Free-Climbing with a Multi-Use Robot

Tim Bretl[1], Stephen Rock[1], Jean-Claude Latombe[2], Brett Kennedy[3], and Hrand Aghazarian[3]

[1] Department of Aeronautics and Astronautics, Stanford University
{tbretl,rock}@stanford.edu
[2] Computer Science Department, Stanford University
latombe@cs.stanford.edu
[3] Jet Propulsion Laboratory, California Institute of Technology
{bkennedy,hrand}@helios.jpl.nasa.gov

Abstract. This paper presents a new four-limbed robot, LEMUR IIb (Legged Excursion Mechanical Utility Rover), that can free-climb vertical rock surfaces. This robot was designed to have a number of capabilities in addition to climbing (e.g., assembly, inspection, maintenance, transport, intervention) and to be able to traverse a variety of other types of terrain (e.g., roads, talus, dirt, urban rubble). To maximize its flexibility in this regard, LEMUR IIb will need to exploit sophisticated control, planning, and sensing techniques in order to climb, rather than rely on specific hardware modifications. In particular, this paper describes a new algorithm for planning safe one-step climbing moves, which has already enabled LEMUR IIb to climb an indoor, near-vertical surface with small, arbitrarily distributed, natural features. To the authors' knowledge, this is the first experimental demonstration of a multi-use, multi-limbed robot climbing such terrain using only friction at contact points (i.e., free-climbing).

1 Introduction

Various types of robots that climb vertical surfaces have been created previously. These include adhesive robots that "stick" to a featureless, flat or smoothly curved surface by using specific end-effectors (e.g., suction cups and pads [8,16,18,21,24], or magnets [9,10]), robots whose end-effectors match engineered features of the environment (e.g. pegs [4], peg-holes [23], fences or porous materials [25], handrails or bars [2,3], and poles [1,19]), and robots designed to climb within pipes and ducts [17,20,26]. Each of these robots was designed for a particular vertical environment, and relies on its specific hardware design in order to climb.

We focus instead on enabling multi-use robots of more general hardware design to climb. We consider robots with a small number of articulated limbs. We do not distinguish between limbs, and call the end-point of each one a *hand*. To climb vertical terrain, the robot must go through a continuous sequence of configurations satisfying certain constraints (e.g., equilibrium, collision, joint-torque limits). At each configuration, some of the robot's hands are in contact with the terrain – a surface with small, arbitrarily distributed

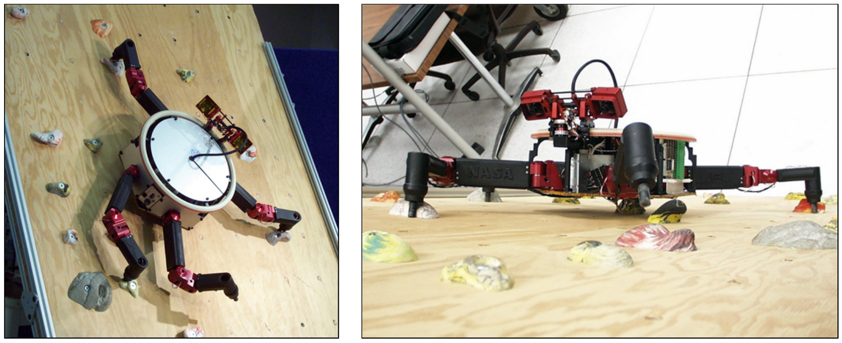

Fig. 1. LEMUR IIb, a multi-use robot capable of free-climbing.

features (e.g., protrusions or holes) called *holds*. During a *one-step* motion, the robot brings one hand to a new hold while using frictional contacts at other hands and internal degrees of freedom (DOF's) to maintain equilibrium. A *multi-step* motion is a sequence of one-step motions.

Our motivation for this approach is the ultimate development of flexible, intervention-capable, multi-limbed robots that can navigate through many different types of terrain. Potential applications include search-and-rescue, surveillance, personal assistance, and planetary exploration. The design of these robots is still critical; however, problems of motion and manipulation in specific environments are addressed by control, planning, and sensing techniques rather than hardware modifications.

