feedback controllers, the software stack can use a command vector that includes command parameters for each component in the lower levels, e.g., a position, torque, and velocity, for each motor in the system. To expose status from the joint feedback controllers, the software stack can use a status vector that includes status information for each component in the lower levels, e.g., a position,
25 velocity, and torque for each motor in the system. In some implementations, the command vectors also include some limited information regarding constraints to be enforced by the controllers in the lower levels.

At a next-highest level, the software stack can include joint collection controllers. A joint collection controller can handle issuing of command and status vectors that are exposed as a set of part abstractions. Each part can include a kinematic model, e.g., for
30 performing inverse kinematic calculations, limit information, as well as a joint status vector and a joint command vector. For example, a single joint collection controller can be used to apply different sets of policies to different subsystems in the lower levels. The joint collection controller can effectively decouple the relationship between how the

motors are physically represented and how control policies are associated with those parts. Thus, for example if a robot arm has a movable base, a joint collection controller can be used to enforce a set of limit policies on how the arm moves and to enforce a different set of limit policies on how the movable base can move.

5 At a next-highest level, the software stack can include joint selection controllers. A joint selection controller can be responsible for dynamically selecting between commands being issued from different sources. In other words, a joint selection controller can receive multiple commands during a control cycle and select one of the multiple commands to be executed during the control cycle. The ability to dynamically
10 select from multiple commands during a real-time control cycle allows greatly increased flexibility in control over conventional robot control systems.

 At a next-highest level, the software stack can include joint position controllers. A joint position controller can receive goal parameters and dynamically compute commands required to achieve the goal parameters. For example, a joint position
15 controller can receive a position goal and can compute a set point for achieving the goal.

 At a next-highest level, the software stack can include Cartesian position controllers and Cartesian selection controllers. A Cartesian position controller can receive as input goals in Cartesian space and use inverse kinematics solvers to compute an output in joint position space. The Cartesian selection controller can then enforce limit
20 policies on the results computed by the Cartesian position controllers before passing the computed results in joint position space to a joint position controller in the next lowest level of the stack. For example, a Cartesian position controller can be given three separate goal states in Cartesian coordinates x , y , and z . For some degrees, the goal state could be a position, while for other degrees, the goal state could be a desired velocity.

25 These functionalities afforded by the software stack thus provide wide flexibility for control directives to be easily expressed as goal states in a way that meshes naturally with the higher-level planning techniques described above. In other words, when the planning process uses a process definition graph to generate concrete actions to be taken, the actions need not be specified in low-level commands for individual robotic
30 components. Rather, they can be expressed as high-level goals that are accepted by the software stack that get translated through the various levels until finally becoming low-level commands. Moreover, the actions generated through the planning process can be specified in Cartesian space in a way that makes them understandable for human operators, which makes debugging and analyzing the schedules easier, faster, and more

intuitive. In addition, the actions generated through the planning process need not be tightly coupled to any particular robot model or low-level command format. Instead, the same actions generated during the planning process can actually be executed by different robot models so long as they support the same degrees of freedom and the appropriate control levels have been implemented in the software stack.

Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non-transitory storage medium for execution by, or to control the operation of, data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus.

The term “data processing apparatus” refers to data processing hardware and encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can also be, or further include, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can optionally include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

A computer program which may also be referred to or described as a program, software, a software application, an app, a module, a software module, a script, or code) can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A program may, but need not, correspond

to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub-programs, or portions of code. A computer program
5 can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a data communication network.

For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware,
10 hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions.

As used in this specification, an “engine,” or “software engine,” refers to a software implemented input/output system that provides an output that is different from the input. An engine can be an encoded block of functionality, such as a library, a platform, a software development kit (“SDK”), or an object. Each engine can be implemented on any appropriate type of computing device, e.g., servers, mobile phones,
20 tablet computers, notebook computers, music players, e-book readers, laptop or desktop computers, PDAs, smart phones, or other stationary or portable devices, that includes one or more processors and computer readable media. Additionally, two or more of the engines may be implemented on the same computing device, or on different computing devices.

The processes and logic flows described in this specification can be performed by
25 one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by special purpose logic circuitry, e.g., an FPGA or an ASIC, or by a combination of special purpose logic circuitry and one or more
30 programmed computers.

Computers suitable for the execution of a computer program can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of

a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. The central processing unit and the memory can be supplemented by, or incorporated in, special purpose logic circuitry. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

Computer-readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks.

To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and pointing device, e.g., a mouse, trackball, or a presence sensitive display or other surface by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's device in response to requests received from the web browser. Also, a computer can interact with a user by sending text messages or other forms of message to a personal device, e.g., a smartphone, running a messaging application, and receiving responsive messages from the user in return.

Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface,

a web browser, or an app through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a
5 communication network. Examples of communication networks include a local area network (LAN) and a wide area network (WAN), e.g., the Internet.

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs
10 running on the respective computers and having a client-server relationship to each other. In some embodiments, a server transmits data, e.g., an HTML page, to a user device, e.g., for purposes of displaying data to and receiving user input from a user interacting with the device, which acts as a client. Data generated at the user device, e.g., a result of the user interaction, can be received at the server from the device.

15 While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or on the scope of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in
20 combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially be claimed as such, one or more features from a claimed combination can in some cases be excised from the
25 combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to
30 achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described

program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain cases, multitasking and parallel processing may be advantageous.

