Randomized Parallel Motion Planning for Robot Manipulators

Caigong Qin
Computing Laboratory
University of Oxford
Oxford OX1 3QD, Britain
qin@comlab.ox.ac.uk

Dominik Henrich
IPR, Dept. of Computer Science
University of Karlsruhe
D-76128 Karlsruhe, Germany
dhenrich@ira.uka.de

February 9, 1996

Abstract

We have presented a novel approach to parallel motion planning for robot manipulators in 3D workspaces. The approach is based on a randomized parallel search algorithm and focuses on solving the path planning problem for industrial robot arms working in a reasonably cluttered workspace. The path planning system works in the discretized configuration space, which needs not to be represented explicitly. The parallel search is conducted by a number of rule-based sequential search processes, which work to find a path connecting the initial configuration to the goal via a number of randomly generated subgoal configurations. Since the planning performs only on-line collision tests with proper proximity information without using pre-computed information, the approach is suitable for planning problems with multirobot or dynamic environments.

The implementation has been carried out on the parallel virtual machine (PVM) of a cluster of SUN4 workstations and SGI machines. The experimental results have shown that the approach works well for a 6-dof robot arm in a reasonably cluttered environment, and that parallel computation increases the efficiency of motion planning significantly.

1 Introduction

This report documents parts of the work in the framework of the HEROS Project[1], which has been carried out at the Institute for Real-Time Computer Systems and Robotics at the University of Karlsruhe. We have come up with a novel approach to parallel motion planning for robot manipulators in the three dimension case. The implementation has been conducted on a parallel virtual machine (PVM) of a cluster of SUN4 workstations and SGI machines available in the institute.

The issue of robot motion planning has been studied for more than a few decades and many important contributions to the problem have been made ([Lat91, Hwa92]). One of the most important results is the application of the concept of configuration space ([LP83]). However, it has been shown that the complexity of the generalized movers problem is exponential with respect to the configuration space dimension ([SS83]) and is PSPACE hard ([Rei79]). Although the configuration space (C-space) approach provides a good framework for theoretical research, motion planning purely based on the approach normally results in a non-practical planner for real-life situations. The C-space approach is well-suited for robots of few DOF (< 4) in a static environment. But it is unfavourable for problems with multiple robots or dynamic environments, since in these cases, C-obstacles and C-free-space have to be recomputed, which is computationally very expensive. In order to avoid the complexity of the explicit computation of the configuration space (*i.e.*, C-free-space and C-obstacles) or its approximation, as was done in [LPJMO92], our method works implicitly in the discretized configuration space with a number of explicit and implicit constraints. The explicit constraints result from the mechanical consideration of the robot, such as the limitations of joint motions. The implicit ones are derived from collision avoidance between the robot and obstacles. In this way, whether a configuration q of the robot is in C-free space or C-obstacle space is determined through collision detections between the robot and the obstacles around, which is calculated in the workspace.

Since the procedure of collision detection could be heavily invoked throughout motion

[1]HEROS - Hazardous Environment RObot Systems, ERBCHRXCT930086, a Human Capital and Mobility Project of the European Union.

planning, an efficient collision detection scheme is essential. We have adopted the *Oxsim* framework ([QCM95]) to build up the robot and world models for the simulation of the parallel motion planning. In this framework, the collision detection is based on the modified GJK distance algorithm ([GJK88]) with some heuristics. Nevertheless, if the collision detection is called stepwise during planning in the discretized C-space, there can be a very large number of detections to consume significant computation time. For the extreme case of the robot in an environment without any obstacles, the stepwised collision detections along a planned path are actually not necessary. To minimize the number of collision detection invocations, we have utilized the heuristic of distance information to increase efficiency.

Algorithms for motion planning can be generally classified in terms of whether the algorithm is complete. Complete algorithms are only practical for simple situations with few degrees of freedom. The main reason for that is due to the faact that the search space for a path is so large. Therefore, in this report, we will restrict ourselves to considering incomplete (but useful) algorithms working in the discretized configuration space.

In an effort to decrease the computation time, some researchers have worked on parallel computations of motion planning ([Hen96]). With parallel processing, not only can some existing sequential algorithms be parallelized but some new parallel algorithms can be designed based on the characteristics of parallelism. We propose a parallel algorithm that would not make much sense for a sequential machine but is fairly effective with parallel processing. The algorithm is implemented using a parallel virtual machine (PVM), which is a software package that allows a heterogeneous network of parallel and serial computers to work as a single concurrent computational resource. Along with the package, a number of routines are provided with the support of a user interface. The main advantages of PVM are that it provides a set of user interface primitives that may be incorporated into existing procedural languages and that it is available on most of the network and/or parallel architectures. For further details about PVM, refer to [GBD+93].

The rest of the report is organized as follows: Section 2 relates our work to previous

one in path planning and parallel processing approaches to motion planning. Section 3 presents our approach to motion planning in detail. Section 4 introduces some heuristics for improving the performance of the motion planner. Section 5 details the implementation of the algorithm and the experimental results. Section 6 gives the conclusion and possible extensions.

2 Related Work

With regard to autonomous robot systems, motion planning is one of the most important aspects in robotics. It has been attracting a great deal of interest over the last 20 years. Generally, motion planning problems can be described in this way: given initial and goal configurations, we need to find a path for a robotic device, which is to move along the planned path without colliding with obstacles around it. The path connects the initial and the goal configurations.

Glavina [Gla90] proposed an algorithm to solve the findpath problem by combining a goal-directed straight-and-slide search and a randomized generation of subgoals. The idea is to conduct the search by following the straight line till the searching point reaches a C-obstacle in the discretized configuration space. Then, the searching configuration point slides along the obstacle boundary only if it reduces the configuration distance, which is a function of a suitably weighted combination of the configuration variables. The sliding process continues until the point gets stuck at a local minimum with respect to the configuration distance function. Then, a new subgoal is generated randomly. The reachability of the subgoal is tested by the same straight-and-slide searching method from all introduced points (start and goal, and previous subgoals). Eventually, a site graph can be constructed as an abstract representation of the C-free-space. During the entire process, stepwise collision tests are carried out in order to detect whether the point is running into a C-obstacle. The algorithm was implemented using a moving polygonal object and the environmental polygonal obstacles in the 2D case.

Our work differs from the previous research in various ways. Rather than working on the moving polygonal object in a 2D case, we consider motion planning for robot manipulators in 3D workspaces. We employ a complete domain-dependent rule base

to guide path searching. The number and depth of local minima are reduced through a number of subgoals randomly or intentionally generated in parallel processing. In addition, we utilize some heuristics to reduce the number of collision detections instead of conducting stepwise collision tests during planning.

Qin [QCM95] presents a solid modelling scheme which is useful for efficient 3D path planners. The scheme makes use of the enhanced version of Gilbert, Johnson and Keerthi's minimum distance algorithm ([GJK88]). This provides an efficient prototype to be tailored to collision detection and to distance computation respectively. We have adopted this scheme in this work. In addition, other heuristics for speeding up collision detection are also being investigated by Henrich and Cheng [HC92], who introduced hierarchical representations of both obstacles and the robot working favourably for complex environments.

