

(51) International Patent Classification:
B25J 9/16 (2006.01)(21) International Application Number:
PCT/US2016/012204(22) International Filing Date:
5 January 2016 (05.01.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/107,755 26 January 2015 (26.01.2015) US

(71) Applicant: DUKE UNIVERSITY [US/US]; 2812 Erwin Road, Suite 306, Durham, North Carolina 27705 (US).

(72) Inventors: KONIDARIS, George D.; c/o Duke University, 2812 Erwin Road, Suite 306, Durham, North Carolina 27705 (US). SORIN, Daniel J.; c/o Duke University, 2812 Erwin Road, Suite 306, Durham, North Carolina 27705 (US).

(74) Agent: KNIGHT, Sarah J.; Talem IP Law, LLP, P.O. Box 358880, Gainesville, Florida 32635-8880 (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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(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))

[Continued on next page]

(54) Title: SPECIALIZED ROBOT MOTION PLANNING HARDWARE AND METHODS OF MAKING AND USING SAME




FIG. 6

(57) Abstract: Specialized robot motion planning hardware and methods of making and using same are provided. A robot-specific hardware can be designed using a tool that receives a robot description comprising a collision geometry of a robot, degrees of freedom for each joint of the robot, and joint limits of the robot; receives a scenario description; generates a probabilistic roadmap (PRM) using the robot description and the scenario description; and for each edge of the PRM, produces a collision detection unit comprising a circuit indicating all parts of obstacles that collide with that edge. The hardware is implemented as parallel collision detection units that provide collision detection results used to remove edges from the PRM that is searched to find a path to a goal position.



-
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*
- *of inventorship (Rule 4.17(iv))*
- Published:**
- *with international search report (Art. 21(3))*

SPECIALIZED ROBOT MOTION PLANNING HARDWARE AND METHODS OF MAKING AND USING SAME

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit of U.S. Provisional Application Serial No. 62/107,755, filed January 26, 2015, which is hereby incorporated by reference its entirety, including any figures, tables, and drawings.

BACKGROUND

[0002] Many of the most exciting applications of computing today are in robotics. Recent advances in robotics have greatly expanded the opportunities for robots to change the way people live. Robots are being developed for a wide spectrum of real-world applications including health care, personal assistance, general-purpose manufacturing, search-and-rescue, and military operations. Unlike traditional robotics applications in which a robot operates in a tightly controlled environment (e.g., on an assembly line where the environment is static and known *a priori*), these applications require the robot to perform *motion planning*.

[0003] Motion planning is the process of determining how a robot should move to achieve a goal. For example, a robot may wish to move its arm to reach a desired object without colliding with itself or any other object. The ability to perform motion planning to navigate and manipulate objects is critical to dealing with natural environments that are not heavily controlled or engineered. The robot first perceives information about the physical world around it, and then the computer determines how to control the physical robot so that the robot can navigate in this perceived environment.

[0004] In many situations, a major constraint on the usefulness of robots is their ability to perform real-time motion planning in environments with obstacles. Motion planning algorithms tend to be computationally intensive, requiring a vast number of collision checks for each movement. Existing techniques for collision detection are generally implemented in software running on one or more general-purpose processors and/or graphics processing units (GPUs). However, running collision detection software on general-purpose processors is insufficiently fast for real-time motion planning and/or is too power-hungry (if using GPUs) for many target applications.

BRIEF SUMMARY

[0005] Specialized robot motion planning hardware and methods of making and using same are provided. The described hardware can address the challenge of real-time motion planning. The described systems and methods can take a specific robot and produce a low-power, low-cost hardware architecture capable of performing real-time motion planning for that robot.

[0006] The robot motion planning hardware described herein can be designed using a software tool that uses a received robot description for a specific robot and a received scenario description to generate a probabilistic roadmap (PRM). Then, for each edge of the PRM, the software tool produces a design for a collision detection unit for that edge. The collision detection unit is implemented as a circuit that computes whether the motion corresponding to that edge collides with obstacle points or surfaces. The specific robot is defined by its collision geometry, degrees of freedom for each of its joints, and its joint limits; and such information is provided in the robot description. The scenario description can include parameters for a fixed scene, a distribution of obstacles, a distribution of starting robot poses, and a distribution of goal poses.

[0007] For a specific robot, a motion planning hardware can be provided that includes a plurality of collision detection units configured to process an input in parallel, each collision detection unit having a circuit for its corresponding edge of a probabilistic roadmap. The circuit can include Boolean logic gates and/or lookup tables. To perform robot motion planning, the robot receives, at its motion planning hardware, data of an environment captured from a sensor system of the robot and data of a goal position, and outputs at each collision detection unit of the motion planning hardware, an indication as to whether an edge of a probabilistic roadmap corresponding to the collision detection unit has a collision with a part of an obstacle identified in the data of the environment. The output of the motion planning hardware is used to remove edges of the probabilistic roadmap that are likely to result in collision in the environment. Then, a graph search of the probabilistic roadmap can be performed to identify a shortest path between the current position of the robot and the goal position along remaining edges of the probabilistic roadmap.

[0008] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Many aspects of the disclosure can be better understood with reference to the following drawings. While several implementations and examples are described in connection

with these drawings, the disclosure is not limited to the implementations and examples disclosed herein.

[0010] **Figures 1A-1C** illustrate motion planning using a probabilistic roadmap.

[0011] **Figure 2** illustrates a swept volume created by a movement from a first position to a second position.

