Safe MultiRobot Localization

John Langford

Problem: Given multiple communicating robots attempting to discover their location within a shared map, what is the best way to combine the information from every robot's sensors to localize each robot? Solving this problem in closed form is intractable so approximations are necessary.

Impact: If this problem is solved well it will give large improvements in robot localization which is a fundamental problem for robots. If a robot knows where it is in the world then it will be able to make better decisions and better understand the outcomes of actions.

State of the Art: There is an early approach to multirobot localization implemented by Dieter and Hannes. The implementation lacks several optimizations and is not probabilistically sound when used repeatedly.

Approach: The approach we take explicitly avoids using the same piece of sensor information twice by keeping seperate parallel belief states. When one robot detects another the detection model is used to construct a belief state for the other robot using a local belief state based only on sensor readings taken since the last detection occurred. This approach scales linearly with the number of robots in general but some techniques can be used to reduce the the amount of computation in the best case to just propagating one belief state as for standard robot localization.

Future Work: The above technique is not optimal because it doesn't get the maximal use out of the a sequence of observations like robot A detects robot B then robot B detect robot C. The problem seems generally difficult because keeping track of the correct belief space may require propagating a number of belief states equal to the powerset of the number of robots in the worst case. This is pretty bad, but it may be possible to improve both the average and worst cases with some approximation.

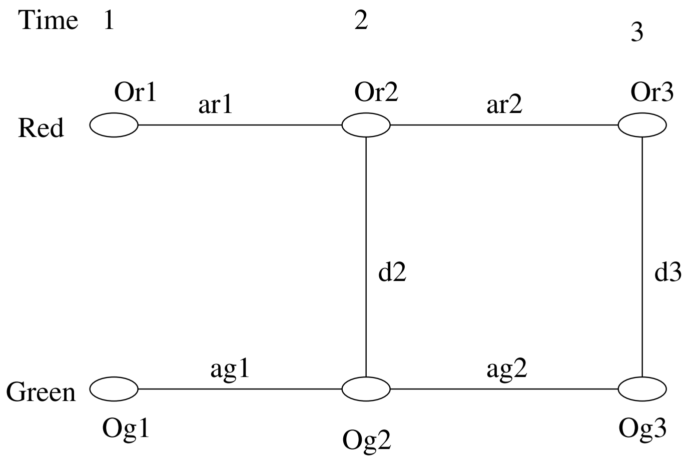

Figure 1: A diagram showing one robot detecting another twice in a row.