Kavraki and Latombe [KL94] proposed an approach concerning the randomized preprocessing of the configuration space to built up a global network of connected configurations. The idea is to use *generate-and- test* method to construct a network of randomly but well selected collision free configurations. However, the algorithm requires that each generated configuration be checked to see if it is in C-*free* space, which could be computationally very expensive.

Challou *et al.* [CGK93] and [CGKO95] presented a parallel motion planner using the parallel formulation of a randomized heuristic search. The algorithm is based on the parallelisation of the randomized robot planning method proposed by [BL91].

Henrich [Hen96] presented an extensive overview of the parallel approaches to robot motion planning. The parallel approaches are divided into four classes: grid-based, graph-based, potential field, and mathematics programming. The method presented in this report can be generally classified as grid-based.

3 Outline of the Approach

Our work focuses on developing an effective approach towards solving the path planning problem for industrial robot arms operating in a reasonably cluttered environment. This means that the workspace of the robot arm is not maze-like. So it is basically

assumed that the path planning problem has a number of acceptable solutions, though the assumption is not compulsory.

The approach is applied in the discretized configuration space of the robot arm. The C-space needs not be represented explicitly with the support of a fast collision test and geometric reasoning. This avoids problems of computational complexity and memory requirement for the transformation between the robot worldspace and the C-space. Given the start and goal configurations, a conventional search (such as depth-first search, best- first search, *etc.*) within the C-free space will readily involve a great deal of backtracking, and will therefore require a large amount of computation. The main reason is that the search space is normally very large and has some local minima with respect to some heuristic function of configuration parameters. One of the human-inspired pathfinding strategies is to divide the whole complicated search task into several simpler sub- tasks by setting up proper subgoals between the start and the goal configurations. However, the problem of how to set up subgoals properly is usually not trivial, and it may help to make use of the generalized Voronoi diagrams (GVD) of the robot's free workspace or other global information.

Figure 1 shows an example of a path search with dead-end obstacles (*cul-de-sac* type). It illustrates that while the direct search from the start to the goal, or vice verca, is not cheap, an indirect search with some via points may be very helpful. This phenomenon, along with the availability of parallel processing, encourages us to come up with a straight *randomized parallel search* algorithm. The general idea of this algorithm is to randomly generate a number of subgoals in the discretised C-free space. Then, parallel searching with each subgoal attempts to find a path connecting the initial configuration with the goal via the subgoal configuration. The purpose of the approach is for the motion planning system to cope with some deep local minima (see Figure 1). In this sense, it is a two-phase search which tries to find a subpath connecting the initial with the randomly generated subgoal and a subpath connecting the subgoal with the final goal. The reasons for not using a three- or multi- phase search are (1) that a path from the initial to the goal via more than two randomly generated subgoals will usually be longer in length and (2) the planner may correspondingly take a longer time than that in a two phase search (see Section 5). But we claim that a three- or multi-phase search will be very effective

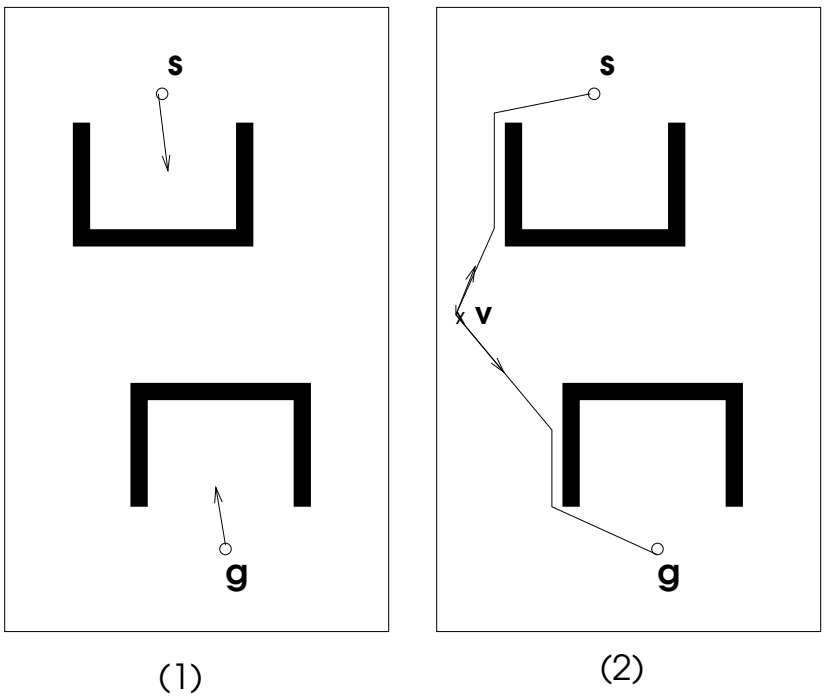

Figure 1: An illustrative diagram for path search. (1) Searching for a path is costly no matter whether from the start to the goal or vice versa. (2) With a via point $\mathbf{v}$, the combination of searching from $\mathbf{v}$ to the start and to the goal, respectively, is much easier.

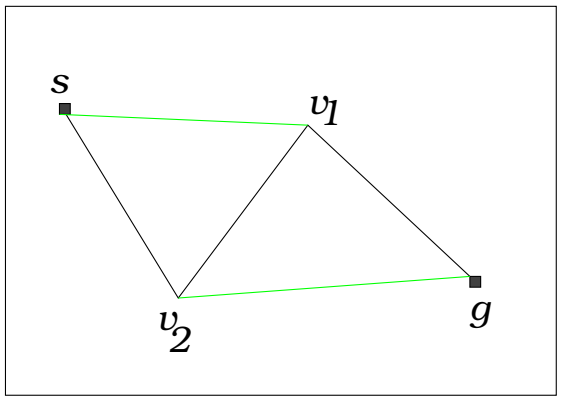

Figure 2: Compared to the two 2-phase search paths: s- v_1-g and s-v_2-g, the 3-phase search may result in paths: s-v_2- v_1-g or s-v_1-v_2-g, both of which are longer than the 2-phase search results (Arbitrary obstacles are omitted for clarity).

to help the robot avoid some very deep local minima if all subgoals are generated under the guidance of some global information. For the current implementation, we restrict ourselves to the two-phase randomized parallel search. Experiments have shown that the motion planning system works well for a 6-*dof* robot arm in a reasonably cluttered workspace (see Figure 7).

Our motion planning simulation system consists of several modules. One module, called the *master module*, is to build up the robot and the workspace model with the help of a geometric modeller, known as ROBMOD ([CA88]). Another module for visualisation, known as the *display module*, has been directly imported from the OxSim robot simulation system. The master module is to send information concerning the planned path to the display module for updating the robot position. Another important module, called the *pathfind module*, is designed to connect to the master module. The master module provides data concerning the initial and the goal configurations to the pathfind module. In return, the latter reports path information to the former. The pathfind module is designed to search a path for the given initial and goal configurations in parallel. Each task within the module is to find a path connecting the initial configuration to the goal via a randomly generated subgoal. Whenever such a path has been found by one of the parallelly running processes, the path will be reported to the master module and then all tasks will be terminated immediately. The interrelation of the modules is shown in Figure 3.

3.1 Sequential search

For each path search task from one configuration to the other, we have employed an expert system method to perform path searching with the guidance of a set of rules. We use CLIPS [2] as the expert system shell, which is embeded as a module into C++ programs. For generality, the following discussion is for a general n-dof robot arm.

