

Multiagent Rollout with Reshuffling for Warehouse Robots Path Planning^{*}

William Emanuelsson^{*} Alejandro Penacho Riveiros^{*}
Yuchao Li^{*} Karl H. Johansson^{*} Jonas Mårtensson^{*}

^{*} *Division of Decision and Control Systems, KTH Royal Institute of Technology, Sweden, (e-mail: wem,alejpr,yuchao,kallej,jonas1@kth.se).*

Abstract: Efficiently solving path planning problems for a large number of robots is critical to the successful operation of modern warehouses. The existing approaches adopt classical shortest path algorithms to plan in environments whose cells are associated with both space and time in order to avoid collision between robots. In this work, we achieve the same goal by means of simulation in a smaller static environment. Built upon the new framework introduced in (Bertsekas, 2021a), we propose multiagent rollout with reshuffling algorithm, and apply it to address the warehouse robots path planning problem. The proposed scheme has a solid theoretical guarantee and exhibits consistent performance in our numerical studies. Moreover, it inherits from the generic rollout methods the ability to adapt to a changing environment by online replanning, which we demonstrate through examples where some robots malfunction.

Keywords: Reinforcement learning control, multi-agent systems applied to industrial systems, industrial applications of optimal control

1. INTRODUCTION

In recent years, the trade of goods over the internet, commonly referred to as ‘e-commerce,’ has been consistently increasing its influence due to its convenience to consumers and the large variety of options available. Companies involved in this business model have experienced massive growth, but must also manage increasingly complex logistical infrastructures. Among the most critical parts of them are fulfillment centers. These are massive warehouses tasked with the storage and handling of a large number of goods at the same time. The efforts to automatize these centers must deal with the challenge of coordinating hundreds of robots to transport items inside the warehouse efficiently. The coordination problem is typically decomposed into several sub-problems [see, e.g., (Wurman et al., 2008)], and in this work, we focus on path planning problem, which can be modeled adequately as computing shortest paths for a large number of robots in a grid environment, as in, for example, (Li et al., 2021).

Research on shortest path problems has a long and rich history. Some of the most well-established algorithms include Bellman-Ford (Bellman, 1958), Dijkstra (Dijkstra, 1959), A* (Hart et al., 1968) and Auction (Bertsekas, 1991). These algorithms exploit different properties of the problem to minimize the computational cost, and collectively address the problem sufficiently well. However, they only tackle problems involving a single agent, thus making them not directly applicable to warehouse path planning. More recent efforts have built upon previous methods to tackle problems with several agents. Algorithms of this kind include cooperative A* (Silver, 2005), M* (Wagner

and Choset, 2011) and conflict based search (Sharon et al., 2015). They offer different possibilities in the trade-off between computational cost and optimality for multiagent path finding. The common feature of those approaches is to augment time to space and create an enlarged grid environment, with each cell associated with both physical space and time. By searching through this enlarged environment, agents can avoid collisions between each other.

Still, when considering the path planning problem in a warehouse, the aforementioned algorithms are not readily applicable. This is because they only deal with one-shot problems, where each agent only has to move from its initial position to a single target position. In contrast, in a real warehouse, a robot is expected to be assigned endless tasks, picking and delivering items one after the other. In other words, it is more suitable to regard the problem as one involving *infinite horizon*, i.e., planning infinitely long paths. Some recent works have designed algorithms that are built upon the previous results and are tailored for this feature, and (Ma et al., 2017) and (Li et al., 2021) are notable successes among them.

Parallel to the algorithmic progress discussed above, some high-profile successes in the field of reinforcement learning (RL), such as (Silver et al., 2018), have taken the world by storm. Among the tools offered by RL, a broad class of algorithms is couched on the theory of dynamic programming, which in turn laid the foundation for the Bellman-Ford algorithm. Therefore, it is no surprise that tools developed in RL can be brought to bear to address the problem considered here. The key distinguishing factor of RL methods from those discussed above is its use of a simulator, where simulations of future results are applied to facilitate decision making at the current stage. In particular, through the use of simulation in a smaller grid

^{*} This work was supported by the Swedish Foundation for Strategic Research, and the Swedish Research Council.

environment than those applied in the above schemes, it is possible to detect potential collisions between agents, thus leading to proper decisions at the present stage. For a textbook account of the subject, see (Sutton and Barto, 2018; Bertsekas, 2019).