In particular, our recent work has focused on the problem of careful foot-placement and trajectory generation for multi-limbed robots, which is necessary on steep, irregular terrain. Other works that have addressed the foot-placement problem make assumptions not valid for climbing (e.g., massless legs, frictionless surfaces, strictly horizontal foot-placements) [5,13,15]. Previously, we presented a fast planner to compute one-step climbing moves for multi-limbed robots, and demonstrated this planner in simulation [7]. Here, we apply our planner to enable a real, multi-use robot (LEMUR IIb) to climb a near-vertical, artificial rock surface. Our experimental results demonstrate the feasibility of free-climbing with such a robot, and have a number of implications for future development.

2 Experimental setup

2.1 Robot and terrain

LEMUR IIb consists of four identical limbs attached to a circular chassis, with a total mass of 7 kg (Figs. 1 and 2). Each limb contains three revolute joints, providing two in-plane (yaw) and one out-of-plane (pitch) degrees

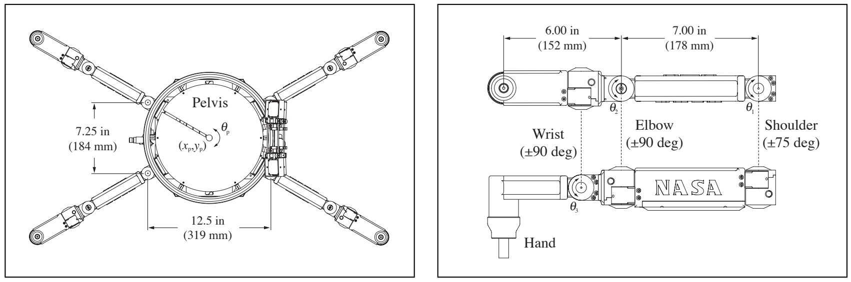

Fig. 2. A schematic diagram of LEMUR IIb, showing dimensions and notation.

of freedom (DOF's). Each joint has an identical drive-train, capable of a maximum continuous torque of 5.0 N-m and a maximum speed of 45 deg/s. Each end-effector is a single peg wrapped in high-friction rubber. LEMUR IIb can be field-operated, with on-board batteries, processing (using a PC104 architecture), and sensors (including a swiveling stereo camera pair, a 6-axis force/torque sensor at each shoulder, a 3-axis accelerometer, and joint angle encoders). When untethered, the robot's symmetry allows it to climb at an arbitrary orientation.

The terrain climbed by LEMUR IIb in our tests is an indoor, near-vertical, planar surface. This surface is covered with small, artificial rock features (holds) exactly as are indoor "climbing gyms" for human climbers. These holds are of arbitrary size and shape. The robot only uses friction to keep contact with holds – this requires careful placement of its center of mass.

2.2 Problem statement

In each experiment, LEMUR IIb is initially placed at an arbitrary, statically-stable configuration on the climbing surface. The robot is then commanded to grasp a particular, distant hold. Typically, this goal hold will be unreachable while maintaining its initial set of contact points, so the robot will have to make a multi-step climbing motion. The challenge is to complete such a motion (autonomously) without falling.

2.3 Scope and limitations

In order to focus on the one-step planning algorithm, we make several simplifying assumptions. First, the location and friction characteristics of each hold are identified manually and pre-surveyed. Also, in our current implementation we maintain LEMUR IIb's chassis parallel to and at a fixed distance from the climbing surface, and use the out-of-plane DOF in each limb only to make or break contact with features. We do not exploit momentum or dynamic movements – at this stage in our research, the robot's motion is slow enough to be assumed quasi-static. We allow the robot to contact the terrain only with its hands (i.e., no "whole-arm" manipulation). Finally, although each one-step motion is planned autonomously (see Section 3), in the experiments described in this paper the user must provide each one-step motion goal (i.e., which hold to grab or release next) along the multi-step path.

