FFRob: An efficient heuristic for task and motion planning

Caelan Garrett, Tomás Lozano-Pérez, and Leslie Kaelbling

LIS Learning & Intelligent Systems

MIT CSAIL

Mobile Manipulation

Mobile Manipulation - robot moving and manipulating world to achieve goals

Applications - shipping, debris cleanup, household chores, search and rescue, …

Pick and Place Problem - grasp and transport objects to goal regions while avoiding collisions with obstacles

Geometric and kinematic constraints introduce significant complication over purely symbolic blocks world

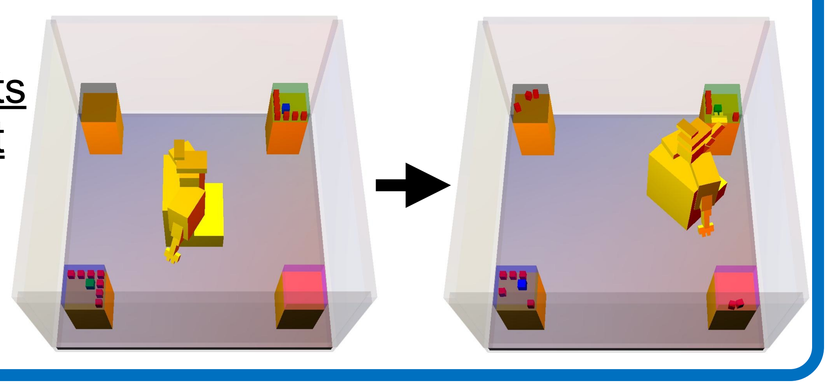

Background

Motion Planning – planning motor specific movements to navigate between two robot configurations

Task Planning – abstract, symbolic planning through states described by propositional formulas

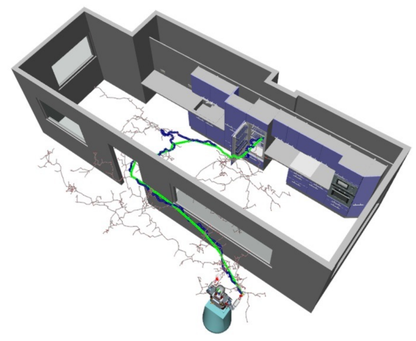

[http://h2t.anthropomatik.kit.edu/english/376.php]

Combination enables solving complex mobile manipulation problems.

Prior work combines but treats each as a "black box"

- Motion planner evaluates state as semantic attachment
- Symbolic planner heuristic guides motion planner

Our system integrates both symbolic and geometric information in heuristic search with a shared roadmap

Actions and CRG

Discretize planning problem by task driven sampling

Pick(C1, O, G, P, C2):
pre: HandEmpty, Pose(O, P), RobotConf(C1), CanGrasp(O, P, G, C2), Reachable(C1,C2)
add: Holding(O, G), RobotConf(C2)
delete: HandEmpty, RobotConf(C1)

Place(C1, O, G, P, C2):
pre: Holding(O, G), RobotConf(C1), CanGrasp(O, P, G, C2), Reachable(C1, C2)
add: HandEmpty, Pose(O, P), RobotConf(C2)
delete: Holding(O, G), RobotConf(C1)

Conditional Reachability Graph (CRG)

- Connects sampled action configurations with a Probabilistic Roadmap (PRM)
- Updates reachable configurations given robot & object positions and grasp
- Provides additional geometric inference to evaluate *Reachable(C1, C2)*
- Lazy collision caching efficiently reuses expensive geometry computations
- Shared between search and heuristic

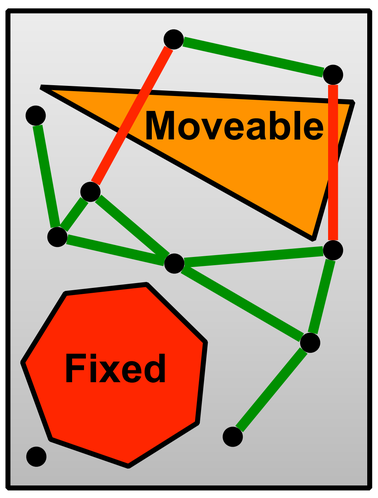

FFRob

FastForward (FF) Planner

- Enforced hill-climbing forward heuristic search
- Relaxed PlanGraph gives H_{add} and H_{max} heuristics
- Action linearization extracts relaxed plan to give H_{FF}
- Helpful actions further prune search space

FFRob - FF Planner applied to robotics

- Captures both geometric and symbolic action guidance
- Relaxation of continuous variables implies reachability can only increase
- *Reachable(C1, C2)* has messy conditional effects complicating action linearization
- Minimum Constraint Removal finds minimum cost achievers

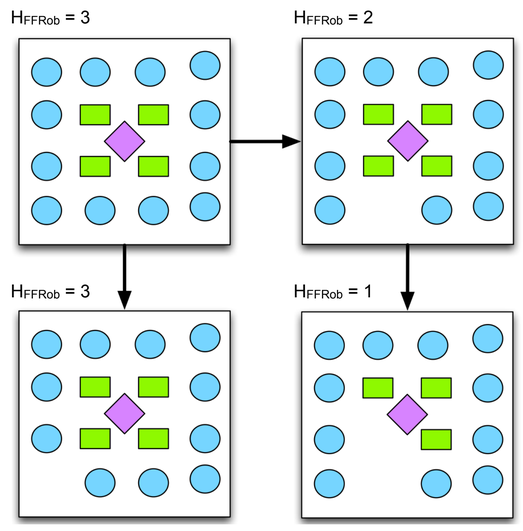

- Low amortized cost
- Geometric biases break ties by choosing states with more reachable placements