Among a great variety of methods in RL, *multiagent rollout* is well-suited for the warehouse path planning problem. Like the rollout method first developed in (Tesauro and Galperin, 1996) for the game of backgammon, it relies on real-time simulation results of a known heuristic for decision making. An appealing consequence of this nature is its strong ability to adapt to a changing environment. The additional feature that is particularly tailored for multiagent problems is that it focuses on the decision selection process one agent at a time, thus avoiding the computational burden rendered by large decision spaces. The method was outlined in (Bertsekas and Tsitsiklis, 1996, Section 6.1.4), and has been substantially extended in (Bertsekas, 2021a). Further details of the method will be given in the next section.

In this work, we build upon the multiagent rollout algorithm, and propose a variant of it that involves random reshuffling of the orders of agents. The proposed method is applied to warehouse robot path planning problems, where up to 200 robots are involved. In our numerical study, the proposed scheme demonstrates a better rate of success compared with both standard rollout, as well as a scheme adapted from cooperative A*. With suitable pre-computation of shortest paths for a static environment, the average computation time can be as little as 50 ms.

The remainder of the paper is organized as follows. Section 2 introduces the mathematical problem, and describes the proposed scheme for this problem. In Section 3, we provide further details on the modeling of warehouse path finding. In Section 4 we demonstrate the results of our numerical studies, where our approach is compared with standard multiagent rollout as well as a variant of cooperative A*.

2. MULTIAGENT ROLLOUT WITH RESHUFFLING

In this section, we describe the multiagent rollout with reshuffling algorithm, which is later applied to plan paths for warehouse robots. We will first introduce the multiagent optimal control problems with which we model the path planning problem of our interest, and then provide details of the rollout method. Our presentation is somewhat abstract, leaving further details specific to warehouse path planning in the next section.

2.1 Multiagent Deterministic Optimal Control

The problem considered here involves stationary dynamics

$$x_{k+1} = f(x_k, u_k), \quad k = 0, 1, \dots, \quad (1)$$

where x_k and u_k are state and control at stage k , which belong to state and control spaces X and U , respectively, and f maps $X \times U$ to X . Both the state and control spaces contain finitely many elements. The control u_k must be chosen from a nonempty constraint set $U(x_k) \subset U$ that may depend on x_k . The cost of applying u_k at state x_k is denoted by $g(x_k, u_k)$, and is assumed to be real-valued:

$$g(x_k, u_k) \in \mathfrak{R}, \quad (2)$$

where \mathfrak{R} denotes the real line. We consider feedback policy μ that maps X to U and satisfying $\mu(x) \in U(x)$ for all x .

The *cost function* of a policy μ , denoted by J_μ , maps X to \mathfrak{R} , and is defined at any initial state $x_0 \in X$, as

$$J_\mu(x_0) = \sum_{k=0}^{\infty} \alpha^k g(x_k, \mu(x_k)), \quad (3)$$

where $\alpha \in (0, 1)$ is called *discount factor*, which reflects the emphasis on the current cost over future ones, and $x_{k+1} = f(x_k, \mu(x_k))$, $k = 0, 1, \dots$. The limit in (3) is well-posed in view of the finiteness of X and U . The optimal cost function J^* is defined pointwise as

$$J^*(x_0) = \inf_{\substack{u_k \in U(x_k), k=0,1,\dots \\ x_{k+1}=f(x_k, u_k), k=0,1,\dots}} \sum_{k=0}^{\infty} \alpha^k g(x_k, u_k). \quad (4)$$