There are two other current limitations of our hardware system. First, although joint-angle encoders allow measurement of the robot's internal configuration, the vision sensors are not yet able to provide the location of the robot relative to holds (i.e., global motion is executed in open-loop). Second, for convenience, position-based control rather than hybrid force-motion control (e.g., as in [12]) is currently used to control the robot's configuration (i.e., no attempt is made to sense or control contact forces). We are in the process of correcting these two limitations; in the meantime, the early success of our climbing experiments is a testament both to the importance and effectiveness of our planning algorithm, and to the quality of the existing position-based control system.

3 One-step motion planning

3.1 Model and notation

We call the robot's circular chassis the *pelvis*. In each limb, the first joint (nearest the pelvis) is called the *shoulder*, the second joint is called the *elbow*, and the third (out-of-plane) joint is called the *wrist*. Assuming that the pelvis moves at a fixed distance parallel to the wall, any configuration of the robot is defined by 15 parameters: the position/orientation (x_p, y_p, θ_p) of the pelvis and the joint angles $(\theta_1, \theta_2, \theta_3)$ of each limb (Fig. 2).

Any point on the terrain is a potential contact – either on the contour of a continuous rock feature, or on the planar climbing surface. We assume that a discrete number of useful contacts have been identified; these points are called *holds*. For LEMUR IIb, all holds lie on an inclined plane but can have arbitrary orientation, so each is defined by a 2-D point (x_i, y_i) and a 3-D direction $\boldsymbol{\nu}_i$. Holds at which hands are in contact are the *supporting holds*. We model friction at contacts by Coulomb's law.

When climbing, LEMUR IIb always maintains either three or four supporting holds. The set of supporting holds is a *stance*, denoted σ – to differentiate between 3-hold and 4-hold stances, we write $\sigma 3$ and $\sigma 4$. The linkage between the supporting holds – containing the pelvis and either three or four limbs – is called the *contact chain*. When only three supporting holds are used, the fourth limb is the *free limb*.

Because of the closed-chain constraint, the robot's continuous motion with four supporting holds occurs on a 3-D manifold $C_{\sigma 4}$ in the robot's configuration space. With three supporting holds, motion occurs on a 6-D manifold $C_{\sigma 3}$. This motion is subject to four additional constraints: quasi-static equilibrium, joint angle limits, joint torque limits, and collision. The *feasible space* at a stance σ is the subset F_σ of C_σ satisfying each of these constraints. The limbs have non-negligible mass, so their motion affects the robot's equilibrium. If two points in F_σ are connected by a continuous path in F_σ, we say they are in the same *component* of F_σ. (Henceforth, a "continuous path" will always be taken to mean a continuous motion of the robot at some fixed stance, i.e., with fixed supporting holds.)

To climb upward, the robot must switch between 3-hold and 4-hold stances. Two stances $\sigma 3$ and $\sigma 4$ are *adjacent* if $\sigma 4 = \sigma 3 \cup \{i\}$ for some hold i. The robot can only switch between adjacent stances σ and σ' (i.e., place or remove a hand) at points $q_t \in F_\sigma \cap F_{\sigma'}$. We call such points *transition points*. Given a start configuration $q_s \in F_\sigma$ at a stance σ, we say that a component of the feasible space $F_{\sigma'}$ at an adjacent stance σ' is *reachable* if there is a continuous path connecting q_s to a transition point in that component. This path is a *one-step motion*. Examples of one-step motions are shown in Figs. 3 and 4.

In the experiments described in this paper, the user specifies each one-step motion goal, i.e., for each stance σ, the user specifies a desired adjacent stance σ'. The *one-step planning problem* is to determine whether this stance σ' is reachable, and if so, to construct a continuous path to reach it.

3.2 Algorithm

As described above, the robot's motion takes place in either a 3-D or 6-D space subject to multiple constraints. Since many one-step motion queries will be made along a multi-step path, it is computationally impractical to determine exactly whether an adjacent stance σ' is reachable from each start configuration $q_s \in F_\sigma$. Therefore, we use an approximate method: first, we sample transition points $q_t \in F_\sigma \cap F_{\sigma'}$, then we try to construct a continuous path from q_s to each q_t.

