

M*: A Complete Multirobot Path Planning Algorithm with Performance Bounds

Glenn Wagner, Howie Choset

Multirobot systems can provide greater sensor coverage and reliability than single robot systems. However, planning for multirobot systems is difficult due to the curse of dimensionality, which states that the size of a system’s configuration space grows exponentially with the dimensionality of the system.

Most existing algorithms can be categorized as coupled approaches, which search the full configuration space of a system, and decoupled approaches, which search some low-dimensional representation. Coupled approaches can be guaranteed to find an optimal path, but are computationally infeasible for large numbers of robots. Decoupled approaches scale better with the number of robots, but cannot generally be guaranteed to find a path. We seek to obtain the benefits of both coupled and decoupled approaches by dynamically determining which robots must be coordinated. Existing dynamic coordination approaches use constraints on communication or sequential motion, or machine learning to determine the necessary coordination between robots. Our approach, named *subdimensional expansion*, explicitly constructs a low dimensional search space, using the results of partial searches to guide the construction.

Subdimensional expansion is inspired by the intuition that robots that are well separated spatially probably do not need to be coordinated. Therefore, the majority of the time, robots can be planned for independently. To exploit this intuition, subdimensional expansion first generates an individually optimal plan for each robot, ignoring the other robots. The initial search space is formed by combining the individually optimal paths for each robot to form a one-dimensional structure embedded in the full configuration space. As the search space is explored by some path planning algorithm, robot-robot collisions are found. The search space is then *locally* increased in dimensionality along any path found by the planning algorithm leading to the collision. Doing so constructs a variable dimensional search space of minimal size which will contain the optimal path.

We then implemented subdimensional expansion for the case where the configuration space of each robot can be represented as a graph, using A* as the underlying path planning algorithm. We name the resulting algorithm M*. M* can be proven to find an optimal path in finite time, or to terminate in finite time if no path exists.

We developed two extensions to M*: rM* and inflated M*. rM* decouples planning between robots involved in simultaneous but physically separated collisions, improving performance while maintaining all the performance guarantees of basic M*. Inflated M* is a ϵ -suboptimal variant, which is guaranteed to find a path that costs no more than ϵ times




Fig. 1: A conceptual visualization of a variable dimensionality search space for five robots. Initially each robot is constrained to its individually optimal path, represented by a single line, but when robots a pair of robots collide, the dimensionality of the search space must be increased, as represented by a square labeled with the robots in collision. When three robots collide, the local dimensionality of the search space must be further increased to include all local paths of the three robots, represented by the cube.

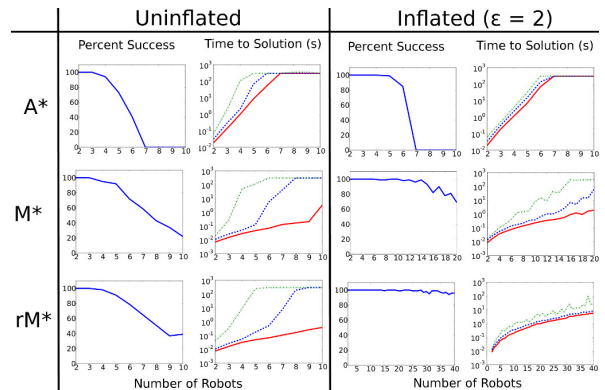


Fig. 2: We plot the percent of trials in which each algorithm was able to find a solution within 5 minutes, and the 10th, 50th, and 90th percentile of times required to find a solution for A*, M*, and recursive M* with both uninflated and inflated heuristics. The plateauing that is apparent in many of the time plots are the result of the algorithm timing out in increasingly large percentages of trials. Note that for the uninflated cases, we show results to 10 robots, while we present results for 10, 20 and 40 robots for inflated A*, M* and rM* respectively

the optimal cost. Inflated M* works by increasing the weight given to the A* heuristic by a factor of ϵ .

Simulations demonstrate that inflated rM* can reliably solve problems involving 40 robots in under 5 minutes, while fully coupled A* search is unable to solve problems involving 7 or more robots. Furthermore, inflated rM* demonstrates sub-exponential scaling of average case computation time with the number of robots, largely avoiding the curse of dimensionality.