


A Additional Analysis of CS-PIBT

A.1 Quantitative Analysis


725 Converting the probability distribution to an action ordering
 using sampling was substantially better than using a strict
 ordering of probabilities. One possibly hypothesis is that
 the distributions induced are different as the strict ordering
 could over-prefer higher probability actions. Figure A1 vi-
 730 sualizes the action distribution when converting probabili-
 ties into action orders using strict ordering and sampling. We
 notice they are very similar, indicating this is not the cause
 of the performance difference within CS-PIBT. This specific
 distribution is for 50 agents not at their goal location. Note
 735 agents at their goal had different overall distributions (i.e.
 “stop” probabilities were 1st nearly all the time), but strict
 and sampled ordering produced similar representative distri-
 butions there too.

Action Ordering Distribution using Strict Ordering



(a) The value of cell [action, k] is the probability that action is the k^{th} preference when converting probability distributions to action orderings using strict ordering.

Action Ordering Distribution using Sampled Ordering



(b) The value of cell [action, k] is the probability that action is the k^{th} preference when converting probability distributions to action ordering using sampled ordering.

Figure A1: Action ordering distributions

A.2 Qualitative Analysis

We provide <https://shorturl.at/mtuF0> with visualization ani- 740
 mations of MAGAT with CS-Naive, CS-PIBT strict ordering
 (without randomness), and CS-PIBT with sampling. Obsta-
 cles are black and free space is white. Each agent is given
 a unique color, with their corresponding goal location visu- 745
 alized in a slightly darker color. When an agent reaches its
 goal location it turns grey. Note that since goals are visu-
 alized similarly to agents, there may be moments it looks
 like one agent is passing over/through another agent. This
 is not the case, those agents are passing over/through an-
 other agent’s goal location. The goal coloring is useful near 750
 the end of animations of failures cases as it highlights how
 agents get stuck due to other agents resting near their goal
 location.

We observe most failure cases occur due to a few agents 755
 getting stuck in deadlock or live-lock near their goal due
 to other agents resting at their nearby goal locations. Fail-
 ure cases with CS-PIBT with strict ordering show live-lock
 where agents alternate between the same positions, while
 CS-PIBT with sampling does not show it as often. Thus, 760
 sampling is effective in reducing these instances by promot-
 ing trying different actions at the same location. We also note
 that with CS-Naive, agents in crowds take quite long (or pos-
 sibly never) to leave/travel through the crowd, while with
 CS-PIBT they are able to travel effectively through them.