Let $\mathcal{C}^n$ be the n-dimension configuration space. In order to reason and search, we discretize the configuration space $\mathcal{C}^n$ into a rectangloid grid $\mathcal{G}C^n$ (with appropriate modular

[2] C Language Integrated Production System, developed in NASA.

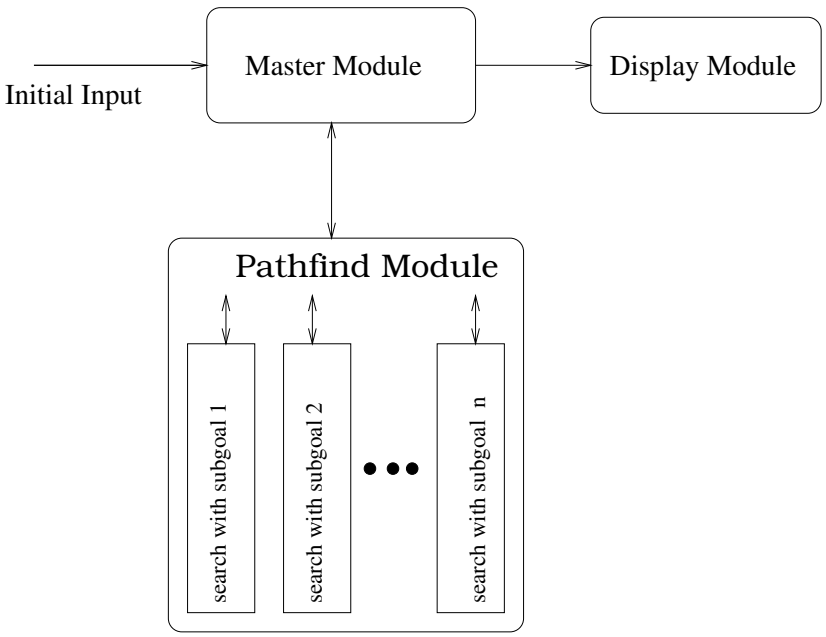

Figure 3: The outline of the robot motion planning system using the randomized parallel search algorithm.

arithmetic for the angular joint parameters). For each configuration $q = (q_1, q_2, ..., q_n)$ (in $\mathcal{C}^n$) at which the robot is within its workspace, we have

$$q_i^{min} \leq q_i \leq q_i^{max} \ , \qquad i = 1, 2, ..., n. \tag{1}$$

which specifies the limitations of joint motions.

For convenience of reference, we define

$$\mathcal{S} = \{q : \mathcal{G}C^n | q = (q_1, q_2, ..., q_n) \text{ satisfies } (1)\}$$

as the set of configurations at which the robot is within its workspace.

In $\mathcal{G}C^n$, each grid node can be indexed by a point in $\mathcal{Z}^n$, where $\mathcal{Z}$ is the set of integers through a mapping:

$$\mathcal{F} : \mathcal{G}C^n \rightarrow \mathcal{Z}^n \ , \quad i.e. \ \mathcal{F}(q) = (i_1, i_2, ...i_n) \ . \tag{2}$$

where $q \in \mathcal{S}$ and $i_\ell \in \mathcal{Z}$ $(1 \leq \ell \leq n)$. In this sense, $\mathcal{Z}^n$ is regarded as a symbolic abstraction of $\mathcal{G}C^n$. The reasoning system runs directly on $\mathcal{Z}^n$ with other symbolic facts. At the lower level, reasoning is conducted directly on geometric data. The two levels are connected through a number of geometric primitives.

Similar to [Lat91], for a given point $x \in \mathcal{Z}^n$, its p- **neighbors** ($1 \leq p \leq n$) are defined as all points in $\mathcal{Z}^n$ having at most p coordinates differing from those of x by the amount of exactly one increment in absolute value. That is, x has six 1- neighbors in 3D and $2n$ 1-neighbors in n dimensional space.

In our reasoning system, path planning considers only the 1-neighbors of the current configuration as candidates to move to in each step. Then a number of rules are designed to select an optimal candidate according to the following *cost function* $\mathcal{P}$:

$$\mathcal{P}(q) = \begin{cases} +\infty & \text{if } q \text{ does not satisfy constraints (1)} \\ +\infty & \text{if } q \text{ is not in } \textit{C-free} \text{ space} \\ f(q) & \text{otherwise} \end{cases}$$

where $f(q)$ must be defined to be decreasing as q approaches the goal configuration g_{goal} and $f(g_{goal}) = 0$. In this sense, $f(q)$ is similar to an attractive force in the artificial potential method [Lat91], the force which draws the robot towards the goal configuration. Therefore, we can define $f(q)$ as

$$f(q) = \sum_{i=1}^{n} \omega_i((\mathcal{F}(q_{goal}))_i - (\mathcal{F}(q))_i)^2 \tag{3}$$

where $\omega = (\omega_1, \omega_2, ...\omega_n)$ is a weight vector with $\omega_i > 0$, which is determined heuristically using the model of the robot manipulators.

With $f(q)$ defined in (3), the symbolic reasoning can be directly performed since reasoning at this level can be conducted in $\mathcal{Z}^n$.

In support of some procedural geometrical primitives, a number of rules have been introduced up to guide the search for a path in $\mathcal{G}C^n$. In addition to rules for initialisation, for the detection of goal reaching and for the generation of candidates for the next movement, there are several other rules devoted to decision-making for the next step. At each step, all the candidates which have been previously visited will be removed first; then the candidate with the lowest cost with respect to $\mathcal{P}(q)$ is checked to see if it is out of the robot workspace or in C- obstacles. If it is in the C-free space and within the robot workspace, the candidate will be accepted and all others removed; otherwise, it will be removed and the next best candidate will be checked. Provided the selection procedure results in an empty set of candidates, a rule will force the search to backtrack. For more details on the rule-based search, refer to [QC96]. In this implementation, we have used 12 rules with the support of 8 procedural geometrical primitives.

In addition to the rule-based search method, other conventional search algorithms, such as A^* search, Iterative Deepening A^*, or other depth-first search ([KGGK94]), may be used as an alternative. But these search mechanisms are usually either computationally expensive or require high memory, which is unfavourable in motion planning.

3.2 Parallel search

Given the start and goal configurations, the sequential search algorithm alone may take too much computation time to find a path. To improve the efficiency of the path search, we take advantage of parallel processing to conduct the path search using the randomized parallel search algorithm. In the example shown in Figure 4, a number of processes concurrently conduct a search for a path connecting the start and the goal configurations. Each process generates one subgoal randomly in the C-free space and then starts searching for a subpath from the subgoal to the start configuration and for another subpath from the subgoal to the goal configuration. The final path is the proper concatenation of the two subpaths. Whenever a process returns such a final path, all search processes will be terminated.

The termination criterion used here is very simple and easy to implement under the PVM. Other termination criteria, such as imposing a time bound or the cost limit of an incumbent path in terms of length, may not work well for our algorithm, because time and path length are hard to be estimated in advance.

For each process, the search always starts from the subgoal to the start and to the goal configurations, rather than from the start configuration to the subgoal and then from the subgoal to the goal configuration or vice versa. It is out of concern that if the start or the goal configuration is in a deep local minimum, it can be hard to jump out of such a minimum by starting the search from the start or the goal configuration even with help of subgoals. In this case, it may be made easier by searching from a subgoal to the start and the goal configurations. This effect is especially exploited by the parallel algorithm where multiple subgoals are used.

