# Reactive Mobile Manipulation with Legged Robots

Vasileios Vasilopoulos
GRASP Lab, University of Pennsylvania
vvasilo@seas.upenn.edu

## I. Motivation

Recent advances in the field of legged robotics [23, 19, 55, 33, 7], including several demonstrations from companies [10, 16, 1, 42], show that legged robots are becoming better at traversing rough terrains and environments. Despite these advances, legged robots are still mostly used as locomotion research platforms [40], and their limited commercial applications are restricted to inspection [2], security, and “last-meter” delivery [3], where interaction with the environment is not needed and rather avoided. Given the inherent ability of legged robots to use their limbs as general-purpose manipulators, this research seeks to demonstrate ways of accomplishing *tasks* with legged robots that require interaction with their surroundings, such as rearrangement planning [47, 18], or navigation among movable obstacles [36] to escape a dangerous situation or help trapped people in search-and-rescue missions.

Focusing on the task planning literature, it can be seen that existing solutions are either *task-specific*, *environment-specific* or *platform-specific*, and are typically not accompanied by any formal proofs of correctness. For example, Task-and-Motion-Planning (TAMP) methods [21, 39] or classical AI methods can find a particular (often optimal [50]) solution to a task at hand, but require good prior knowledge [28], and do not generalize well in the presence of unanticipated conditions. Similarly, recent developments in Deep Reinforcement Learning [37] have yielded impressive results [40, 29], but are tied to a specific platform for which an abundance of data is needed.

Instead, as shown in Fig. 1, we seek to come up with a modular, and task, environment and platform independent architecture (inherently unavailable in end-to-end deep learning schemes), with formal correctness conclusions based on some underlying assumptions about the environment, where an offline deliberative layer for task planning works closely with an online reactive module, that uses exteroception and handles environment uncertainties. This reactive module communicates with a platform-specific gait layer, comprised of a set of simple dynamical primitives, that realizes the commands from the reactive layer in a way that is meaningful for the robot. Each of these independent layers comes with provable guarantees of optimality (for the deliberative layer), collision avoidance and convergence (for the reactive layer) or low-level performance, expressed as “symbols” of energy landscapes composed either in parallel [14, 41] or sequentially [26, 13] (for the gait layer), offering the chance of generalization across multiple mobile manipulators (legged or wheeled).

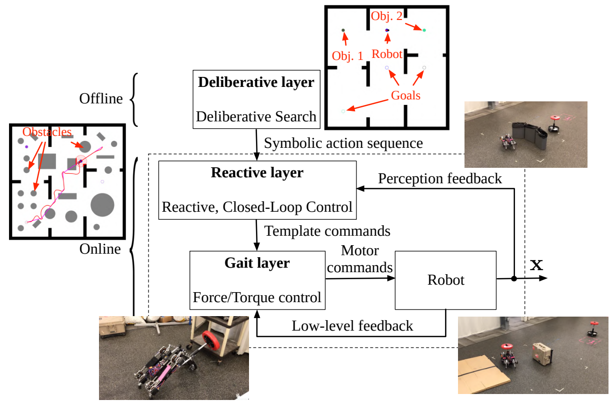


Fig. 1. The proposed hierarchical control structure. In the deliberative layer, an offline high-level planner outputs a sequence of symbolic actions, that are executed online using a reactive controller that incorporates perception, modifies the high-level plan appropriately to account for unanticipated conditions and obstacles, and issues abstract velocity and gripper commands (see Section III). The low-level gait layer uses these commands to call out appropriately parameterized joint-level feedback controllers for the robotic platform.

## II. Prior Work

Although the problem of using a higher-level planner to inform subgoals of a lower-level planner has been examined previously, most work has focused on ad hoc abstractions that perform well empirically. For example, Wolfe et al. [54] use a task hierarchy to guide the search for a low-level plan by expanding high-level plans in a best-first way. Berenson et al. [6] and Konidaris et al. [25] use specific formulations of hierarchy without guaranteeing optimality. Kaelbling and Lozano-Perez [21] avoid the computational cost by committing to decisions at a high level of abstraction. On the other hand, Vega-Brown and Roy [50] provided a further step towards tractable planning that incorporated complex kinematic constraints, and showed how to use angelic semantics [27] to guarantee hierarchical optimality [51].