[0012] **Figure 3** illustrates an exemplary motion planning process.

[0013] **Figure 4** illustrates a process for generating a motion planning hardware design that may be carried out by a software tool executing on a computing device.

[0014] **Figure 5** illustrates a high level overview of a robot motion planning hardware architecture.

[0015] **Figure 6** is a high level schematic of an example motion planning hardware.

[0016] **Figure 7** is the symbolic text for an example edge configuration generated to feed to the circuit generating software.

[0017] **Figure 8** illustrates a portion of an example generated circuit for robot motion planning hardware.

DETAILED DESCRIPTION

[0018] Specialized robot motion planning hardware and methods of making and using same are provided. The described hardware can address the challenge of real-time motion planning. The described systems and methods can take a specific robot and produce a low-power, low-cost hardware architecture capable of performing real-time motion planning for that robot.

[0019] Figures 1A-1C illustrate motion planning using a probabilistic roadmap. As mentioned in the background, motion planning refers to the processes used to identify the motions a robot will undergo in order to move from one position to another position until the robot achieves a desired position. Since an articulated robot is formed of a set of rigid bodies connected by joints, motion planning determines the movement from one pose to another pose until a goal pose is met.

[0020] The pose of an articulated robot can be described by assigning positions to all of its joints. The space of all such joint positions is called the robot's *configuration space* (often called *C-space*), denoted as Q . Some configurations (of the joint positions) in Q would result in a collision with an obstacle in the robot's environment or the robot's own body; the remainder of Q is called *free space*.

[0021] The robot motion-planning problem consists of finding a path through free space from some start configuration q_0 to any configuration within a set of goal configurations, G , such as shown in Figure 1A. Referring to Figure 1A, a C-space **100** can include obstacle collision

regions **102**, **104** and free space **110**. A path **120** can be found through the free space **110** from a start configuration **122** to a configuration within the set of goal configurations **124**.

[0022] A common approach to solving motion planning problems is to build a *probabilistic roadmap* (PRM), such as described by Kavraki *et al.* in “Probabilistic roadmaps for path planning in high-dimensional configuration spaces,” (IEEE Transactions on Robotics and Automation, vol. 12, no. 4, pp. 566-580, Aug. 1996), which is hereby incorporated by reference in its entirety. A PRM consists of a graph where each node is a point in C-space, and a pair of points is connected if a direct movement between them is possible without a collision, such as shown in Figure 1B. Referring to Figure 1B, a PRM of nodes **130** and edges **140** (the lines connecting two nodes) is shown in the free space **110** of configuration space **100**.

[0023] Given a (q_0, G) pair, a motion plan can be found by attempting to find a collision-free connection from each node to q_0 and G , and then finding a path in the graph between them. This is shown in Figure 1C, where a path **150** from q_0 to G can be found that will be collision-free.

[0024] As mentioned above, a PRM is constructed by determining whether a straight-line path between two points in configuration space (representing the direct movement of each of the robot’s joints from one position to another) collides with an obstacle, referred to as “collision detection”. Collision detection can be carried out by building a geometric model of the *swept volume* of each robot movement. A swept volume refers to the space (in three dimensions) through which a robot moves when moving from one position to another position.

[0025] Figure 2 illustrates a swept volume created by a movement from a first position to a second position. In the example movement, an arm with a first rigid body **201** connected to a base **202** at a first joint **203** and a second rigid body **204** connected to the first rigid body **201** by a second joint **205** forms a volume, referred to as a “swept volume” **210**, as the arm moves between position 1 and position 2.

[0026] Figure 3 illustrates an exemplary motion planning process. Referring to Figure 3, the entire motion planning process **300** has two parts: an offline portion and a real-time portion. The initial, offline portion, begins with creating a PRM (**301**). This step can be carried out before a robot’s motion planning hardware is designed, but is specific to a particular robot. To create the PRM in operation **301**, nodes are sampled from the robot’s C-space and edges are added between sampled nodes if the edges are determined to not cause self-collisions. That is, an edge is not included in the initial PRM graph if, for example, that movement would cause the arm to collide with the robot when the robot moves from the configuration represented by one node to the configuration represented by the other node.

[0027] The PRM, which can be stored in the form of a graph data structure, can optionally be pruned (or “trimmed”) to remove certain edges based on some utility function. The initial (and optionally pruned) PRM has edges only in free space for a world without obstacles other than the robot itself. This PRM can be stored in a storage system of the robot for use during motion planning. This PRM can be used in the design of the motion planning hardware for that particular robot (described in more detail below).

[0028] The PRM created offline (and which can be stored on a storage system of the robot) is then used in real-time, where edges that are not in free space due to obstacles in the perceived environment are disabled (302). During operation of the robot, the robot analyzes and transmits the robot’s environment to the motion planning hardware. It should be understood that the environment captured and analyzed by the robot is not necessarily that only in a direct path between the robot’s current position and a goal position. A robot may have a valid path only in a circuitous route (and that path may not even be in a total line of sight of the goal pose). The robot can include or communicate with any of a variety of sensor systems through which an environment can be detected and analyzed. For example, the robot can have one or more camera sensors capable of receiving red, green, blue, and depth information (e.g., via infrared). As another example, cameras (or other sensors) may be located at one or more positions in the environment such as on walls, ceilings, floors and/or objects in the environment (like on a table, box, etc.). Environmental information captured from the sensor(s) in the environment can be transmitted to the robot via wired or wireless means.