As we know, the final path obtained in this method tends not to be optimal even if each subpath is optimal. But we claim that it is probabilistically optimal provided that the number of random subgoals tends to be infinite.

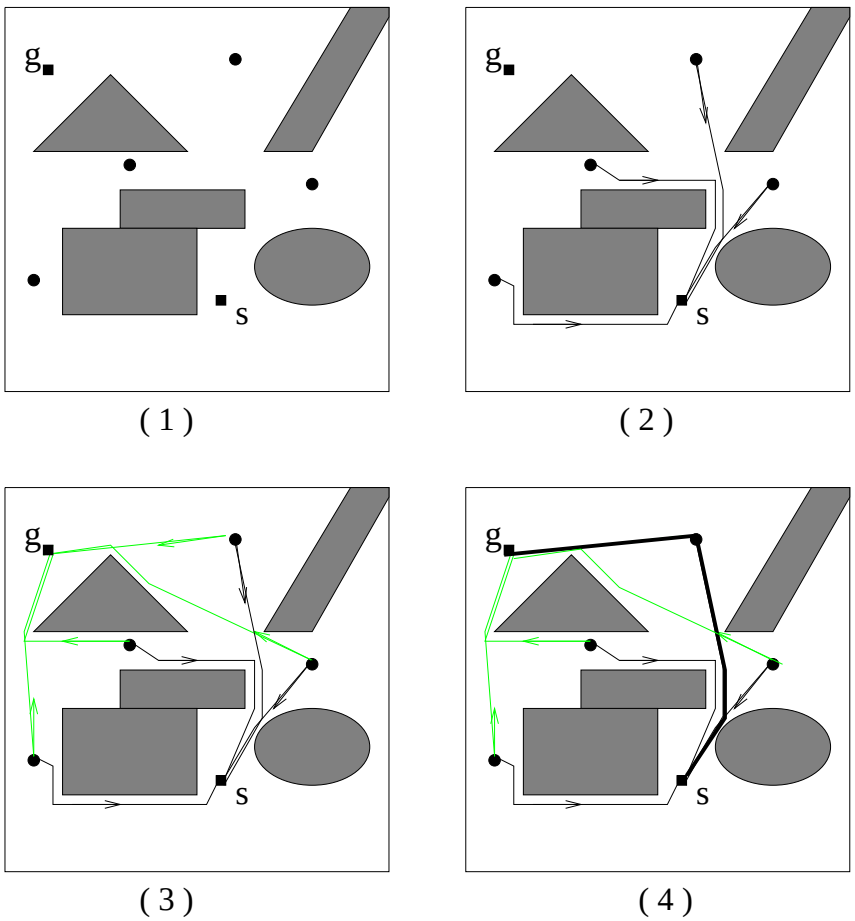

Figure 4: An illustrative diagram of the randomized parallel search in the 2D case for simplicity. (1) Randomly generated subgoals in C-free space. (2) Each process conducts a search for a subpath from the subgoal to the start configuration. (3) After finishing the search at step 2, each process starts searching for another subpath from the subgoal to the goal configuration. (4) The first path a process returns will be accepted as the final path and then all search processes will be terminated, though some processes are on the way of searching.

In some situation with or without sparsely-cluttered obstacles, the search for a path from the start to the goal configuration or vice versa may be more efficient. To make the method more robust, two extra processes are designated to conduct the search directly from the start to the goal configuration and from the goal to the start configuration respectively, while other processes conduct the search through subgoals. Therefore, it holds that the time for planning with one process is not less than that with multi-processes.

Overall, our approach to parallel processing has some advantages:

- no heavy communication load. The only communications are for problem broadcast and solution collection between the master module and the pathfind module.
- no load balancing necessary, since all search processes will be terminated whenever a process returns a final path.
- independent of parallel architecture, such as processor interconnection topologies.
- low memory space requirements, as each sequential search is like a depth-first search with proper backtracking.

4 Heuristics

Human beings are very good at using heuristics in everyday life. Employing proper heuristics can generally provide a *short-cut* in problem- solving. This is the same in robot motion planning systems. In this section, we will introduce some heuristics that have been used in our algorithm.

4.1 Discretisation resolution

As previously mentioned, the motion planning takes place in a discretised configuration space of the robot manipulator. The resolution settlement of discretisation is also an important issue. There is a trade-off in the granularity of discretisation or resolution: too fine will increase the search space exponentially and too coarse may result in failing to find a path even if there one exists.

We have adopted a heuristic to help set up the discretisation resolution of the C-space. Instead of having a uniform resolution along each configuration coordinate, we set up the resolution along each coordinate differently by estimating the maximum movements of the robot's endeffect at each step the robot moves along the coordinate. The resolution should be so fine that the maximum movement of the robot endeffect is not more than a pre-set distance at each step the robot moves along the coordinate. In this way, generally, the nearer a joint is to the base, the finer the discretisation resolution is for the corresponding joint angle.

Formally, for the i-th coordinate q_i of the C-space, let N_i be the number of intervals along q_i. Then,

$$N_i = \lceil \frac{q_i^{max} - q_i^{min}}{\Delta\theta_i} \rceil$$

and

$$\Delta\theta_i = 2\arcsin(\frac{MaxMove}{2l}) \quad ,$$

where q_i^{max} and q_i^{min} are the limits of joint motions (see Formula (1)), l is the length between the center of the i-th joint to the farthest point the endeffect can reach, and $MaxMove$ is a pre-set distance the robot moves along the coordinate at one step (see the picture).

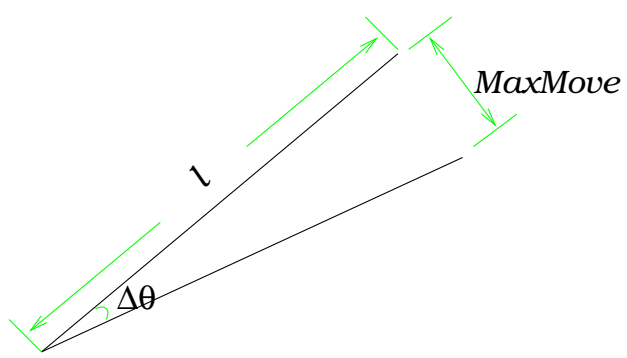

4.2 Prediction of maximum movement

Geometric reasoning plays an important role in motion planning. [LP87] used it to build up the approximation of the configuration space by calculating the maximum movement of links. Similarly, [Bee92] used the idea to provide analytical formulae for a specific robot model to reason about the occupancy of C-obstacles.

This heuristic is used to improve the rule-based sequential search algorithm by minimizing the number of collision detections. As introduced in § 3.1, the rule-based search performs a collision test at least once during each movement step. So, even in the extreme case that there is no obstacle around, the search still must conduct unnecessary collision tests at each step along the line connecting the start and the goal configurations. To cope with the problem, we make use of a heuristic called *prediction of maximum movement.*

The idea of the heuristic is to utilize proximity information in path searching rather than simply conducting collision tests at each step. As we know, the movement of the robot endeffect is not more than the pre-set $MaxMove$ at each step that the robot moves along any one of the coordinate axis directions of the C-space. This also implies that the movement of any other points on the robot arm is not more than $MaxMove$ at each step. Let d be the minimum distance between the robot and obstacles around. Let $FreeSteps$ be

$$FreeSteps = \lfloor \frac{d}{MaxMove} \rfloor \ .$$

Then, we can conclude that the robot arm is ensured to be collision-free for $FreeSteps$ consecutive steps. This helps release the searching from collision tests for $FreeSteps$ steps. Figure 5 shows an example with $FreeSteps = 3$.