A stationary policy μ^* is called optimal if

$$J_{\mu^*}(x) = J^*(x), \quad \forall x \in X.$$

The multiagent nature of the problem manifests itself through the Cartesian product structure of $U(x_k)$, i.e.,

$$U(x_k) = U^1(x_k) \times \dots \times U^m(x_k), \quad (5)$$

where m is an integer representing the number of agents involved in the problem. As a result, u_k can be written as

$$u_k = (u_k^1, \dots, u_k^m),$$

with $u_k^i \in U^i(x_k)$, $i = 1, \dots, m$. Accordingly, the system dynamics (1) and stage cost (2) can be written as

$$x_{k+1} = f(x_k, u_k^1, \dots, u_k^m), \quad g(x_k, u_k^1, \dots, u_k^m),$$

respectively, with $u_k = (u_k^1, \dots, u_k^m)$ when stressing the multiagent structure of the problem. In addition, policy itself μ can be decomposed similarly as

$$\mu(x) = (\mu^1(x), \dots, \mu^m(x)),$$

with $\mu^i(x) \in U^i(x)$, $i = 1, \dots, m$.

For the problem considered here, it can be shown that there exists a stationary optimal policy μ^* ; see, i.e., (Bertsekas, 2019, Props. 4.3.2 and 4.3.4). Two exact DP algorithms for addressing the problem are value iteration (VI) and policy iteration (PI). However, these algorithms may be intractable, due to the *curse of dimensionality*, which refers to the explosion of the computation as the cardinality of X increases. Thus, various approximation schemes have been proposed, where one aims to obtain some suboptimal policy $\tilde{\mu}$ such that $J_{\tilde{\mu}} \approx J^*$. These schemes often involve minimizing some objective over the set $U(x)$. Yet, due to the structure (5) of $U(x)$, even those suboptimal schemes can be rather challenging as searching through the set $U(x)$ can be too costly as well. To see this, suppose each set $U^i(x)$ contains at most C elements, then the total number of elements in $U(x)$ can be as large as C^m , which is prohibitively expensive to search through even for modest m . The multiagent rollout algorithm proposed in (Bertsekas, 2021a) addresses exactly this problem, as we will discuss next.

2.2 Multiagent Rollout

The standard rollout scheme relies on the use of a policy, which is called *base policy*. It improves upon the base policy by searching through the set $U(x)$ in real-time, and produces a new policy which is called the *rollout policy*.

It can be shown that the rollout policy is guaranteed to outperform the base policy in the sense that its cost function values are no more than that of the base policy at every state; see (Bertsekas, 2019, Prop. 4.6.1).

In particular, given a base policy μ and the current state x , the rollout policy is defined by performing minimization

$$\min_{u \in U(x)} \{g(x, u) + \alpha J_\mu(f(x, u))\} \quad (6)$$

in real-time and applies the control that attains the minimum. Since the cost function of the base policy μ is typically not in closed form, obtaining their values involves certain forms of real-time computation as well, such as simulation. Therefore, the above minimization requires computing as many values $J_\mu(f(x, u))$ as the number of elements in $U(x)$, and comparing just as many numbers in the form of (6). In the case of a multiagent system, the computation amount can be daunting in face of hard time constraints, which makes the scheme impractical.

Multiagent rollout circumvents the challenge by improving the policy *one-agent-at-a-time*. Specifically, given the base policy $\mu = (\mu^1, \dots, \mu^m)$, it performs a sequence of minimization