[0029] The data received from the sensor inputs can be analyzed to identify objects in the environment, which can be treated as obstacles. The object identification can be included in the environmental information provided to the motion planning hardware. In some cases, objects are represented by the values of a voxel (e.g., 3D pixel). In another case, objects detected in the environment can be represented as mesh models made out of triangles, where each vertex of a triangle is described with its 3D coordinates. Of course, other representations may be used in certain implementations, such as a probabilistic occupancy grid or an octomap.

[0030] Once the environmental information is provided to the motion planning hardware, the motion planning hardware can begin disabling the edges of the PRM that would result in a collision with an obstacle by first detecting collisions of a part of an obstacle (represented in some manner by the robot) with all robot movements in the PRM. For example, an image capture of an environment can identify a lamp (or other object) in the robot’s environment. The 3D coordinates for the lamp (or at least its surfaces) can be fed into the hardware for parallel analysis.

[0031] In the hardware, a collision detection unit is provided for each edge of the PRM. Each collision detection unit can consider the swept volume as having a *single Boolean function* that determines whether the swept volume collides with a particular point in 3D space. That is, each collision detection unit can consist of the logic for one Boolean function, $F(\text{object_edge})$, that evaluates to true if a point on the object's surface (as perceived by the sensor system) collides with that robot movement. An example symbolic text of this Boolean function is shown in Figure 7. The Boolean function has a number of inputs equal to the number of bits used to describe a point in 3D space. In the example implementation, 15 bits are used to describe a point in 3D space (with 5 bits per axis) and therefore 15 inputs are provided to the circuitry for performing the Boolean function.

[0032] For typical environments and applications, the coordinates describing a point in 3D space do not have to be specified extremely finely. Robotic motion planning has less stringent requirements than physics simulations that determine whether two objects will collide; the physics simulation needs to know exactly whether the objects collide, whereas robotic motion planning simply needs to avoid collisions, even if this is done conservatively (e.g., eliminating a path out of conservatism, even if that path might potentially be collision-free). By being conservative in predicting when two objects will collide, calculations may be performed with few numbers of bits. For example, calculations can be performed with 5-bit fixed-point arithmetic instead of far more expensive 32-bit floating point arithmetic. To illustrate how the lower resolution remains sufficient for typical environments, consider a robot with arms that can navigate around a 1 m^3 box. With only, 5 bits to describe a coordinate, a point's position in each dimension can still be resolved down to $1 \text{ m}/2^5$, which is about 3.1 cm.

[0033] Once collision detection has been performed, the edges of the PRM that would result in a collision if the robot were to move between the node pairs connected by those edges can be disabled using the outcome of collision detection. For example, the output of each collision detection unit can be a single bit, a flag, or other indicia as to whether there is or is not a collision between the edge represented by that collision detection unit and a portion of an obstacle. The output of the motion planning hardware is used to remove edges of the probabilistic roadmap that are likely to result in collision in the environment.

[0034] By removing edges likely to cause a collision, the remaining edges can be used to determine a legal motion plan (304). For example, the robot can determine a legal motion plan by identifying an edge of a graph data structure of the probabilistic roadmap as being valid if the circuit corresponding to that edge indicates that there is no collision. Then, a graph search can be performed on the graph. The motion planning hardware processor or the software running on the

robot can select a path using the valid edges (indicated as collision-free from the output of the collision detection units). For example, Dijkstra's shortest path algorithm can be used, or the Bellman-Ford algorithm can be used.

[0035] As illustrated in Figure 3, a portion of the motion planning process **300** is preliminarily carried out offline so that decisions with respect to motion planning can be made real-time while a robot is operating in an environment. This real-time decision making capability for motion planning is due to the use of the special robot motion planning hardware.

[0036] Figure 4 illustrates a process for generating a motion planning hardware design that may be carried out by a software tool executing on a computing device (such as a general purpose computer). The robot motion planning hardware described herein can be designed using a software tool that uses a received robot description for a specific robot **401** and a received scenario description **402** to generate a PRM (**403**). Then, for each edge of the PRM, the software tool produces a design for a collision detection unit for that edge (**404**). An example design produced for a collision detection unit for an edge is shown in Figure 7. The collision detection unit is a circuit, composed of logic gates and/or lookup tables, that detects a collision between an edge in the PRM (corresponding to a swept volume) and a part of an obstacle (expressed as, for example, a 3D point in space) present in the environment. The software tool can output a circuit design file **405** describing the circuitry of the collision detection units. The circuit design file **405** may be embodied in a variety of file formats depending on the implementation so that the file is suitable for being used to program a field programmable gate array (FPGA), for being used directly or indirectly in a layout software for integrated circuit fabrication, or for being used in conjunction with other design software used to create a hardware processor.

[0037] Through the two inputs of robot description **401** and scenario **402**, the motion planning hardware design takes into consideration, and can accommodate, a wide range of possible robots and classes of environments.

[0038] The specific robot is defined by its collision geometry, degrees of freedom for each of its joints, and its joint limits; and such information is provided in the robot description **401**. In some cases, the robot description **401** can be provided in the form of an industry-standard Unified Robotic Description Format (URDF) file. The URDF file specifies the collision geometry of the robot, its display geometry (these are often different for the purposes of visualization), and the degrees of freedom for each joint, along with joint limits, and can almost always be obtained from the robot manufacturer. Of course, other standardized forms now or later known may be used.

