

Figure 6: The Kitchen scene. From left to right. The reference ray-traced image, the scene rendered using our approach searching for all the vertices, the scene rendered using our approach, searching for only 10% of the vertices and approximating the others and the scene using dynamic environment maps (i.e., recomputed at each frame). Our approach for full search as well as with the approximation is clearly much closer to the ray-traced solution than DEM.

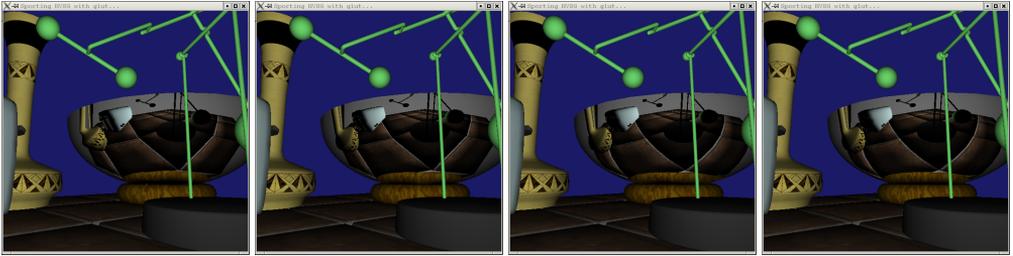


Figure 7: The Kitchen scene. Same view with different densities for the grids. The number of grid vertices is a fixed percentage of the number of mesh vertices. From left to right: 4.1%, 5.5%, 9.6%, and 14.3% of mesh vertices.

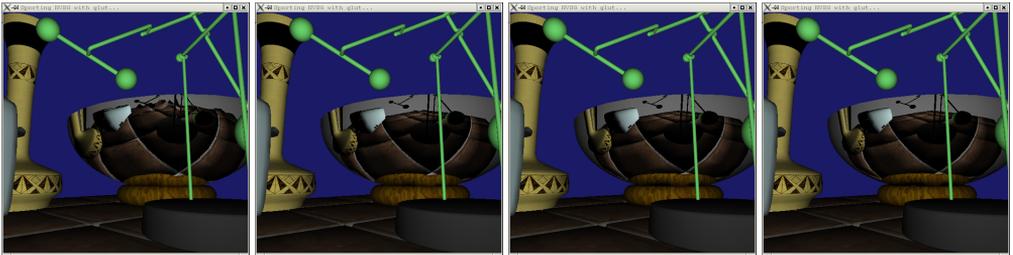


Figure 8: The Kitchen scene. Same view with different number of iterations used for the local search. No grid used, all vertices are reflected. From left to right: 3, 5, 7, and 9 iterations. Using more than 9 iterations gives no noticeable improvements.