Summary/critique of:

A general framework for sampling on the medial axis of the free space
J. Lien, S. Thomas, N. Amato, ICRA 2003

This paper presents a generalized version of the authors' previous work on the "Medial Axis PRM" (MAPRM). The basic idea behind this technique is that traditional PRMs have trouble sampling narrow spaces in constrained environments. This technique overcomes this problem by retracting randomly selected configurations to the medial axis of the configuration space. This improves sampling in narrow spaces where randomly selected configurations are likely to be contained in C-obstacles; the retraction of these points onto the medial axis increases the "basin of attraction" for narrow corridors.

The authors' previous work presented the details of this approach for two and three dimensional workspaces (that result in two to six dimensional configuration spaces). This previous work used a combination of exact and approximate methods to do the retraction without explicitly constructing the C-obstacles or the medial axis. The present work extends this approach for high-dimensional C-spaces and describes a framework consisting of three algorithms: MAPRM, MAPRM$\sim$, and MAPRM$\approx$. The different algorithms can handle different cases of dimensionality and robot convexity. The primary difference is whether they use exact or approximate methods to compute the clearance of a free configuration from the closest C-obstacle and the penetration depth of a non-free obstacle.

The basic idea of MAPRM is to take any configuration, find the closest point on the boundary of the C-obstacles, and move either away from it for free configurations or past it for non-free configurations until there is a different closest point, indicating that configuration is on the medial axis. A clearance measurement method is required to find the medial axis for a configuration in free space, and a penetration depth measurement method is needed for finding the shortest way out of the non-free configuration space.

The simplest algorithm (MAPRM) uses exact methods for 2D worlds consisting of polygonal or polyhedral obstacles. For 3D worlds, MAPRM$\sim$ uses an exact method (provided by collision detection packages) is used to measure clearance for free configurations, but an approximate method (the Lin-Canny closest features algorithm for convex robot and obstacles; otherwise the authors' brute force method) is used to determine penetration distance.

For high dimensional C-spaces, the authors propose a method that approximates the clearance or penetration depth. This method picks some number of random directions and iteratively steps forward in each direction until the boundary of the C-obstacles is reached.

The authors present experimental results on three different worlds. In two of the worlds, the traditional PRM was unable to solve the problem in 11 hours of execution time, whereas the MAPRM algorithms found solutions in 22–213 seconds. In the third environment, MAPRM found solutions in 40-99 seconds whereas the traditional PRM took 121-162 seconds.

In the simplest example, it was clear that when an exact method can be used, it results in better performance than an approximate method. The iterative methods used during sampling significantly increase the computation during this phase. One thing not explained in the paper is why MAPRM$\approx$ usually takes longer than MAPRM$\sim$ to connect nodes. Presumably this would be because the samples from MAPRM$\sim$ are more accurately on the medial axis due to the exact clearance measurement method used there.

The paper also presents analysis based on the two parameters of the approximate methods: the number of sample directions (N) chosen for calculating clearance and penetration depth. The results of their analysis are that the accuracy of the approximate method is fairly flat after N=4–100. For MAPRM$\sim$, the total solution time increases approximately linearly with N, but trends for MAPRM$\approx$ are less clear. It seems that optimal selection of these parameters depends upon the nature of the environment. For environments with many lange C-obstacles, higher values of N to calculate penetration depth are beneficial to find the closest boundary, but lower values (e.g. N=4) are sufficient for lower C-obstacle density or for smaller C-obstacles.