[0039] Because the motion planning hardware is designed for a specific robot, an efficient processor can be achieved. Some of this efficiency arises because motion planning hardware does

not need to perform general motion planning for an arbitrary robot. For example, if a scenario involves a robot with two arms that are each 3 feet long, the processor needs to plan motions for just those two specific arms and does not need to consider objects farther than 3 feet from the robot. Furthermore, by tailoring the processor to a specific robot, many expensive (e.g., high power usage or slow speed) calculations can be performed offline, prior even to processor design. For example, results of equations that depend solely on the robot itself (e.g., calculations to determine whether moving one robot arm will cause a collision with the other robot arm) can be pre-computed and the results of these pre-computations embedded in the processor. The ability to perform these calculations offline avoids having to perform these time-intensive and/or processing intensive computations in the field.

[0040] The scenario description **402** can include parameters for a fixed scene (e.g., floor, wall, workspace size, a desk at a certain height), a distribution of obstacles (locations, sizes, and shapes), a distribution of starting robot poses, and a distribution of goal poses. As used herein, “scenario” refers to a class of environments, such as “desk with objects on it” or “bookshelf with objects on it”. A scenario includes distributions of the number of obstacles and their locations, sizes, and shapes. A scenario also includes distributions of robot starting poses and goal poses with positions of various robot parts. At any given time, the robot needs to perform motion planning *for one task drawn from one scenario*. For example, if the scenario is “3-foot high desk with objects on it”—with given distributions of the number of obstacles, obstacle locations, obstacle sizes/shapes, and robot start and goal poses—it is expected that the robot does not need to be able to move its arm beneath the table. As another example, in many scenarios, it is unlikely that the robot needs to be able to move its arm behind itself. This a priori knowledge of the scenario conserves resources (and dedicates more hardware) for solving the expected motion planning problems. If more than one scenario is possible in a given environment, then a composite scenario that is effectively the union of the possible scenarios is utilized or the robot motion planning hardware is reconfigured when the scenario changes.

[0041] The software tool can perform certain optimizations to output a design that is compact enough to fit, for example, on a FPGA, but is also highly likely to produce plans for problems drawn from the scenario. It should be understood that implementations of the robot motion planning hardware can be implemented as an application-specific integrated circuit (ASIC), a FPGA, or other circuitry.

[0042] For example, the software tool can optionally trim edges from the PRM generated in step **403** to reduce the total number of edges in the PRM. This optional process (**406**) can include creating problem instances and trimming the edges according to a contribution factor. The

contribution factor is the extent to which an edge in the initial PRM contributes to finding the minimum solution path for a problem instance. If an edge is not frequently or ever involved in the minimum solution path, then it could be removed.

[0043] In some cases, the scenario description from the scenario description **402** is used to generate a plurality of concrete problem instances. The number of problem instances can be selected based on the tolerance for failure to find a path in a PRM. For example, tens, hundreds, thousands or more of problem instances may be generated. A problem instance refers to a scene/obstacle/start/goal tuple generated from the scenario description in the scenario description **402**. The problem instance indicates a motion planning problem that satisfies the constraints of a given scenario. For each problem instance, that instance is evaluated to determine the extent to which each edge in the PRM contributes to finding the minimum solution path (e.g., the usefulness of the edge defined by the estimated probability that the edge is necessary for solving a problem drawn from the scenario). In effect, the software tool runs a simulation of the motion planning hardware using an environment/scenario described by each problem instance.

[0044] In one case, the trimming of the edges can be based on the average contribution of an edge so that those edges that meet or exceed a threshold are kept and those that don't are removed. In another case, the edges can be ranked by average contribution and only the top ten, hundred, thousand or other number of edges are kept (the amount of the edge budget for the processor). In many cases, the least "useful" edges can be iteratively removed until a total edge budget is reached. Available algorithms that sparsify and prune roadmaps without any background task knowledge can also be incorporated. Although the trimming process can be very computationally expensive, it need only occur once per robot/scenario pair, and is performed offline.

[0045] If the points in C-space are sampled randomly, the use of PRM is probabilistically complete. As the algorithm samples more points in C-space, the probability of finding a path solution tends to 1. However, by fixing the PRM in the processor design, particularly when given a fixed budget of hardware and PRM edges, it is possible to encounter a situation where no path solution can be found. To minimize this occurrence when using a fixed PRM processor design, the processes that generate the PRM (e.g., the software tool) can incorporate simulations of sample motion planning problems that satisfy the constraints of a given scenario. In addition, scenario-specific sampling strategies can be incorporated. For example, samples that are expected to be helpful in a particular scenario (such as indicating to sample right above a desk for a task that may involve an action on or near a desk) can be specifically incorporated. These simulations are described with respect to optional process **406**.

[0046] Another optimal optimization can include reducing the number of gates implementing a collision detection unit (407). For example, Boolean minimization can be used.

[0047] In one implementation, the Klampt open-source software for modeling robotic motion is used to produce a robot- and scene-specific PRM (in operation 403). This PRM can be optimized to fit on an FPGA using the optimizing steps 406 and/or 407. As mentioned above, once the PRM has been trimmed (e.g., in process 406), the trimmed PRM can be used to produce a set of collision detection unit circuits (as described in operation 404). The collision detection unit circuits can then be optimized to reduce the number of gates (407) using Boolean minimization (e.g., with an ESPRESSO logic minimizer). The FPGA can be configured (and reconfigured) for a desired robot/scenario pair based on the design produced by the software tool.