What is claimed is:

CLAIMS

1. A method performed by one or more computers, the method comprising:
 - receiving a process definition graph, the process definition graph having a plurality of task nodes that represent respective tasks to be performed by a respective robot of a plurality of robots, wherein each task node is associated with a location at which the task will be performed;
 - generating, from the process definition graph, an initial modified process definition graph that adds constraints for respective swept volumes occupied by each task represented by the plurality of task nodes; and
 - generating, from the initial modified process definition graph, a refined process definition graph, wherein the refined process definition graph includes respective motion plans for robots moving between tasks, wherein the motion plans define transitions that avoid the swept volumes occupied by each task represented by the plurality of task nodes.
2. The method of claim 1, wherein generating the refined process definition graph comprises generating motion plans for a plurality of combinations of the plurality of robots performing the plurality of tasks.
3. The method of claim 2, wherein generating the motion plans for the plurality of combinations of the plurality of robots comprises:
 - selecting an initial ordering of the plurality of tasks; and
 - iteratively generating respective motion plans for a respective robot to perform each subsequent task in the initial ordering until no additional motion plans can be generated.
4. The method of claim 3, further comprising:
 - selecting a different ordering of the plurality of tasks; and
 - iteratively generating respective motion plans for each subsequent robot in the different ordering to perform a respective task until no additional motion plans can be generated.

5. The method of any one of claims 1-4, wherein generating, from the initial modified process definition graph, a refined process definition graph occurs before any tasks have been assigned to any of the plurality of robots.
6. The method of any one of claims 1-5, wherein generating, from the initial modified process definition graph, a refined process definition graph occurs before any ordering of execution has been assigned to the plurality of tasks.
7. The method of any one of claims 1-6, further comprising generating, from the refined process definition graph, a schedule for the plurality of robots that specifies executing motion actions that avoid the swept volumes occupied by tasks represented by the plurality of task nodes in the process definition graph.
8. A system comprising one or more computers and one or more storage devices on which are stored instructions that are operable, when executed by the one or more computers, to cause the one or more computers to perform operations comprising:
 - receiving a process definition graph, the process definition graph having a plurality of task nodes that represent respective tasks to be performed by a respective robot of a plurality of robots, wherein each task node is associated with a location at which the task will be performed;
 - generating, from the process definition graph, an initial modified process definition graph that adds constraints for respective swept volumes occupied by each task represented by the plurality of task nodes; and
 - generating, from the initial modified process definition graph, a refined process definition graph, wherein the refined process definition graph includes respective motion plans for robots moving between tasks, wherein the motion plans define transitions that avoid the swept volumes occupied by each task represented by the plurality of task nodes.
9. The system of claim 8, wherein generating the refined process definition graph comprises generating motion plans for a plurality of combinations of the plurality of robots performing the plurality of tasks.
10. The system of claim 9, wherein generating the motion plans for the plurality of combinations of the plurality of robots comprises:

selecting an initial ordering of the plurality of tasks; and
iteratively generating respective motion plans for a respective robot to perform each subsequent task in the initial ordering until no additional motion plans can be generated.

11. The system of claim 10, the operations further comprising:
selecting a different ordering of the plurality of tasks; and
iteratively generating respective motion plans for each subsequent robot in the different ordering to perform a respective task until no additional motion plans can be generated.
12. The system of any one of claims 8-11, wherein generating, from the initial modified process definition graph, a refined process definition graph occurs before any tasks have been assigned to any of the plurality of robots.
13. The system of any one of claims 8-12, wherein generating, from the initial modified process definition graph, a refined process definition graph occurs before any ordering of execution has been assigned to the plurality of tasks.
14. The system of any one of claims 8-13, further comprising generating, from the refined process definition graph, a schedule for the plurality of robots that specifies executing motion actions that avoid the swept volumes occupied by tasks represented by the plurality of task nodes in the process definition graph.
15. A computer storage medium encoded with instructions that, when executed by one or more computers, cause the one or more computers to perform operations comprising:
receiving a process definition graph, the process definition graph having a plurality of task nodes that represent respective tasks to be performed by a respective robot of a plurality of robots, wherein each task node is associated with a location at which the task will be performed;
generating, from the process definition graph, an initial modified process definition graph that adds constraints for respective swept volumes occupied by each task represented by the plurality of task nodes; and

generating, from the initial modified process definition graph, a refined process definition graph, wherein the refined process definition graph includes respective motion plans for robots moving between tasks, wherein the motion plans define transitions that avoid the swept volumes occupied by each task represented by the plurality of task nodes.

16. The computer storage medium of claim 15, wherein generating the refined process definition graph comprises generating motion plans for a plurality of combinations of the plurality of robots performing the plurality of tasks.

17. The computer storage medium of claim 16, wherein generating the motion plans for the plurality of combinations of the plurality of robots comprises:

selecting an initial ordering of the plurality of tasks; and

iteratively generating respective motion plans for a respective robot to perform each subsequent task in the initial ordering until no additional motion plans can be generated.

18. The computer storage medium of claim 17, the operations further comprising:
selecting a different ordering of the plurality of tasks; and