5 Experimental Results

The motion planning system is implemented in C++ with the embedded rule- based expert system. Experiments have been carried out on a parallel virtual machine of a cluster of SGIs and SUN4 workstations. The number of machines or processors available is limited (up to 45 workstations). In addition, the heterogeneneous machines range from SUN Sparc classic to Sparc20 and from SGI IRIX5.3 R3000 to IRIX5.3 R4000. In this section, we will present some experimental results and show the performance of the planning system with some examples.

Due to the randomness of our parallel algorithm, the time taken to solve one problem may vary more or less from one run to another[3]. All the data presented in this section

[3]Under the PVM, other users on the network may also affect the experimental results. But the

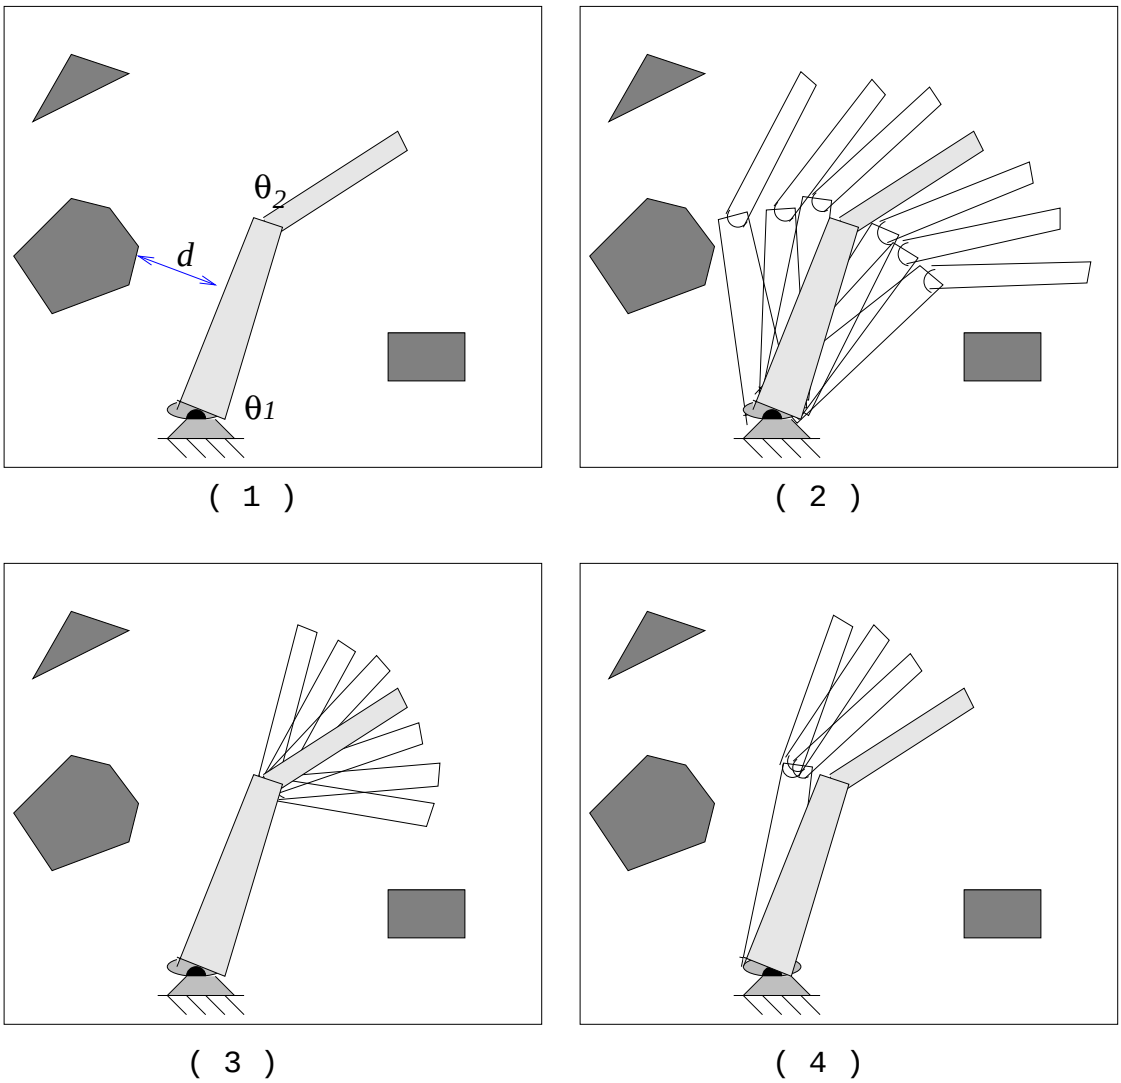

Figure 5: An illustration of using the heuristic of prediction of maximum movement with $FreeSteps = 3$. (1) The arm's current configuration. (2) The arm is ensured to be collision-free while moving along θ_1 for 3 consecutive steps in either the positive or negative direction (with θ_2 fixed). (3) Similar as the above with θ_1 and θ_2 exchanged. (4) The arm is collision-free while moving along the coordinate axis (either θ_1 or θ_2) directions of the C- space for 3 consecutive steps.

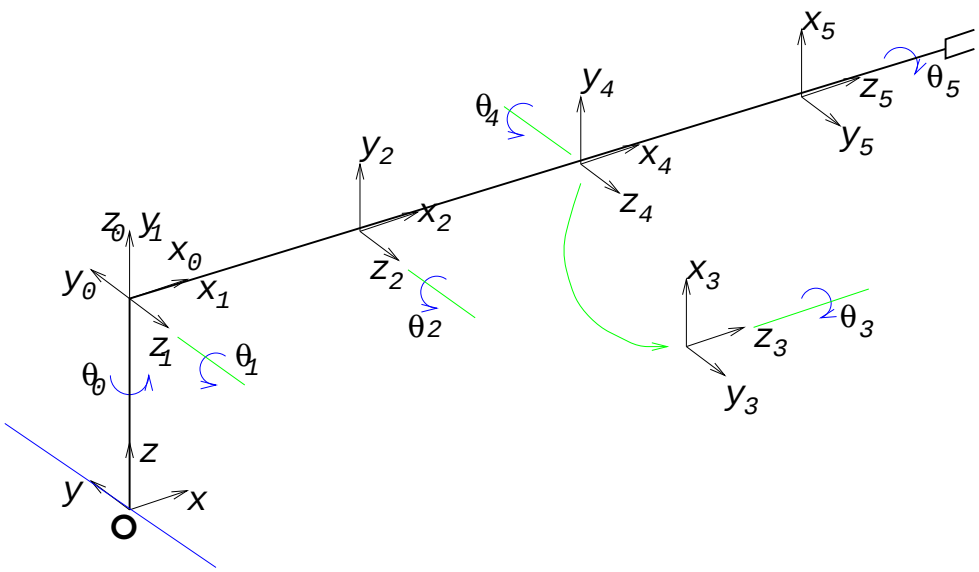

Figure 6: The diagram of the kinematic description of the 6-*dof* robot model.

are calculated by taking the mean value of 10 numbers which are selected from 12 experimental results by deleting the maximum and the minimum values, so that the data collected can be more representative.