Also, recent advances in the theory of sensor-based reactive navigation [4] and its application to wheeled [5] and legged [45] robots promote its central role in provably correct architectures for complicated mobile manipulation tasks [46, 47]. The advance of the new theory [4] over prior sensor-based collision avoidance schemes [44, 43, 9, 8, 12, 15, 20, 30, 38] was the additional guaranteed convergence to a designated goal which had theretofore only been established for reactive planners possessing substantial prior knowledge about the environment [35, 34]. We seek to build on such methods that trade away prior knowledge for the presumption of geometric simplicity, expand them to geometrically more interesting

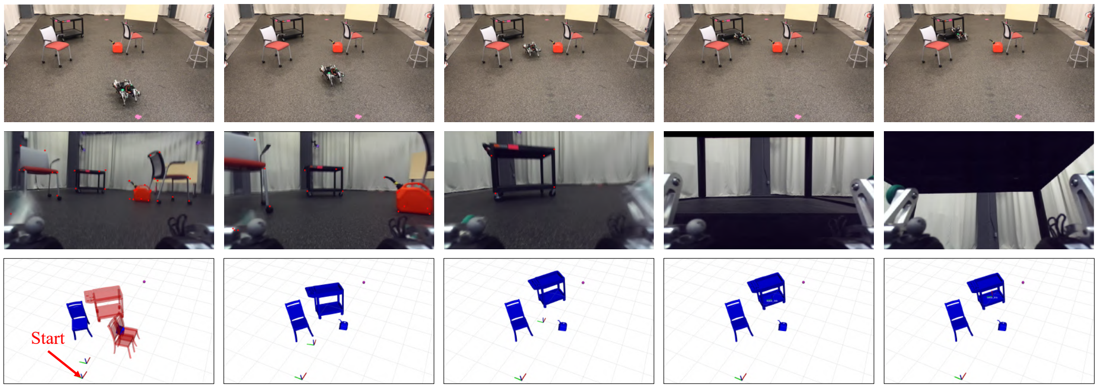


Fig. 2. Minitaur using the reactive control architecture in [49], also shown in the reactive layer of Fig. 1, and its onboard sensors, to avoid semantically tagged and other unknown obstacles, successfully localize an object of interest (cart) and use mobile manipulation primitives [41] to jump and mount it.

environments, and use them in parallel with recent methods that show how to compositionally perform complex mobile manipulation maneuvers with legged robots [41].

## III. Problem Statement

For our work, we use the Minitaur [23] robot and assume it operates in a closed and compact workspace whose boundary is known. The robot is tasked to either move to a predefined location that is not accessible without manipulating its environment, or move each of $n$ movable objects from their initial configuration to a user-specified goal configuration. We assume that both the initial configuration and the target configuration are known. In addition to the known boundary of the workspace, the workspace is cluttered by an unknown number of fixed, disjoint, potentially non-convex obstacles.

For (reactive) planning purposes, Minitaur is modeled as a first-order, nonholonomically-constrained, disk-shaped robot. The robot is assumed to have access to its state (e.g., through legged state estimation methods [17]), and to possess a LIDAR for local obstacle avoidance and a camera for familiar object/obstacle recognition, using either deep learning perception schemes [31] or conventional methods like AprilTags [53]. It is also assumed to use a gripper for moving objects, which can be either engaged or disengaged. Of course, Minitaur is only an imperfect unicycle [45] and does not actually possess a gripper; it has to coordinate its limbs and walk while following a path, avoid an obstacle, jump, or lock an object in place. Hence, the reactive planner's commands must be translated to suitable low-level commands on the robot's joints.

The aforementioned description imposes the hierarchical structure shown in Fig. 1 and the following problem decomposition into the complementary sub-problems:

1) In the *deliberative layer*, find a *symbolic plan*, i.e., a sequence of symbolic actions whose successful implementation is guaranteed to complete the task, assuming idealized perfect prior knowledge.
2) In the *reactive layer*, implement each of the symbolic actions by finding appropriate commands according to the robot's equations of motion, while avoiding the perceived obstacles (unanticipated by the deliberative planner) encountered along the way.
3) In the *gait layer*, use a hybrid dynamical systems framework with simple guard conditions to choose between constituent gaits, providing an abstract interface to the reactive layer, regardless of the state of the robot/objects.