Assume that $\sigma' = \sigma \cup \{i\}$ for some hold i, so σ is a stance with three supporting holds. Then to sample q_t, we search for points in $F_{\sigma'}$ – a 3-D space – and check that these points are also in F_σ. Note that if $q \in F_{\sigma'}$, then we can verify $q \in F_\sigma$ only by checking the equilibrium constraint at stance σ. Therefore, feasible q_t can be found very quickly. Instead, it is the path-planning problem that dominates computation time.

Because of its flexibility and speed, the Probabilistic-RoadMap (PRM) approach, or one of its variants, is widely used for this type of problem [11]. We use a lazy, bi-directional PRM (as in [22]), exploiting the efficient check of quasi-static equilibrium with arbitrary frictional contacts we described in [7]. However, these planners tend to lose efficiency when the feasible subset of the configuration space contains narrow passages or is subject to closed-chain constraints, as these features are difficult to sample. In [7], we showed that equilibrium constraints for simple climbing robots create narrow passages that occur in a low-dimensional subspace, which can be sampled separately to speed up planning. We have used these results to generate efficient PRM sampling strategies for LEMUR IIb, allowing fast one-step planning.

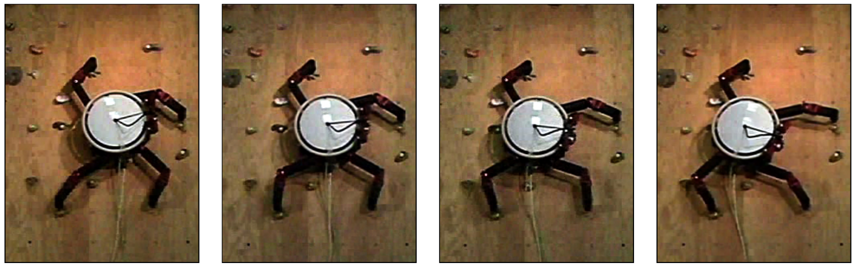

Fig. 3. A one-step motion with a 4-hold stance, to remove the bottom right hand.

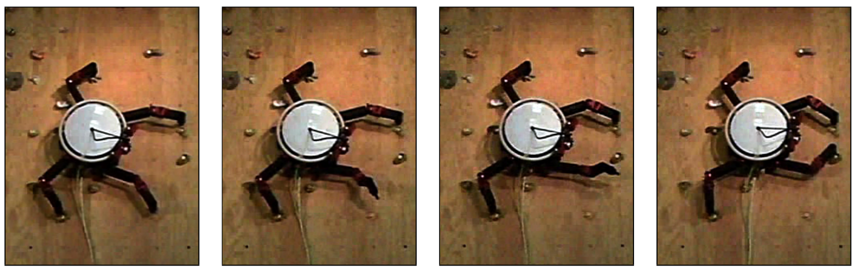

Fig. 4. A one-step motion with a 3-hold stance, to place the bottom right hand.

These heuristics primarily consist of two modifications to the algorithm presented in [7]. First, configurations of the robot with $\theta_2 = 0$ in any limb of the contact chain (straight-limb configurations) are initially sampled explicitly, since the feasible space F_σ is most likely to be disconnected on these manifolds. During subsequent exploration of F_σ, we only attempt to connect those pairs of configurations with identical elbow-bends (i.e., when θ_2 in each limb of the contact chain has identical sign for both configurations).

Second, at stances with three supporting holds, we add several deterministic configurations corresponding to each sampled one. Given a sampled configuration of the contact chain, we explicitly calculate the configurations of the free limb that bring the robot closest to infeasibility with respect to the equilibrium constraint (in general, there are two). If these points are in F_σ, they are added to the roadmap as well. This strategy tends to approximate the analytical decomposition technique described in [7]. (Note the interesting similarity between our heuristics and those based on "manipulability" [14].)