In the experiment, we used a 6-*dof* robot manipulator of PUMA 200 type as the robot model. Figure 6 shows the kinematic description of the model. In our experiments on a SGI machine of IRIX5.3 R4000, it takes 3.4 ms to modify the robot from one configuration to another, and averagely about 0.435 ms to conduct a single collision detection for the case shown in Figure 7. The discretisation resolution has been set through the heuristics (see Section 4), so that the joint i of the robot moves $\Delta\theta_i$ at each movement step, where $(\Delta\theta_0, \Delta\theta_1, \Delta\theta_2, \Delta\theta_3, \Delta\theta_4, \Delta\theta_5) = (2.23°, 2.23°, 3.47°, 6.88°, 6.88°, 9.56°)$. The weight vector ω for the cost function $f(q)$ in Formula (3) is set to be $(6, 5, 4, 3, 2, 1)$, which gives higher priority to the first few joints.

Experiments have also shown that employing the heuristic of prediction of maximum movement improves the performance of the planning system. In the example shown in Figure 7, averagely, the planning takes about 22.69 seconds and conducts 622 collision tests without using the heuristic. Nevertheless, it only takes about 9.64 seconds and conducts 371 collision tests by use of the heuristic on the PVM of a cluster of ten SUN4 and four SGI machines.

experiments took place at a time when this effect could be mimimal.

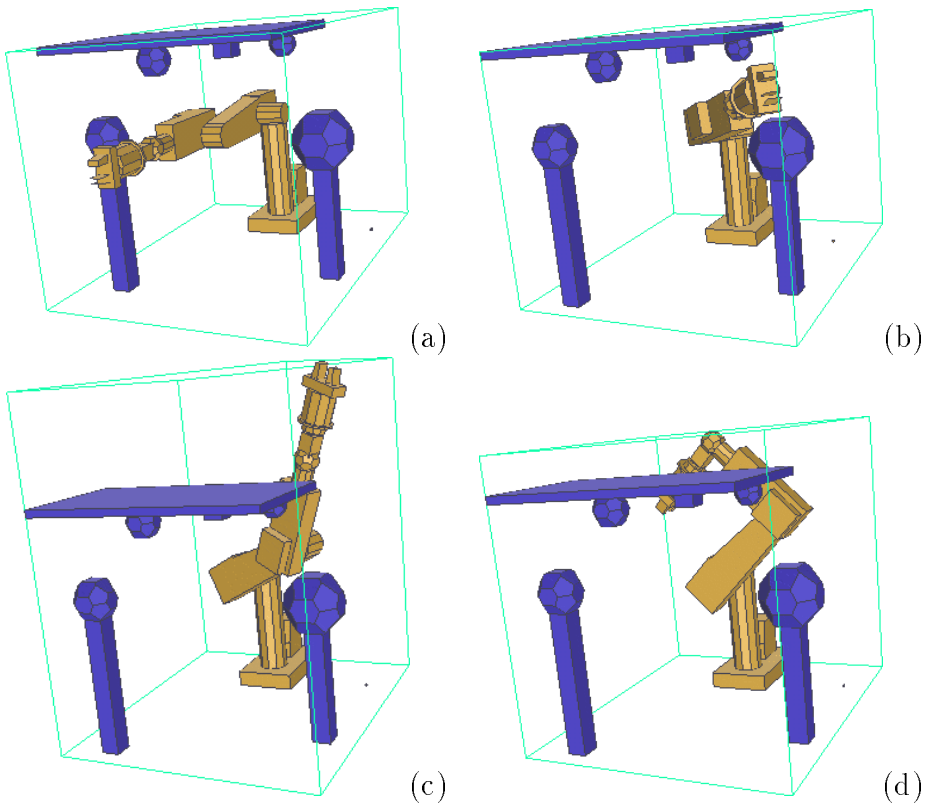

Figure 7: Snapshots of the robot in motion. The planning takes 9.64 seconds averagely on the PVM of a cluster of ten SUN4 and four SGI machines. (a) The robot is in the initial configuration, (b) and (c) the robot is searching its way towards the goal, (d) the arm reaches the goal.

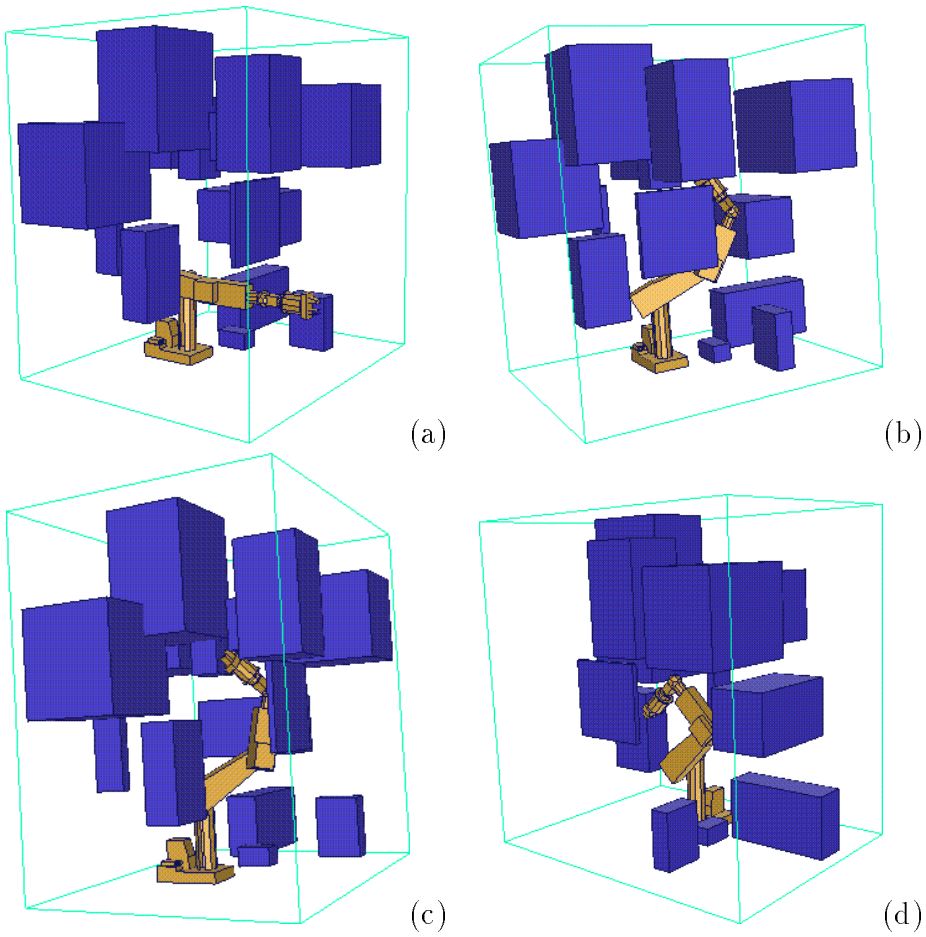

Figure 8: The robot arm is in motion. The planning takes 32.96 seconds averagely on PVM of a cluster of 41 SUN4 and 4 SGI machines.