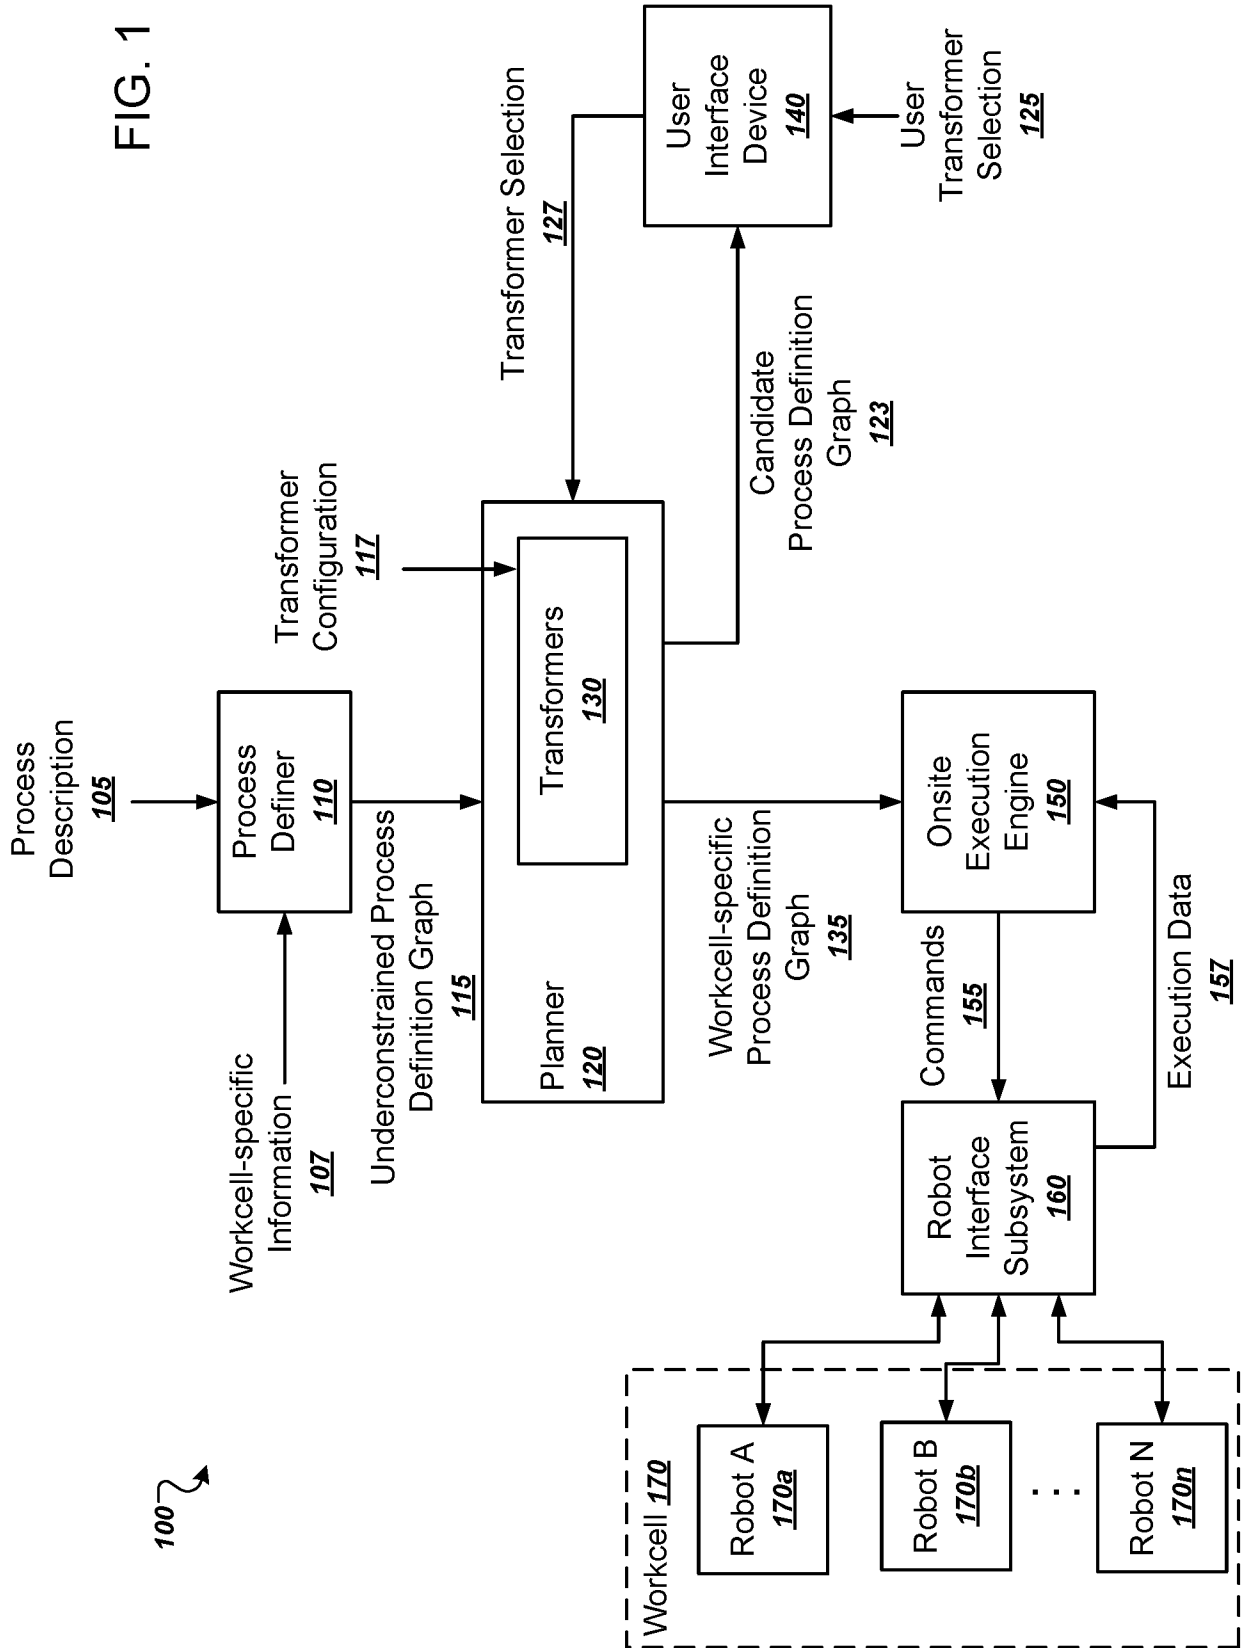
iteratively generating respective motion plans for each subsequent robot in the different ordering to perform a respective task until no additional motion plans can be generated.