4 Experimental Results

4.1 Summary

We applied our one-step planner to generate climbing motions for LEMUR IIb, which were subsequently executed by the real robot. Snapshots from one

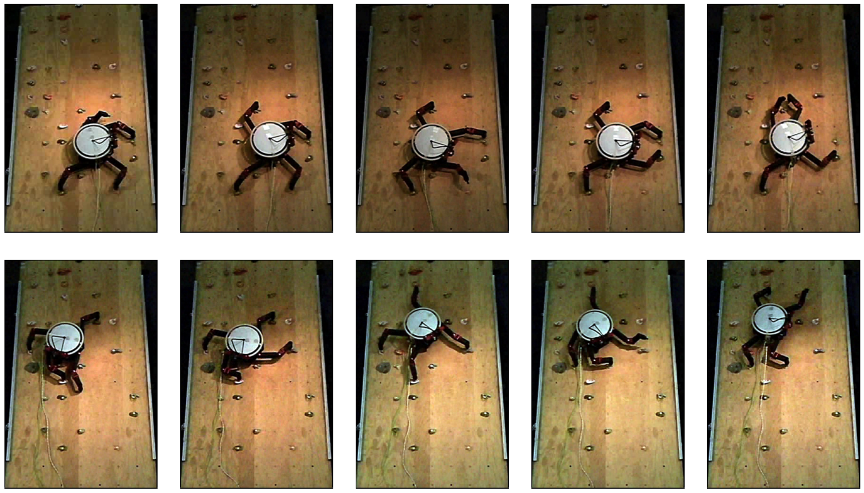

Fig. 5. Snapshots of LEMUR IIb climbing a near-vertical surface covered with artificial rock features.

climb (see also a video at `http://arl.stanford.edu/~tbretl`), taking the robot from bottom to top of the climbing surface, are shown in Fig. 5. In this example, the one-step planner generated, on-line, each of 88 one-step motion trajectories forming the multi-step path. Our experiments demonstrate the feasibility of autonomous free-climbing with a multi-limbed robot, given sophisticated planning and a basic level of control and sensing. Even using tele-operation, these experiments would have been difficult or even impossible without our planner – one-step motions are very challenging to construct by hand, due to the interplay between equilibrium and joint-angle constraints.

4.2 Implications for multi-step planning

Requirements. In addition to demonstrating the usefulness of autonomous one-step planning for free-climbing robots, our experiments clearly show a need for autonomous multi-step planning as well.

For example, it took one of the authors two full days to design the terrain and a sequence of feasible one-step motion goals for the 88-move path shown in Fig. 5, despite the fact that he is a human rock-climber. The reason is that multi-step planning is a hard combinatorial problem, even given a fast one-step planner. There are over 20000 feasible stances for LEMUR IIb in the terrain shown, populating a graph search of moderate breadth and high depth. (The situation becomes even worse if holds are sampled from a continuous environment rather than specified and pre-surveyed.) One might think that because the author, an experienced human climber, is skilled at constructing multi-step paths for himself, and because there is an obvious

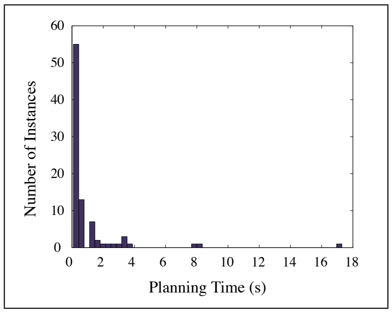

Fig. 6. The distribution of one-step planning times along a multi-step path consisting of 88 one-step motions. Minimum planning time was 0.09 s, maximum was 17.3 s. Mean planning time was 1.02 s, but over 75% of the one-step motions were computed quicker than average.

similarity between LEMUR IIb's motion and his own (also see [7]), he would have an advantage when planning for the robot. However, it is difficult to map climbing motions across morphologies; joint-angle constraints and the effect of body position on the center of mass are both much different.