[0048] There are a number of additional optimizations that can be carried out during the time of steps 406 and/or 407. For example, strategic sampling of a robot's C-space can be included to further optimize design. Typical algorithms choose points at random so as to provide probabilistic guarantees about being able to find paths. However, certain implementations strategically sample points in the C-space to minimize the hardware necessary to perform collision detection and/or take advantage of a specific robot's expected scenario. Given a fixed hardware and power budget, a larger PRM (with more samples and thus more motion options) can be represented in the hardware if the collision detection units are smaller and consume less power. To reduce the size and power of collision detection units, it was observed that not all robot motions will have equal-sized circuits for determining collisions with objects/obstacles. Therefore, it is possible to choose points in C-space so as to obtain swept volumes whose circuits are more amenable to local logic minimization (i.e., require smaller circuits). Furthermore, it is possible to choose points in C-space so as to facilitate global logic optimization across collision detection units. As in the example illustrated in Figure 7 below, 15 bits can be used to describe the point in 3D space, which implies a truth table with 2^{15} entries. Since most of these truth tables will be sparse, the Boolean function can be efficiently computed (e.g., using bounding boxes for the swept volume) and can be expected to result in relatively small circuits (when optimized using, for example the ESPRESSO logic minimizer). Indeed, experimentally, it was found that most truth tables can be optimized down to only dozens of logic gates.

[0049] Beyond locally optimizing the logic for each collision detection unit (e.g., by using ESPRESSO), it is possible to globally optimize *across* collision detection units. Many movements will have truth tables that are similar, and this similarity can be exploited to optimize globally. Preliminary experiments with ESPRESSO and a very simple robot arm in 3-space found that many common sub-expressions can be factored out for global optimization. Although the experimental

configuration is limited to optimizing across 16 collision detection units (due to the ESPRESSO logic minimizer used in the experiments), global optimization can be carried out across any number of collision detection units in practice

[0050] Another opportunity for strategic sampling (and thereby optimization) targets the case where the scenarios the robot will face are known, in the form of distributions over start points, goals, and obstacles. For example, in many scenarios, the robot is much more likely to be asked to generate a plan to position its gripper in the workspace in front of it than in the area just above and behind its own shoulder. Similarly, in many scenarios, obstacles are much more likely to appear in the workspace area, or on the far sides of the workspace area, than directly in front of the robot's elbows. In some scenarios, obstacles are likely to be small, and in others they are likely to be large. This kind of coarse prior information about scenarios is easy to describe in many real-world applications. With this scenario information, we can place our samples adaptively to, for example, provide sparser coverage of configurations that have the robot's arm behind its shoulder.

[0051] Figure 5 illustrates a high level overview of a robot motion planning hardware architecture. Referring to Figure 5, the high level architecture **500** for motion planning involves robot specifications **501**, motion planning goals **502**, obstacle data **503**, collision detection results **504**, and motion plans **505**. In particular, a robot motion planning hardware **510** uses robot specifications **501**, motion planning goals **502**, and obstacle data **503** to generate collision detection results **504** and motion plans **505**. As discussed above, the robot specifications **501** (and even motion planning goals **502**) can be used offline to generate the hardware used in real-time motion planning. The output of the robot motion planning hardware **510** can be provided to the robot via an interface **520** so that the robot can act on the motion plans **505**.

[0052] Figure 6 is a high level schematic of an example motion planning hardware. A robot motion planning hardware can be a special-purpose motion planning processor **600**. The motion planning processor includes a collision detection unit (e.g., **601**, **602**, **603** ..., **60X**) for every possible motion of the robot (or at least every motion in the PRM produced by the software tool) and a path computation unit **610**.

[0053] The processor's inputs describe the obstacles **611** in the environment and the path goals (motion planning goals) **612**. The obstacle information **611** is received by the collision detection units. In a preferred embodiment, all of the collision detection units run in parallel. However, in some implementations, the collision detection units may be time-multiplexed to serve multiple robot motions. The motion planning goals **612** and the output of the collision detection units **601**, **602**, **603** ..., **60X** are provided to the path computation unit **610**, which performs a

graph search to identify the path used to achieve a goal. The processor's outputs are the path **613** and, optionally, collision detection information.

[0054] Referring to Figure 6, for a specific robot, a motion planning hardware can be provided that includes a plurality of collision detection units configured to process an input in parallel, each collision detection unit having a circuit, with logic gates and/or lookup tables, for its corresponding edge of a probabilistic roadmap. To perform robot motion planning, the robot receives, at its motion planning hardware, data of a goal position and data of an environment captured from a sensor system of the robot, and outputs at each collision detection unit of the motion planning hardware, an indication as to whether an edge of a probabilistic roadmap corresponding to the collision detection unit has a collision with a part of an obstacle identified in the data received. The output of the motion planning hardware is used to remove edges of the probabilistic roadmap that are likely to result in collision in the environment. Then, a graph search of the probabilistic roadmap can be performed to identify a shortest path along remaining edges of the probabilistic roadmap.

[0055] Additional inputs to the collision detection circuits may be included, for example to handle certain robot motions such as, but not limited to, grasping, gripping, pushing, moving, and a combination thereof. Depending on the robot specifications, certain actions may be carried out via various parts of the robot body. The possible robot motions can be modeled and incorporated into the hardware design. For example, for a case where a specific robot can grasp an object and move with the object (such as in the common "pick-and-place" task), the object in the robot's grasp may be modeled very coarsely as a cube. Then, the processor can be configured to activate one of a plurality of object sizes contained in each hand. This would add a small number of inputs to the collision detection circuit, but greatly enhance the flexibility of the planner.