## IV. Contributions

Based on the aforementioned description, we suggest with formal arguments and empirical demonstration [47] the effectiveness of a hierarchical control structure for a highly dynamic physical system, shown in Fig. 1. We believe this is the first provably correct deliberative/reactive planner to engage an unmodified general purpose mobile manipulator in physical rearrangements of its environment. We are able to accomplish a variety of tasks, including desired assemblies of objects with size comparable to the robot's size among unanticipated conditions and obstacles [46], navigation among movable obstacles, and strategic escapes by exploiting and manipulating the robot's environment [52]. To this end, we develop the mobile manipulation maneuvers to accomplish each task at hand [47], successfully anchor the useful kinematic unicycle template to control the highly dynamic Minitaur robot [45] and integrate perceptual feedback with low-level control to coordinate the robot's movement [47], as shown in Fig. 2.

At the same time, this research also exploits recent developments in semantic SLAM [11] and object pose and triangular mesh extraction using convolutional neural net architectures [31, 22, 24] to provide an avenue for incorporating partial prior knowledge within a deterministic framework well suited to existing vector field planning methods [4]. In this way, we are able to guarantee collision avoidance and convergence to the designated goal for both a differential drive robot and a differential drive robot gripping and manipulating objects, in a workspace cluttered with completely unknown convex obstacles [46], "familiar", online recognizable non-convex obstacles [49, 48], or completely unknown non-convex obstacles [47] that obey specific "length-scale" geometric assumptions [32].

Finally, in order to encourage the application of our methods, we are planning to release accompanying software with an open-source implementation of our reactive mobile manipulation algorithms in C++ and Python, with ROS wrappers.

## References


[1] ANYbotics. ANYmal Let Robots Go Anywhere, 2018. URL https://www.youtube.com/watch?v=m1-s8iOJaI4.

[2] ANYbotics. World's First Autonomous Offshore Robot, 2018. URL https://youtu.be/DzTvIPrt0DY.

[3] ANYbotics. Last-Meter Robotic Package Delivery with ANYmal (CES 2019, ANYbotics & Continental), 2019. URL https://www.youtube.com/watch?v=v3g5xp5Kr2g.

[4] O. Arslan and D. E. Koditschek. Sensor-Based Reactive Navigation in Unknown Convex Sphere Worlds. In *The 12th International Workshop on the Algorithmic Foundations of Robotics*, 2016.

[5] O. Arslan and D. E. Koditschek. Sensor-Based Reactive Navigation in Unknown Convex Sphere Worlds. *International Journal of Robotics Research*, 38(2-3):196–223, 2018.

[6] D. Berenson, S. Srinivasa, D. Ferguson, and J. Kuffner. Manipulation planning on constraint manifolds. In *IEEE International Conference on Robotics and Automation*, 2009.

[7] G. Bledt, M. J. Powell, B. Katz, J. D. Carlo, P. M. Wensing, and S. Kim. MIT Cheetah 3: Design and Control of a Robust, Dynamic Quadruped Robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2245 – 2252, 2018.

[8] J. Borenstein and Y. Koren. Real-time obstacle avoidance for fast mobile robots. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):1179–1187, 1989.

[9] J. Borenstein and Y. Koren. The vector field histogram-fast obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation*, 7(3):278–288, 1991.

[10] Boston Dynamics. SpotMini Autonomous Navigation, 2018. URL https://www.youtube.com/watch?v=Ve9kWX_KXus/.

[11] S. L. Bowman, N. Atanasov, K. Daniilidis, and G. J. Pappas. Probabilistic data association for semantic SLAM. In *IEEE International Conference on Robotics and Automation*, pages 1722–1729, 2017.

[12] O. Brock and O. Khatib. High-speed navigation using the global dynamic window approach. In *IEEE International Conference on Robotics and Automation*, pages 341–346, 1999.

[13] R. Burridge, A. Rizzi, and D. Koditschek. Sequential Composition of Dynamically Dexterous Robot Behaviors. *The International Journal of Robotics Research*, 18:534–555, 1999.