Likewise, consider the large number of steps needed for LEMUR IIb to climb a relatively short distance: the robot has a reach of 1 m, but to climb a distance of only 2 m (as in Fig. 5) took 88 moves. The shortest possible path with the same multi-step goal consists of about 76 moves, still quite large. In fact, it generally takes 4-12 one-step motions for the robot even to climb to the limit of its tactile sensor range, given an initial stance. This suggests that multi-step planning, rather than reactive motion, is required.

Approach. Our experiments also have strongly influenced our current work in designing an autonomous multi-step planner. For example, the distribution of planning times for the multi-step path in Fig. 5 is shown in Fig. 6. Most one-step moves were planned very quickly (less than 1.0 s on a 1GHz PowerPC). However, several difficult moves took more time (more than 5.0 s), even as much as 17.3 s. This reveals an important issue when searching for multi-step plans, in which many more potential one-step motions are explored than ultimately used. Since PRM planners lack a formal stopping criterion, how much time T_{max} should be spent on each one-step motion query before it is declared infeasible? In this example, suppose $T_{max} = 2.0\ s$ (twice the mean value) – infeasible one-step motions are rejected quickly, but several difficult moves (15% of the multi-step path) likely are not found. Alternatively, suppose $T_{max} = 20.0\ s$ (greater than the maximum value) – now, all feasible one-step motions likely are found, but every infeasible query takes ten times longer, drastically increasing total search time. A similar problem exists for non-gaited motion of humanoid robots and for manipulation planning. We are currently investigating several possible solutions to this problem (see [6]).

4.3 Other lessons learned

Several other lessons were learned as a result of our experiments, that seem obvious in hindsight. For example, most experimental failures occurred when the robot's hands (rigid pegs) rolled along hold contours. Resulting cumulative errors led either to missed hand placements or to torque overloads from jamming. This problem could be addressed in software, but a hardware solution involving passive, articulated hand endpoints (rather than rigid pegs) should prove more practical.

Additional issues raised include future control and sensing requirements (e.g., the need for hybrid force-motion control, the integration of tactile sensing, or the required precision and range of visual sensors), hardware modifications (e.g., whether skewed rather than symmetric joint-angle limits might increase mobility), and user interface design (e.g., how to communicate risk, particularly when one-step motions are "almost feasible").

5 Conclusion

This paper presented a new multi-limbed, multi-use robot, LEMUR IIb. The ability of this robot to free-climb vertical rock surfaces, using a previously presented one-step planning algorithm, was experimentally demonstrated. These experiments revealed a number of additional issues, many concerned with the need for and requirements of an autonomous multi-step planner. Many other challenges remain to be addressed (e.g., integration of local visual and tactile sensing, implementation of hybrid force-motion control, and consideration of dynamic motion).

Acknowledgments. T. Bretl is supported by a Herbert Kunzel Fellowship. The authors would like to thank T. Miller, D. Black-Schaffer, and K. Hauser for their contributions.

References

1. M. Almonacid, R. Saltarén, R. Aracil, and O. Reinoso. Motion planning of a climbing parallel robot. *IEEE Tr. on Rob. and Aut.*, 19(3):485–489, 2003.
2. H. Amano, K. Osuka, and T.-J. Tarn. Development of vertically moving robot with gripping handrails for fire fighting. In *IEEE/RSJ Int. Conf. on Int. Rob. and Sys.*, Maui, HI, 2001.
3. C. Balaguer, A. Giménez, J. Pastor, V. Padrón, and M. Abderrahim. A climbing autonomous robot for inspection applications in 3d complex environments. *Robotica*, 18:287–297, 2000.
4. D. Bevly, S. Dubowsky, and C. Mavroidis. A simplified cartesian-computed torque controller for highly geared systems and its application to an experimental climbing robot. *ASME J. of Dynamic Systems, Measurement, and Control*, 122(1):27–32, 2000.
5. J.-D. Boissonnat, O. Devillers, and S. Lazard. Motion planning of legged robots. *SIAM J. on Computing*, 30(1):218–246, 2000.
6. T. Bretl, S. Lall, J.-C. Latombe, and S. Rock. Multi-step motion planning for free-climbing robots. To appear in *WAFR*, Zeist, The Netherlands, 2004.
7. T. Bretl, J.-C. Latombe, and S. Rock. Toward autonomous free-climbing robots. In *Int. Symp. of Rob. Res.*, Siena, Italy, 2003.