[0056] Most of the computational effort required for motion planning involves detecting whether a motion will lead to a collision. As previously mentioned, collision detection is computationally expensive. However, by performing certain processes in parallel using the described motion planning hardware, it is possible to simultaneously detect collisions for a large number of robotic motions against a large number of obstacles.

[0057] Advantageously, the motion planning hardware does not require the general-purpose hardware that is found in general-purpose CPUs or GPUs because the motion planning hardware does not need to fetch and decode instructions, perform IEEE-754-compliant floating point arithmetic, or a wide variety of other features that consume hardware and energy in general-purpose processors. Indeed, since the precision needed to achieve motion planning goals is known

a priori, the hardware can be tailored to perform operations at exactly that granularity, even if it happens to be an unusual width (e.g., 9-bit fixed point).

[0058] Figure 7 is the symbolic text for an example edge configuration generated to feed to the circuit generating software. In Figure 7, the logic for a single edge is shown. In this example case, a part of an obstacle (expressed as a 3D point in space) found to have a likelihood of colliding with that edge is represented by five bits for the x-axis, five bits for the y-axis, and five bits for the z axis for a total of 15 bits. Each bit has two possible values: one represented by the letter (e.g., A) and the other represented by the not-letter (e.g., !A). There are 70 collisions reflected in the symbolic text. The circuit generating software then uses AND, OR, NOT, NAND, and NOR gates and/or lookup tables to implement the symbolic text. In some cases, the number of gates can be optimized (reduced) by transforming some or all of the gates into NOR or NAND logic.

[0059] Figure 8 illustrates a portion of an example generated circuit for robot motion planning hardware. Referring to Figure 8, a motion planning hardware edge logic, such as the edge logic for Edge-1 **800** can include a Boolean circuit **810** with logic gates **812** representing each of the parts of an obstacle that a single edge (swept volume) would collide with and OR gate(s) **814** ORing the Boolean logic representations of the parts of obstacles, such as described by the symbolic text of Figure 7, to output a single bit. A flip-flop **820** may be included to store the value of the single-bit output of the Boolean circuit **810** so that the appropriate value can be read at the output when needed. The edge logic is provided for each edge (e.g., edge logic for Edge-1 **800** and edge logic for Edge-2 **830**) so that when an input representing a part of an obstacle is received, the motion planning hardware edge logic for all edges can process that input in parallel.

[0060] Figure 9 shows a block diagram of an example robot that can incorporate the robot motion planning hardware described herein. Referring to Figure 9, a robot **900** can include a sensor system **910**, an actuator system **920**, a processing system **930**, a storage system **940**, and robot motion planning hardware **950**. Although not shown, robot **900** can include a power system for supplying power, as needed to the various components. Batteries may be used or the power system may include wired (e.g., an ac adaptor and plug) and wireless means (e.g., using transducers and charge storage devices) for supplying power. It should be understood that robot **900** can include other systems and components not shown in the figure and that such systems and components have been omitted from the drawing for simplicity.

[0061] The sensor system **910** can be any suitable sensor system through which an environment can be detected and analyzed. Example sensor systems include, but are not limited to, optical, microwave, acoustic and other sensors in the form of video cameras, image sensors,

infrared sensors, ultrasound transceivers, microphones, force sensors, capacitive sensors, accelerometers, gyroscopes, and temperature sensors. Sensor system **910** can include hardware for some processing of a signal from a sensor and/or control of that sensor. Sensor system **910** can include components that may be local to the robot and components that may be remote from the robot. For example, in some cases sensor system **910** may include sensors that can be located in an environment, the output of which may be analyzed locally at the robot or analyzed remote from the robot and communicated to the robot for use in motion planning (among other applications that can be available). In some cases, one or more of the sensors of sensor system **910** are positioned on the robot itself.

[0062] Actuator system **920** can include actuators and motors for moving parts of the robot. Actuator system **920** can involve pneumatic, electric, thermal, magnetic or mechanical actuation as some examples.

[0063] Processing system **930** can include one or more processors to transform or manipulate data according to the instructions of software stored on the storage system **940**. Examples of processors of the processing system **930** include general purpose central processing units, application specific processors, and logic devices, as well as any other type of processing device, combinations, or variations thereof. The processing system **930** may be, or is included in, a system-on-chip (SoC) along with one or more other components such as network connectivity components, sensors, video display components, and the like.

[0064] Storage system **940** may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. Examples of storage media of storage system **940** include random access memory, read only memory, magnetic disks, optical disks, CDs, DVDs, flash memory, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other suitable storage media. In no case is the storage medium (or media) a propagated signal or carrier wave. Storage system **940** may be implemented as a single storage device but may also be implemented across multiple storage devices or sub-systems co-located or distributed relative to each other. Storage system **940** may include additional elements, such as a controller, capable of communicating with processing system **930**.

[0065] Various software applications and data may be stored on the storage system **940**. For example, the robot operating system can be stored on the storage system **940**. In some cases, one or more probabilistic roadmaps can be stored in a graph data structure on the storage system **940**.

[0066] Robot motion planning hardware 950 is a special purpose hardware based on the specifications of robot 900. The motion planning hardware 950 can be designed and embodied, for example, as described with respect to Figures 4-6. In some cases, robot motion planning hardware 950 can be embodied on a FPGA or as an ASIC. In some cases, the motion planning hardware 950 can include sub-systems of the storage system 940.