19. The computer storage medium of any one of claims 15-18, wherein generating, from the initial modified process definition graph, a refined process definition graph occurs before any tasks have been assigned to any of the plurality of robots.

20. The computer storage medium of any one of claims 15-19, wherein generating, from the initial modified process definition graph, a refined process definition graph occurs before any ordering of execution has been assigned to the plurality of tasks.

21. The computer storage medium of any one of claims 15-20, further comprising generating, from the refined process definition graph, a schedule for the plurality of robots that specifies executing motion actions that avoid the swept volumes occupied by tasks represented by the plurality of task nodes in the process definition graph.

FIG. 1



2/11

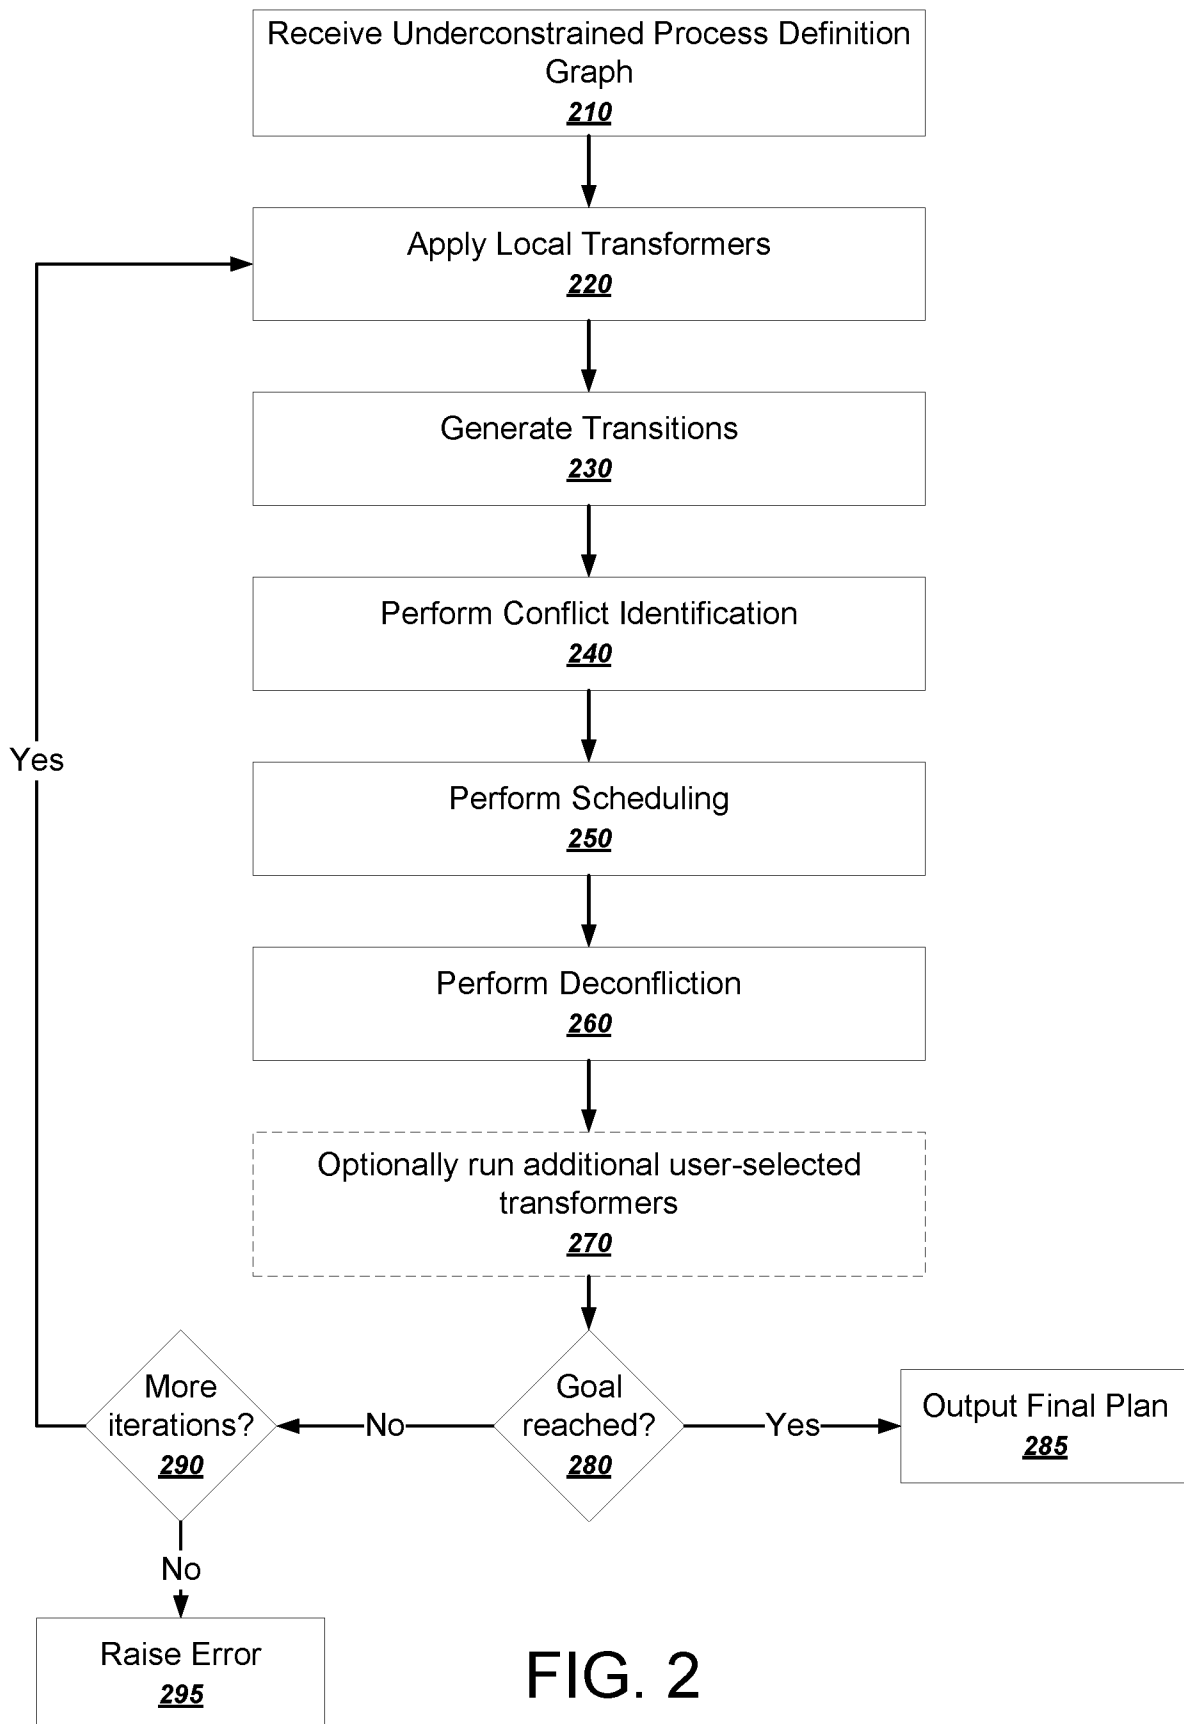


FIG. 2

3/11

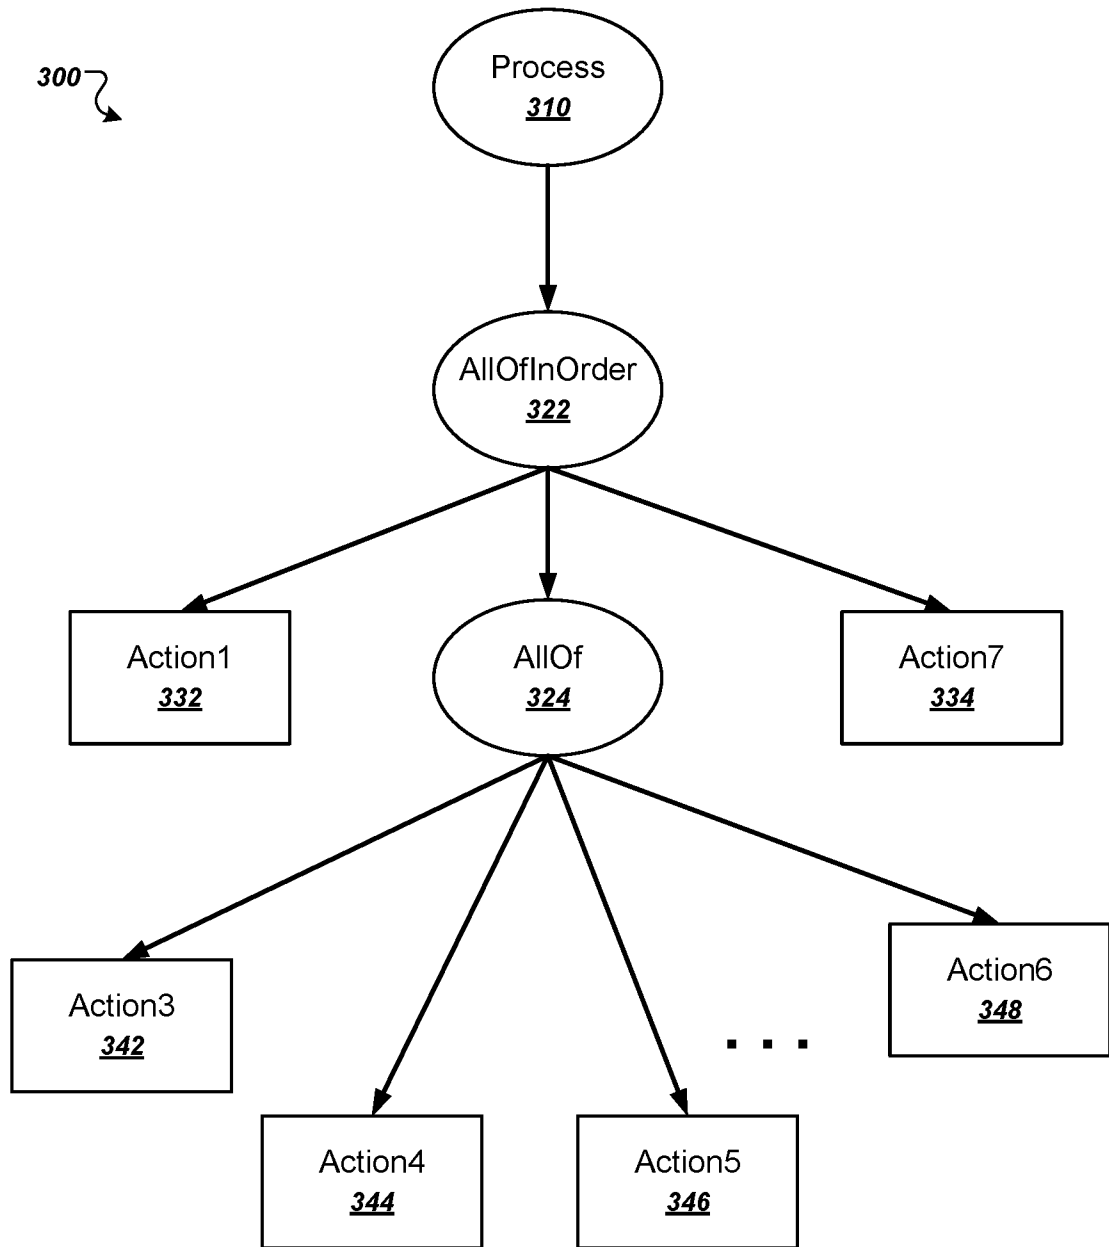


FIG. 3

4/11




FIG. 4A

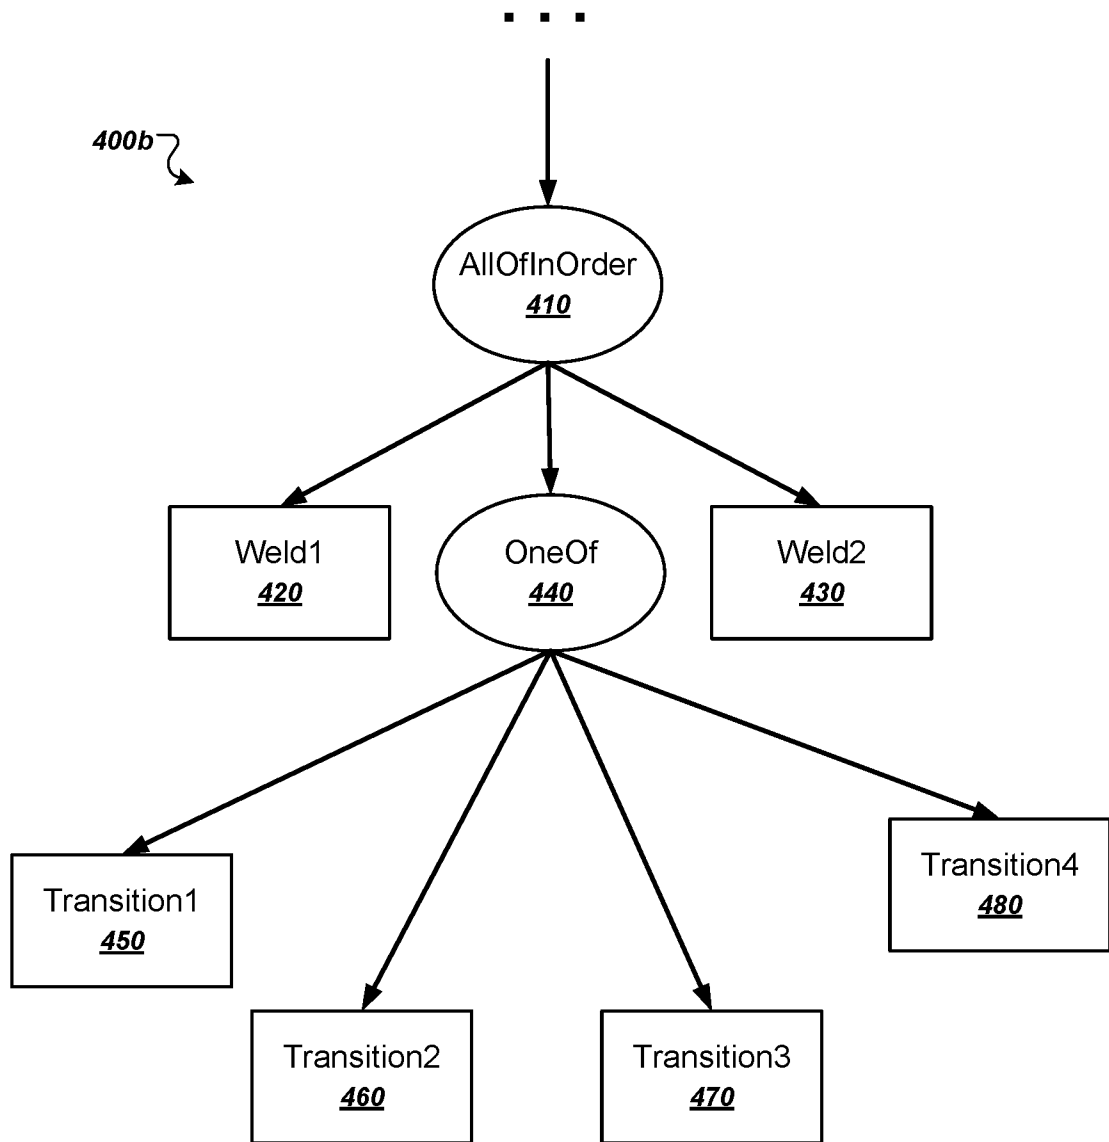


FIG. 4B




FIG. 5

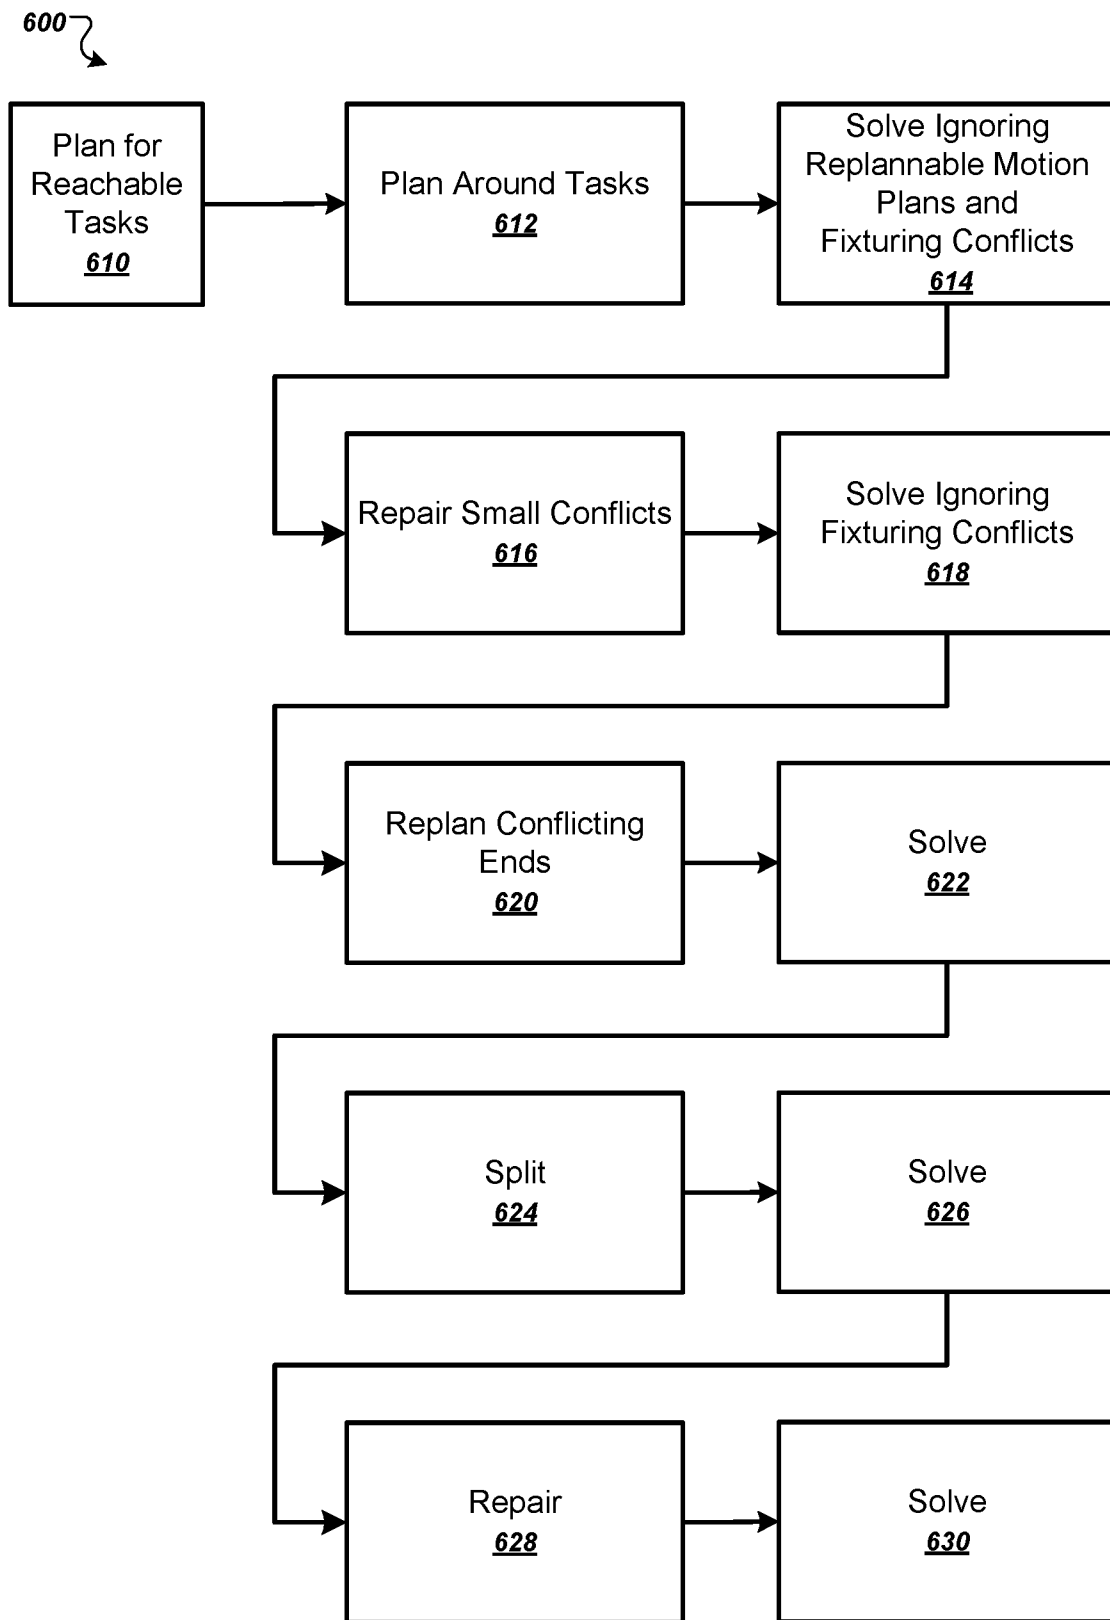


FIG. 6




FIG. 7A




FIG. 7B




FIG. 7C

11/11

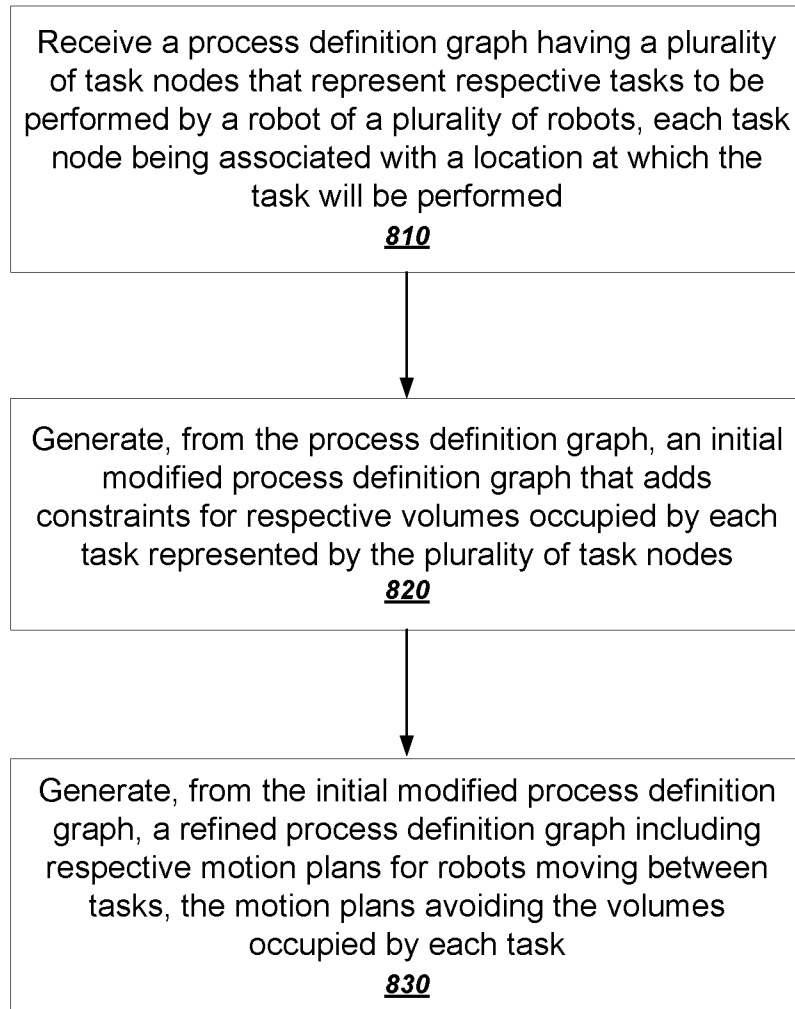


FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2020/047810

A. CLASSIFICATION OF SUBJECT MATTER
INV. B25J9/16
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
B25J
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GLIBERT P-R ET AL: "Scheduling of a multi-robot assembly cell", COMPUTER-INTEGRATED