Another example shown in Figure 8 is to find a path for the robot arm in the workspace occupied with 15 randomly sized and located rectangular obstacles. It takes 32.96 seconds on the PVM of a cluster of available 41 SUN4 and 4 SGI machines. Further experiments show that the efficiency increases with increasing number of processors (see Figure 9). Figure 10 shows that the more subgoals employed, the better a path in length is found in the example which is shown in Figure 8. Nevertheless, for this example, there is not much difference between the ways of searching from subgoals to the start and from the start to the subgoals for the first subpath. The reason is that the start configuration of the robot has a relatively open neighborhood of the C-free-space.

As we know, *speedup* is defined as the ratio of the time taken to solve a problem on a single processor with the best sequential algorithm to the time needed to solve the same problem on a parallel machine with identical processors. The performance of our

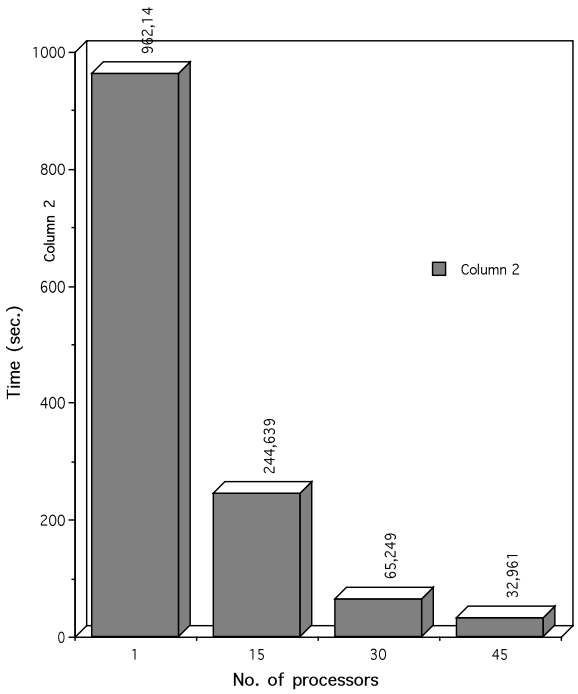

Figure 9: The illustrative graph of the performance of the randomized parallel planning algorithm (each process is assigned to one processor).

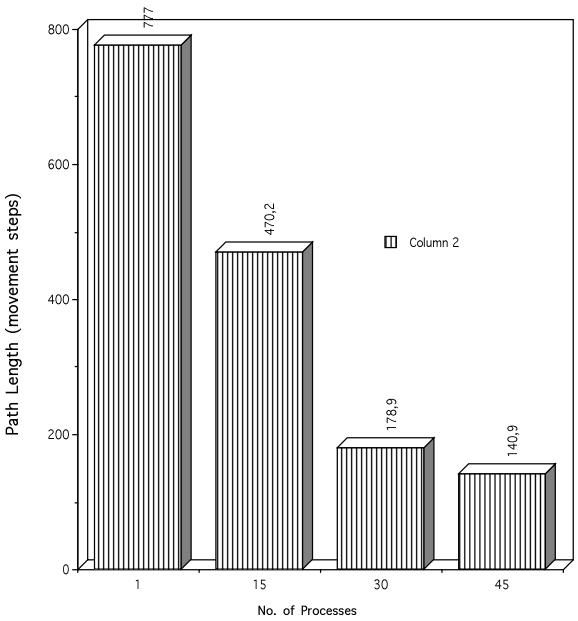

Figure 10: The more subgoals employed, the better a path in length can be found (each process deals with one subgoal).

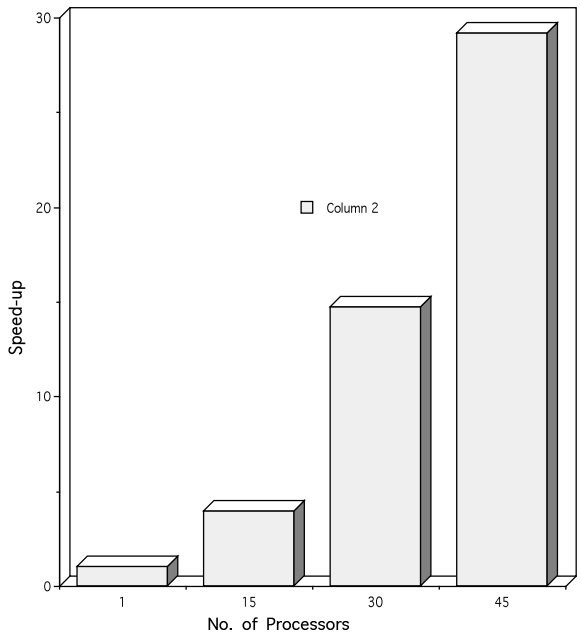

Figure 11: The illustrative graph for speedup of the algorithm for one example.

algorithm on the example shown in Figure 8 implies the speedup achieved is as shown in Figure 11, though the experiments are not conducted on identical processors and the time taken to find a path on a single processor may not be optimal.

Figure 12 gives a further example for finding a path for the robot arm in the workspace occupied with 20 randomly sized and located rectangular obstacles. It takes 34.81 seconds on the PVM of a cluster of available 41 SUN4 and 4 SGI machines.

6 Conclusions

This report presents a novel approach to parallel motion planning for robot manipulators in 3D workspace. The approach is based on the *randomized parallel search* algorithm and focuses upon solving the path planning problem for industrial robot arms working in a reasonably cluttered workspace. The path planning is conducted in the discretized configuration space which needs not to be represented explicitly. The general idea of the randomized parallel search algorithm is to randomly generate a number of subgoals in the discretised C-free space and then do a parallel search with each subgoal. A path is then found connecting the initial configuration to the goal via the subgoal configuration. Searching for subpaths connecting a subgoal to the start and the goal configuration is