[0067] Example Implementation

[0068] A motion planning processor was designed for a 6-degree JACO-2 arm from Kinova. The sensor system for the robot involves multiple Kinect-2 modules to enable a visualization of the environment. For the example implementation, a motion planning processor was designed specifically for the control of the arm and connected to the robot through an I/O port of the robot. The Robot Operating System (ROS) software ecosystem, available from the Open Source Robotics Foundation, was used for controlling the robot and interfacing to the motion planning processor. The ROS ecosystem is used to interpret objects in the environment from the output of the Kinect-2 modules and the representations generated by the ROS software running on the robot are communicated to the motion planning processor.

[0069] With local and/or global optimization (such as described with respect to Figure 4), it was found that on the order of 5000 collision detection units were possible to fit on a high capacity FPGA. In addition, it was possible to stream 3D points (identifying parts of obstacles) taken directly from the Kinect-2 modules (without much further processing) into the processor, further optimizing the hardware.

[0070] Example alternative implementation

[0071] For a mesh model/triangle implementation, since each triangle has three vertices, each vertex can be described with its 3D coordinates.

[0072] In a case where collision detection units include circuitry for computing triangle collisions, a robot can stream object triangles to the motion planning hardware and collision detection units configured to detect whether robotic movements would cause collisions with these object triangles. In this example implementation, each collision detection unit is responsible for one edge in the graph, which represents one possible robot movement. The robot movement corresponds to a volume in 3D-space that is swept in the course of performing that movement. Here, the collision detection unit can represent the swept volume as a hull consisting of triangles and can store nine k -bit numbers for each triangle. Then, to detect collisions with an incoming object triangle, the collision detection unit performs several computations to determine whether triangles intersect each other. For example, the collision detection unit can compare each object triangle edge against each robot triangle and compare each robot triangle edge against each object

triangle. Each comparison between a robot triangle and an object triangle involves a significant amount of computation, including dot-products and cross-products. For a case where 24 dot products and 7 cross products are performed, the hardware uses 114 multipliers and 69 adders. This implementation can be optimized to reduce the computation to a single comparison (one-direction instead of the two direction comparison) by imposing constraints on triangles—a minimum length of robot triangle edges and a maximum length of obstacle edges.

[0073] It is estimated that about 100 of these collision detection units with swept volume mesh triangle storage, multipliers, and adders can be fit to a modern high-capacity FPGA. Since this number is far fewer than the number of edges in the graph, the large-size collision detection units are suitable for robots having more processing capability than found in modern high-capacity FPGAs (or in environments that can use a roadmap with fewer than 100 edges). Where time-multiplexing is used to increase the number of such collision detection units, memory bandwidth would need to be sufficient to handle the context switch between states of the edges so that the new edge's triangle state can be read in fast enough.

[0074] It should be understood that the examples and embodiments described herein are for illustrative purposes only and that various modifications or changes in light thereof will be suggested to persons skilled in the art and are to be included within the spirit and purview of this application. One skilled in the art will readily appreciate that the present invention is well adapted to carry out the objects and obtain the ends and advantages mentioned, as well as those inherent therein. The present examples along with the methods described herein are presently representative of preferred embodiments, are exemplary, and are not intended as limitations on the scope of the invention. Changes therein and other uses will occur to those skilled in the art which are encompassed within the spirit of the invention as defined by the scope of the claims.

CLAIMS

What is claimed is:

1. One or more computer readable storage media having instructions stored thereon that, when executed by a computing system, direct the computing system to at least:
 - receive a robot description comprising a collision geometry of a robot, degrees of freedom for each joint of the robot, and joint limits of the robot;
 - receive a scenario description;
 - generate a probabilistic roadmap (PRM) using the robot description and the scenario description; and
 - for each edge of the PRM, produce a collision detection unit comprising a circuit indicating all parts of obstacles that collide with that edge.
2. The media of claim 1, wherein the instructions to produce the collision detection unit comprising the circuit direct the computing system to at least:
 - identify all possible edges in the PRM that can collide with each part of an obstacle; and
 - generate an OR circuit, for each edge, ORing representations of the parts of obstacles to which that edge can collide.
3. The media of claim 2, further comprising instructions that direct the computing system to: optimize the OR circuit to reduce number of gates forming the OR circuit.
4. The media of claim 1, wherein the scenario description comprises: a fixed scene, a distribution of obstacles, a distribution of starting robot poses, and a distribution of goal poses.
5. The media of claim 4, further comprising instructions that direct the computing system to at least:
 - before producing the collision detection unit for each edge, trim the PRM to reduce a total number of edges in the PRM.
6. The media of claim 5, wherein the instructions to trim the PRM direct the computing system to at least:

generate a plurality of problem instances using a tuple of the fixed scene, a selected number of obstacles from the distribution of obstacles, a selected starting robot pose of the starting robot poses, and a selected goal pose from the distribution of goal poses;

for each problem instance, evaluate a contribution factor of each edge with respect to finding a minimum solution path; and

remove edges of the PRM based on the contribution factors of those edges.

7. The media of claim 1, wherein the robot description comprises a Unified Robotic Description Format (URDF) file.

8. A method of robot motion planning, comprising:

receiving, at a motion planning hardware comprising a circuit for each edge of a probabilistic roadmap, data of an environment captured from a sensor system of the robot and a goal position, the data comprising 3D points expressing parts of obstacles in the environment;

removing edges of a graph data structure of the probabilistic roadmap that are indicated as having a collision with a part of an obstacle; and

determining a legal motion plan.