[14] A. De and D. E. Koditschek. Vertical hopper compositions for preflexive and feedback-stabilized quadrupedal bounding, pacing, pronking, and trotting. *The International Journal of Robotics Research*, 37(7):743–778, 2018.

[15] P. Fiorini and Z. Shiller. Motion Planning in Dynamic Environments Using Velocity Obstacles. *The International Journal of Robotics Research*, 17(7):760–772, 1998.

[16] Ghost Robotics. Ghost Robotics: Q-UGV Capabilities Summary, 2018. URL https://www.youtube.com/watch?v=kfQIFw5kbOY.

[17] R. Hartley, J. Mangelson, L. Gan, M. G. Jadidi, J. M. Walls, R. M. Eustice, and J. W. Grizzle. Legged robot state-estimation through combined forward kinematic and preintegrated contact factors. In *IEEE International Conference on Robotics and Automation*, pages 1–8, 2018.

[18] J. E. Hopcroft, J. T. Schwartz, and M. Sharir. On the Complexity of Motion Planning for Multiple Independent Objects; PSPACE-Hardness of the "Warehouseman's Problem". *The International Journal of Robotics Research*, 3(4):76–88, 1984.

[19] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, and M. Hoepflinger. ANYmal - a highly mobile and dynamic quadrupedal robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 38–44, 2016.

[20] A. M. Johnson, M. T. Hale, G. C. Haynes, and D. E. Koditschek. Autonomous legged hill and stairwell ascent. In *IEEE International Symposium on Safety, Security, and Rescue Robotics*, pages 134–142, 2011.

[21] L. P. Kaelbling and T. Lozano-Perez. Hierarchical task and motion planning in the now. In *IEEE International Conference on Robotics and Automation*, pages 1470–1477, 2011.

[22] A. Kar, S. Tulsiani, J. Carreira, and J. Malik. Category-specific object reconstruction from a single image. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1966–1974, 2015.

[23] G. Kenneally, A. De, and D. E. Koditschek. Design Principles for a Family of Direct-Drive Legged Robots. *IEEE Robotics and Automation Letters*, 1(2):900–907, 2016.

[24] C. Kong, C.-H. Lin, and S. Lucey. Using Locally Corresponding CAD Models for Dense 3D Reconstructions from a Single Image. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 5603–5611, 2017.

[25] G. Konidaris, L. Kaelbling, and T. Lozano-Perez. Constructing Symbolic Representations for High-Level Planning. In *AAAI*, 2014.

[26] A. Majumdar and R. Tedrake. Funnel Libraries for Real-Time Robust Feedback Motion Planning. *The International Journal of Robotics Research*, 36(8):947–982, 2017.

[27] B. Marthi, S. Russell, and J. Wolfe. Angelic Hierarchical Planning: Optimal and Online Algorithms. In *International Conference on Automated Planning and Scheduling*, 2008.

[28] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL: The Planning Domain Definition Language. *Technical*