8. H. Dulimarta and R. L. Tummala. Design and control of miniature climbing robots with nonholonomic constraints. In *World Congress on Intelligent Control and Automation*, Shanghai, P.R.China, Jun 2002.
9. J. C. Grieco, M. Prieto, M. Armada, and P. G. de Santos. A six-legged climbing robot for high payloads. In *IEEE Int. Conf. on Cont. App.*, Trieste, Italy, 1998.
10. S. Hirose, A. Nagabuko, and R. Toyama. Machine that can walk and climb on floors, walls, and ceilings. In *ICAR*, pages 753–758, Pisa, Italy, 1991.
11. L. E. Kavraki, P. Svetska, J.-C. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Tr. on Rob. and Aut.*, 12(4):566–580, 1996.
12. O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE J. of Rob. and Aut.*, RA-3(1), 1987.
13. J. J. Kuffner, Jr., K. Nishiwaki, S. Kagami, M. Inaba, and H. Inoue. Motion planning for humanoid robots under obstacle and dynamic balance constraints. In *IEEE Int. Conf. on Rob. and Aut.*, 2001.
14. P. Leven and S. Hutchinson. Using manipulability to bias sampling during the construction of probabilistic roadmaps. *IEEE Tr. on Rob. and Aut.*, 19(6):1020–1026, 2003.
15. A. Madhani and S. Dubowsky. The force workspace: A tool for the design and motion planning of multi-limb robotic systems. *ASME Journal of Mechanical Design*, 119(2):218–224, 1997.
16. A. Nagakubo and S. Hirose. Walking and running of the quadruped wall-climbing robot. In *IEEE Int. Conf. on Rob. and Aut.*, volume 2, pages 1005–1012, 1994.
17. W. Neubauer. A spider-like robot that climbs vertically in ducts or pipes. In *Int. Conf. on Int. Rob. and Sys.*, pages 1178–1185, Munich, Germany, 1994.
18. M. Rachkov. Control of climbing robot for rough surfaces. In *Int. Workshop on Robot Motion and Control*, pages 101–105, 2002.
19. Z. Ripin, T. B. Soon, A. Abdullah, and Z. Samad. Development of a low-cost modular pole climbing robot. In *TENCON*, volume 1, pages 196–200, Kuala Lumpur, Malaysia, 2000.
20. T. Roßmann and F. Pfeiffer. Control of an eight legged pipe crawling robot. In *Int. Symp. on Experimental Robotics*, pages 353–346, 1997.
21. S. W. Ryu, J. J. Park, S. M. Ryew, and H. R. Choi. Self-contained wall-climbing robot with closed link mechanism. In *IEEE/RSJ Int. Conf. on Int. Rob. and Sys.*, Maui, HI, 2001.
22. G. Sánchez and J.-C. Latombe. On delaying collision checking in PRM planning: Application to multi-robot coordination. *Int. J. of Rob. Res.*, 21(1):5–26, 2002.
23. Y. Xu, H. Brown, M. Friendman, and T. Kanade. Control system of the self-mobile space manipulator. *IEEE Tr. on Cont. Sys. Tech.*, 2(3):207–219, 1994.
24. W. Yan, L. Shuliang, X. Dianguo, Z. Yanzheng, S. Hao, and G. Xuesban. Development and application of wall-climbing robots. In *IEEE Int. Conf. on Rob. and Aut.*, Detroit, MI, 1999.
25. M. Yim, S. Homans, and K. Roufas. Climbing with snake-robots. In *IFAC Workshop on Mobile Robot Technology*, Jejudo, Korea, 2001.
26. A. Zagler and F. Pfeiffer. MORITZ a pipe crawler for tube junctions. In *IEEE Int. Conf. on Rob. and Aut.*, pages 2954–2959, Taipei, Taiwan, 2003.