9. The method of claim 8, wherein determining the legal motion plan comprises:

performing a search of the graph data structure of the probabilistic roadmap for a shortest path between a current position of the robot and the goal position along edges of the probabilistic roadmap remaining after removing the edges that are indicated as having the collision with the part of the obstacle.

10. The method of claim 8, wherein the circuit for each edge of the probabilistic roadmap comprises a Boolean circuit, wherein the removing of the edges of the probabilistic roadmap comprises:

for each 3D point of the 3D points expressing the parts of the obstacles, processing that 3D point in parallel through each Boolean circuit, each Boolean circuit receiving the 3D point and outputting an indication of a collision or an indication of not a collision for the edge corresponding to that Boolean circuit; and

removing, from the graph data structure of the probabilistic roadmap, edges having the indication of the collision from its corresponding Boolean circuit.

11. A robot motion planning hardware comprising:
a plurality of collision detection units configured to process an input in parallel, each collision detection unit comprising a Boolean circuit for its corresponding edge of a probabilistic roadmap.
12. The hardware of claim 11, wherein the Boolean circuit comprises an OR circuit combining a representation, for each edge, of the 3D points to which that edge can collide.
13. The hardware of claim 11, wherein each Boolean circuit takes, in parallel, a same 3D point represented by a set of bits and returns a 1 or a 0 to indicate a collision or a not collision, respectively, with that 3D point.
14. A robot comprising:
a sensor system;
an actuator system;
a processing system;
a storage system; and
a robot motion planning hardware comprising a circuit for each edge of a probabilistic roadmap.
15. The robot of claim 14, wherein the robot motion planning hardware at least receives data of an environment captured from the sensor system, the data comprising 3D points expressing parts of obstacles in the environment, and
for each 3D point of the 3D points expressing the parts of the obstacles, the circuits process that 3D point in parallel to output an indication of a collision or an indication of not a collision with that 3D point for the edge corresponding to that circuit.
16. The robot of claim 15, further comprising a graph data structure of the probabilistic roadmap stored on the storage system, wherein the processing system identifies an edge of the graph data structure as being valid if the circuit corresponding to that edge outputs the indication of not the collision; and performs a graph search of the graph data structure for a shortest path along edges indicated as being valid.

17. The robot of claim 14, wherein the circuit for each edge of the probabilistic roadmap comprises logic gates.

18. The robot of claim 14, wherein the circuit for each edge of the probabilistic roadmap comprises lookup tables.

19. The robot of claim 14, wherein the robot motion planning hardware further comprises a path computation unit, wherein the path computation unit at least:

receives data of a goal position and collision detection results output from the circuits, and

selects a path in the probabilistic roadmap from a current position of the robot to the goal position using valid edges indicated in the collision detection results.




FIG. 1A




FIG. 1B




FIG. 1C







FIG. 2

**FIG. 3**

**FIG. 4**500**FIG. 5**




FIG. 6

[illegible]

FIG. 7




FIG. 8




FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2016/012204

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - B25J9/16 (2016.01)

CPC - B25J9/1664 (2016.01)

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)

IPC(8) - B25J5/00, B25J9/00, B25J9/16, G05B19/04, G05D1/00, G06F17/10, G06F19/00 (2016.01)

CPC - B25J9/1664, B25J9/1666, G05B19/4061, G05B2219/40317, G05B2219/40448, G05B2219/40465, G05D1/0217

G05D2201/0217, G06Q10/047, G06T19/003, G06T2210/21, Y10T307/826 (2016.01)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

US: 307/125, 318/587, 340/436, 700/245, 700/247, 700/248, 700/249, 700/250, 700/255

700/262, 701/25, 701/301, 701/533, 703/2, 703/7, 706/14 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Orbit, Google Patents, Google Scholar, Google

Search terms used: robot, motion, path, plan, collision, contact, impact, detection, probabilistic roadmap, PRM, logic, circuit, gate, boolean, OR, Oring, edge, joint, scenario, description, unified robot description format, URDF

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	KAVRAKI et al. Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION, VOL. 12, NO. 4, AUGUST 1996. [retrieved on 2016.03.03] Retrieved from the Internet: <URL:http://http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=508439&isnumber=11078>. Pages 566-580	1, 8-16, 18, 19
Y		2-7, 17
Y	US 6,049,756 A (LIBBY) 11 April 2000 (11.04.2000) entire document	2, 3, 17
Y	US 2014/0121833 A1 (SAMSUNG) 01 May 2014 (01.05.2014) entire document	4-6
Y	US 2014/0251702 A1 (WILLOW GARAGE, INC) 11 September 2014 (11.09.2014) entire document	7
A	US 8,855,812 B2 (KAPOOR) 07 October 2014 (07.10.2014) entire document	1-19
P,A	US 2015/0266182 A1 (ABB TECHNOLOGY LTD) 24 September 2015 (24.09.2015) entire document	1-19
A	US 8,825,207 B2 (KIM et al) 02 September 2014 (02.09.2014) entire document	1-19
A	US 8,666,548 B2 (LIM) 04 March 2014 (04.03.2014) entire document	1-19
A	US 2011/0153080 A1 (SHAPIRO et al) 23 June 2011 (23.06.2011) entire document	1-19

☐ 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

03 March 2016

Date of mailing of the international search report

21 MAR 2016

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774