*Report CVC TR98003/DCS TR1165*, 1998.
[29] OpenAI. OpenAI Dota 2 1v1 bot, 2017. URL https://openai.com/the-international/.
[30] A. A. Paranjape, K. C. Meier, X. Shi, S.-J. Chung, and S. Hutchinson. Motion primitives and 3D path planning for fast flight through a forest. *The International Journal of Robotics Research*, 34(3):357–377, 2015.
[31] G. Pavlakos, X. Zhou, A. Chan, K. G. Derpanis, and K. Daniilidis. 6-DoF object pose from semantic keypoints. In *IEEE International Conference on Robotics and Automation*, pages 2011–2018, 2017.
[32] R. A. Poliquin, R. T. Rockafellar, and L. Thibault. Local Differentiability of Distance Functions. *Transactions of the American Mathematical Society*, 352(11):5231–5249, 2000.
[33] A. Ramezani, J. W. Hurst, K. A. Hamed, and J. W. Grizzle. Performance Analysis and Feedback Control of ATRIAS, A Three-Dimensional Bipedal Robot. *Journal of Dynamic Systems, Measurement, and Control*, 2013.
[34] E. Rimon. *Exact robot navigation using artificial potential functions*. PhD thesis, Yale University, 1990.
[35] E. Rimon and D. E. Koditschek. Exact Robot Navigation Using Artificial Potential Functions. *IEEE Transactions on Robotics and Automation*, 8(5):501–518, 1992.
[36] J. Scholz, N. Jindal, M. Levihn, C. L. Isbell, and H. I. Christensen. Navigation Among Movable Obstacles with learned dynamic constraints. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3706–3713, 2016.
[37] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel. Trust Region Policy Optimization. *arXiv: 1502.05477*, 2015.
[38] R. Simmons. The curvature-velocity method for local obstacle avoidance. In *IEEE International Conference on Robotics and Automation*, volume 4, pages 3375–3382, 1996.
[39] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *IEEE International Conference on Robotics and Automation*, pages 639–646, 2014.
[40] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke. Sim-to-Real: Learning Agile Locomotion For Quadruped Robots. *arXiv: 1804.10332*, 2018.
[41] T. T. Topping, V. Vasilopoulos, A. De, and D. E. Koditschek. Composition of templates for transitional pedipulation behaviors. In *International Symposium on Robotics Research*, 2019.
[42] Unitree Robotics. Laikago: a four leg robot is coming to you, 2017. URL https://www.youtube.com/watch?v=d6Ja643GqL8.
[43] J. van den Berg, M. Lin, and D. Manocha. Reciprocal Velocity Obstacles for real-time multi-agent navigation. In *IEEE International Conference on Robotics and Automation*, pages 1928–1935, 2008.
[44] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha. *Reciprocal n-Body Collision Avoidance*, pages 3–19. Springer Berlin Heidelberg, 2011.
[45] V. Vasilopoulos, O. Arslan, A. De, and D. E. Koditschek. Sensor-Based Legged Robot Homing Using Range-Only Target Localization. In *IEEE International Conference on Robotics and Biomimetics*, pages 2630–2637, 2017.
[46] V. Vasilopoulos, W. Vega-Brown, O. Arslan, N. Roy, and D. E. Koditschek. Sensor-Based Reactive Symbolic Planning in Partially Known Environments. In *IEEE International Conference on Robotics and Automation*, pages 5683–5690, 2018.
[47] V. Vasilopoulos, T. T. Topping, W. Vega-Brown, N. Roy, and D. E. Koditschek. Sensor-Based Reactive Execution of Symbolic Rearrangement Plans by a Legged Mobile Manipulator. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3298–3305, 2018.
[48] V. Vasilopoulos, G. Pavlakos, S. L. Bowman, J. D. Caporale, K. Daniilidis, G. J. Pappas, and D. E. Koditschek. Technical Report: Reactive Semantic Planning in Unexplored Semantic Environments Using Deep Perceptual Feedback. *Technical Report, arXiv: 2002.12349*, 2020.
[49] V. Vasilopoulos, G. Pavlakos, K. Schmeckpeper, K. Daniilidis, and D. E. Koditschek. Reactive Navigation in Partially Familiar Planar Environments Using Semantic Perceptual Feedback. *Under review, arXiv: 2002.08946*, 2020.
[50] W. Vega-Brown and N. Roy. Asymptotically optimal planning under piecewise-analytic constraints. In *The 12th International Workshop on the Algorithmic Foundations of Robotics*, 2016.
[51] W. Vega-Brown and N. Roy. Admissible abstractions for near-optimal task and motion planning. In *The 27th International Joint Conference on Artificial Intelligence*, pages 4852–4859, 2018.
[52] W. Vega-Brown, V. Vasilopoulos, T. T. Topping, D. E. Koditschek, and N. Roy. Robust Near-optimal Task and Motion Planning. *In Prep*, 2020.
[53] J. Wang and E. Olson. AprilTag 2: Efficient and robust fiducial detection. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.
[54] J. Wolfe, B. Marthi, and S. Russell. Combined task and motion planning for mobile manipulation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 2010.
[55] D. Wooden, M. Malchano, K. Blankespoor, A. Howardy, A. A. Rizzi, and M. Raibert. Autonomous navigation for BigDog. In *IEEE International Conference on Robotics and Automation*, pages 4736–4741, 2010.