MANUFACTURING SYSTEMS, ELSEVIER SCIENCE, GB, vol. 3, no. 4, 1 November 1990 (1990-11-01), pages 236-245, XP024240052, ISSN: 0951-5240, DOI: 10.1016/0951-5240(90)90064-L [retrieved on 1990-11-01] page 241 page 243 ----- -/--	1,7,8, 14,15, 20,21

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search 19 November 2020	Date of mailing of the international search report 30/11/2020
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Lefeure, Guillaume

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2020/047810

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2017/097631 A1 (LINNELL JEFFREY [US] ET AL) 6 April 2017 (2017-04-06)	1,5-8, 12-15, 19-21
Y	paragraph [0007] paragraph [0118] paragraph [0121]	2-4, 9-11, 16-18
Y	----- WO 2018/122567 A1 (SIEMENS IND SOFTWARE LTD [IL]) 5 July 2018 (2018-07-05) paragraph [0067]	2-4, 9-11, 16-18
A	----- SHITAL S. CHIDDARWAR ET AL: "Conflict free coordinated path planning for multiple robots using a dynamic path modification sequence", ROBOTICS AND AUTONOMOUS SYSTEMS, vol. 59, no. 7-8, 1 July 2011 (2011-07-01), pages 508-518, XP055203187, ISSN: 0921-8890, DOI: 10.1016/j.robot.2011.03.006 the whole document	1-21
A	----- PELLEGRINELLI STEFANIA ET AL: "Multi-robot spot-welding cells for car-body assembly: Design and motion planning", ROBOTICS AND COMPUTER INTEGRATED MANUFACTURING, ELSEVIER SCIENCE PUBLISHERS BV., BARKING, GB, vol. 44, 21 August 2016 (2016-08-21), pages 97-116, XP029830767, ISSN: 0736-5845, DOI: 10.1016/J.RCIM.2016.08.006 the whole document	1-21
A	----- US 2019/143524 A1 (TAKAHASHI MAKOTO [JP] ET AL) 16 May 2019 (2019-05-16) the whole document -----	1-21

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2020/047810

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
US 2017097631	A1	06-04-2017	CN 107073708 A	18-08-2017
			CN 111515947 A	11-08-2020
			EP 3145682 A1	29-03-2017
			US 2015336269 A1	26-11-2015
			US 2017097631 A1	06-04-2017
			WO 2015179099 A1	26-11-2015

WO 2018122567	A1	05-07-2018	CN 110099773 A	06-08-2019
			EP 3562629 A1	06-11-2019
			US 2019344443 A1	14-11-2019
			WO 2018122567 A1	05-07-2018

US 2019143524	A1	16-05-2019	CN 109760042 A	17-05-2019
			EP 3482886 A1	15-05-2019
			JP 2019084664 A	06-06-2019
			US 2019143524 A1	16-05-2019