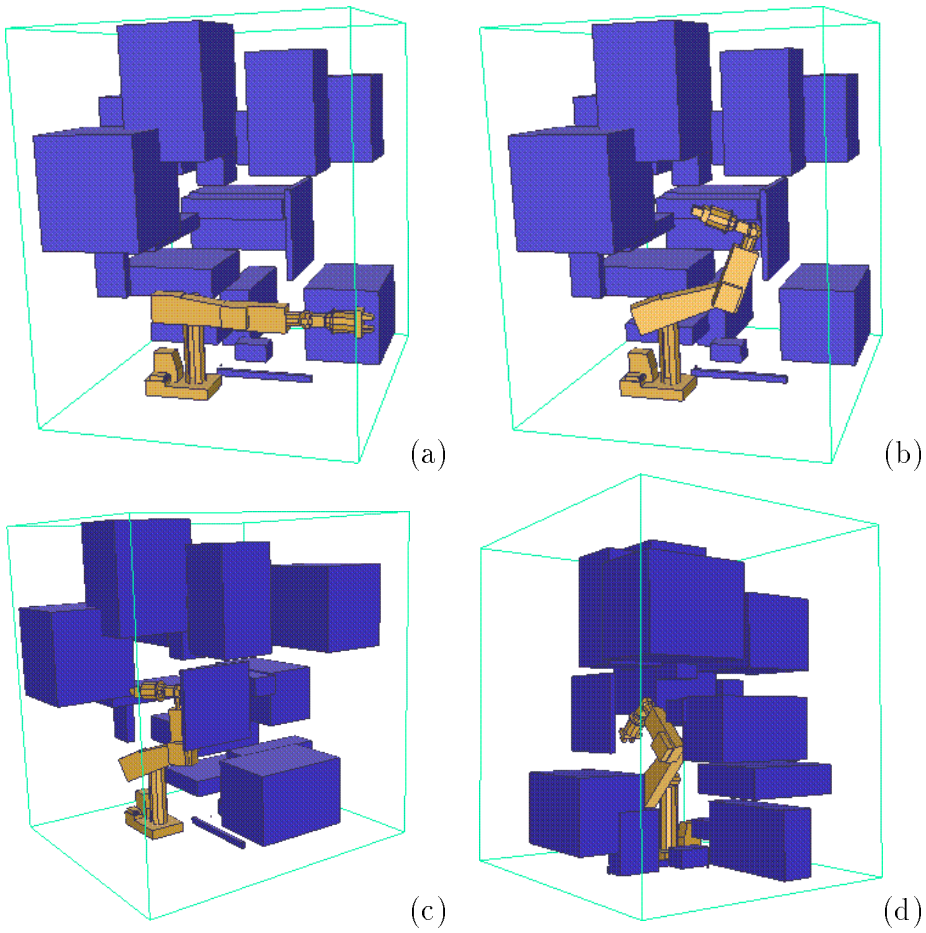

Figure 12: The robot arm is in motion. The planning takes 34.81 seconds averagely on PVM of a cluster of 10 SUN4 and 4 SGI machines.

based on a rule-based sequential search method.

The implementation is carried out on a cluster of SUN4 workstations and SGI machines under the support of the parallel virtual machine. The experimental results have shown that the approach works well for a 6-dof robot arm in a reasonably cluttered environment, and that computation with multi-processors increases the efficiency of motion planning significantly.

However, the method is not recommendable for general redundant robot manipulators, as a too large search space may result in high computation time. In addition, the method of the current implementation makes use of randomly generated subgoals for a two-phase parallel search. This may not work well in a very complex workspace, such as a maze-like one. The major problem lies in the fact that no global information is used for subgoal selection. An extension of the work can be to investigate the method of the subgoal selection under the support of global topological information of the environment, for example, using the generalized Voronoi diagrams or information derived from artificial potential fields in the workspace, to improve the robustness and efficiency of the method.

Since the planning performs only on-line collision tests with proper proximity information without using pre-computed information, the approach is suitable for planning problems with multirobot or dynamic environments. Therefore, in addition to the case of dynamic environments, one extension of the work can be to exploit motion planning for two robot arms working in a shared workspace. The issue of coordinating the two arms can be solved by integrating time scheduling into our current algorithm for path searhing.

Acknowledgments

The research work was carried out at Institute for Real-Time Computer System and Robotics, Prof. Dr.-Ing. U. Rembold and Prof. Dr.-Ing. R. Dillmann, University of Karlsruhe, D-76128 Karlsruhe, Germany. The material is based upon the work under the support of the contract ERBCHRXCT930086 of a Human Capital and Mobility Project of the European Union.

References

[Bee92] E. Beer. Jam - just another algorithm to solve a motion planning problem: Ein geometrisches bahnplanungsverfahren fuer zweiarmroboter. Dissertation, University of Karlsruhe, 1992.

[BL91] J. Barraquand and J.C. Latombe. Robot motion planning: A distributed representation approach. *Int. J. Robotics Res.*, 10(6):628–649, 1991.

[CA88] S.A. Cameron and J. Aylett. Robmod: A geometry engine for robotics. *IEEE Int. Conf. Robotics & Automation*, pages 880–885, April 1988. San Francisco.

[CGK93] D.J. Challou, M. Gini, and V. Kumar. Parallel search algorithms for robot motion planning. In *IEEE Int. Conf. Robotics & Automation*, pages 46–50, Atlanta, May 1993. Vol.2.

[CGKO95] D.J. Challou, M. Gini, V. Kumar, and C. Olson. Very fast motion planning for dexterous robots. In *IEEE Int. Conf. Robotics & Automation*, pages 201–206, May 1995. Nagoya.

[GBD+93] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V Sunderam. Pvm 3 user's guide and reference manual, May 1993. Electronic publish, Math. Sci. Section, Oak Ridge National Laboratory.

[GJK88] E.G. Gilbert, D.W. Johnson, and S.S. Keerthi. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE J. Robotics & Automation*, 4(2):193–203, April 1988.

[Gla90] B. Glavina. Solving findpath by combination of goal-directed and randomized search. In *IEEE Int. Conf. Robotics & Automation*, pages 1718–1723, Cincinnati, Ohio, May 1990.

[HC92] D. Henrich and X. Cheng. Fast distance computation for on-line collision detection with multi-arm robots. In *IEEE Int. Conf. Robotics & Automation*, pages 2514–2519, Nizza, France, May 1992.

[Hen96] D. Henrich. Parallel processing approaches to motion planning – an overview. In *IEEE Int. Conf. Robotics & Automation*, Minnesota, 1996.

[Hwa92] Y.K. Hwang. Gross motion planning - a survey. *ACM Computing Surveys*, 24(3), September 1992.

[KGGK94] V. Kumar, A. Grama, A. Gupta, and G. Karypis. *Introduction to Parallel Computing: Design and Analysis of Algorithms.* The Benjamin/Cummings Publishing Company, Inc., 1994.

[KL94] Lydia Kavraki and Jean-Claude Latombe. Randomized preprocessing of configuration space for fast path planning. In *IEEE Int. Conf. Robotics & Automation*, pages 2138–45, San Diego, May 1994.

[Lat91] Jean-Claude Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, 1991.

[LP83] T. Lozano-Pérez. Spatial planning: A configuration space approach. *IEEE Trans. on Computers*, 32(92), 1983.

[LP87] T. Lozano-Pérez. A simple motion-planning algorithm for general robot manipulators. *IEEE J. of Robotics and Automation*, RA-3(3), June 1987.

[LPJMO92] T. Lozano-Pérez, J. Jones, E. Mazer, and P. O'Donnell. *HANDEY: A Robot Task Planner.* The MIT Press, Cambridge, Massachusetts, 1992.

[QC96] C. Qin and S. Cameron. Motion planning for manipulators using an expert system. In preparation, 1996.

[QCM95] C. Qin, S. Cameron, and A. McLean. Towards efficient motion planning for manipulators with complex geometry. In *Proc. IEEE International Symposium on Assembly an Task Planning*, pages 207–212, Pittsburgh, USA, August 1995.

[Rei79] J.H. Reif. Complexity of the mover's problem and generalizations. In *20th IEEE Symposium on Foundations of Computer Science*, pages 421–427, 1979.

[SS83] J.T. Schwartz and M. Sharir. On the 'piano movers' problem: Ii. general techniques for computing topological properties of real algebraic manifolds. In *Advances in Applied Mathematics*, pages 298–351, Academic Press, 1983.