Simulations

Simulation 4 - dig block out of packed region (20 actions)

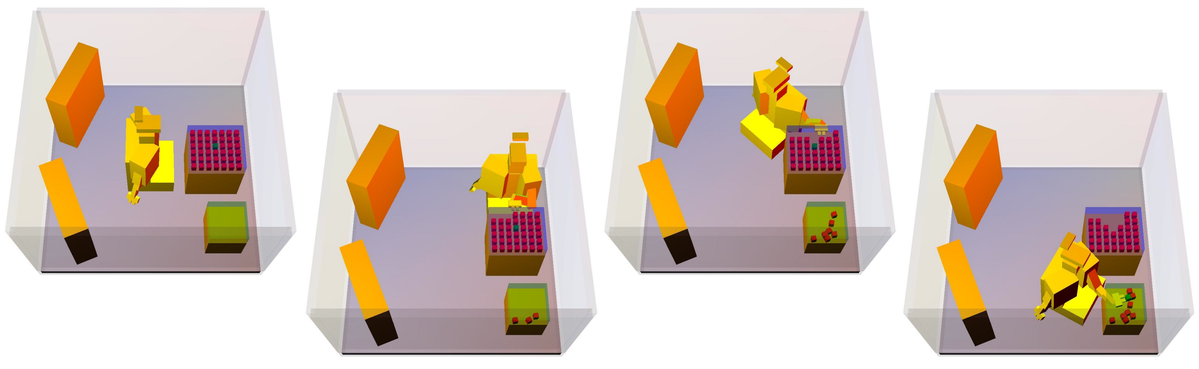

Simulation 5 - clean up messy kitchen (28 actions)

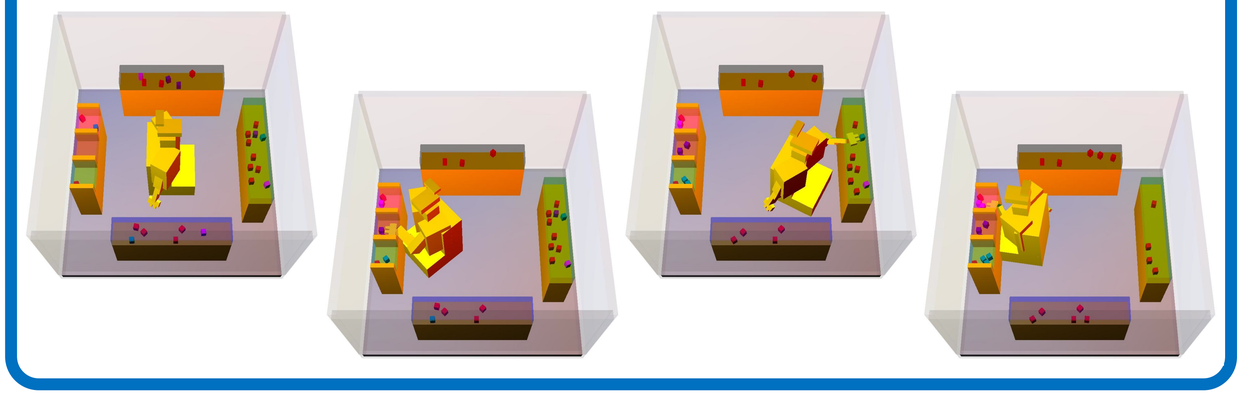

Results and Conclusion

T	Pre	No H			H_{FF}			H_{AddR}			H_{FFR}, HA			H_{FFRB}			H_{FFRB}, HA		
		t	m	s	t	m	s	t	m	s	t	m	s	t	m	s	t	m	s
0	21	265	35	48719	102	72	6123	41	19	536	6	5	78	7	5	87	2	0	23
1	25	300	0	63407	283	17	14300	162	55	2042	3	0	8	16	11	153	4	1	49
2	29	300	0	50903	300	0	8947	300	0	3052	5	1	12	17	13	114	7	2	32
3	23	300	0	39509	300	0	4849	300	0	1767	83	19	464	99	43	523	13	1	69
4	30	300	0	23920	300	0	1574	300	0	1028	300	0	1274	18	3	20	16	3	20
5	51	300	0	9422	300	0	1533	300	0	592	300	1	272	106	17	32	99	14	32

- Performance table of 6 worlds (T) with a 300 sec timeout
- Median time in sec (t), MAD time (m), median states (s)

Conclusion

- Geometric information important for strong heuristic
- CRG reuses expensive geometric work

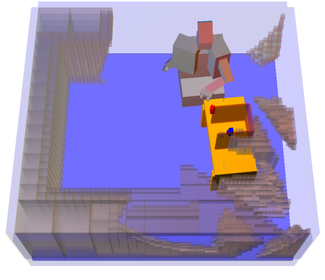

Future Work

- Generalization of approach
- Dynamic roadmap construction
- Uncertainty