$$\begin{aligned} \tilde{\mu}^1(x) &\in \arg \min_{u^1 \in U^1(x)} \left\{ g(x, u^1, \mu^2(x), \dots, \mu^m(x)) \right. \\ &\quad \left. + \alpha J_\mu \left(f(x, u^1, \mu^2(x), \dots, \mu^m(x)) \right) \right\}, \\ \tilde{\mu}^2(x) &\in \arg \min_{u^2 \in U^2(x)} \left\{ g(x, \tilde{\mu}^1(x), u^2, \mu^3(x), \dots, \mu^m(x)) \right. \\ &\quad \left. + \alpha J_\mu \left(f(x, \tilde{\mu}^1(x), u^2, \mu^3(x), \dots, \mu^m(x)) \right) \right\}, \\ &\dots \quad \dots \quad \dots \\ \tilde{\mu}^m(x) &\in \arg \min_{u^m \in U^m(x)} \left\{ g(x, \tilde{\mu}^1(x), \dots, \tilde{\mu}^{m-1}(x), u^m) \right. \\ &\quad \left. + \alpha J_\mu \left(f(x, \tilde{\mu}^1(x), \dots, \tilde{\mu}^{m-1}(x), u^m) \right) \right\}, \end{aligned} \quad (7)$$

and applies the control $\tilde{\mu}(x)$ that is defined as

$$\tilde{\mu}(x) = (\tilde{\mu}^1(x), \dots, \tilde{\mu}^m(x)).$$

Let us assume that each set $U^i(x)$ contains at most C elements as before. Then the above minimization involves evaluating and comparing at most Cm values, which is a substantial improvement compared with C^m values needed in (6). In addition, if we denote the optimal value at the end of the minimization as $\tilde{J}(x)$, i.e.,

$$\begin{aligned} \tilde{J}(x) &= \min_{u^m \in U^m(x)} \left\{ g(x, \tilde{\mu}^1(x), \dots, \tilde{\mu}^{m-1}(x), u^m) \right. \\ &\quad \left. + \alpha J_\mu \left(f(x, \tilde{\mu}^1(x), \dots, \tilde{\mu}^{m-1}(x), u^m) \right) \right\}, \end{aligned}$$

or equivalently

$$\tilde{J}(x) = g(x, \tilde{\mu}(x)) + \alpha J_\mu \left(f(x, \tilde{\mu}(x)) \right), \quad (8)$$

then we can show that

$$J_{\tilde{\mu}}(x) \leq \tilde{J}(x), \quad \forall x \in X. \quad (9)$$

A formal statement as well as its proof is given in Appendix A of the extended version (Emanuelsson et al., 2022). Besides, it is shown in (Bertsekas, 2021a, Prop. 2) that the multiagent rollout policy $\tilde{\mu}$ outperforms the base policy in the same sense as in the standard rollout case, i.e.,

$$J_{\tilde{\mu}}(x) \leq J_\mu(x), \quad \forall x \in X. \quad (10)$$

In fact, rollout and multiagent rollout can be regarded as one step of Newton's method for solving the original optimal control problem and its equivalent reformulation, respectively, as shown in (Bertsekas, 2021a, Section 3.4 and 3.7). Therefore, the produced policies are likely to be much better than the base policy. However, it is not known in general whether the standard rollout policy is better than the multiagent one, despite the much more effort being required. Empirical studies given in (Bhattacharya et al., 2021) addressing a partial information problem have found that these two have quite comparative performance.

Owing to its flexibility, multiagent rollout admits various forms of modifications and extensions. Extensive discussions can be found in (Bertsekas, 2021b, Sections 3.2, 5.3 and 5.6). One extension that is essential for our application involves multiple base policies. In particular, suppose that we have access to policies μ_1 and μ_2 . This form of rollout defines a function $\bar{J} : X \mapsto \mathfrak{R}$ as $\bar{J}(x) = \min\{J_{\mu_1}(x), J_{\mu_2}(x)\}$, and replaces J_μ with \bar{J} in the sequence of minimization (7). It can be shown that the rollout policy $\tilde{\mu}$ is better than both μ_1 and μ_2 , i.e.,

$$J_{\tilde{\mu}}(x) \leq J_{\mu_1}(x), \quad J_{\tilde{\mu}}(x) \leq J_{\mu_2}(x), \quad \forall x \in X.$$

In addition, the bound (9) also holds, with J_μ replaced by \bar{J} in (8) which defines the corresponding bound \bar{J} .

2.3 Multiagent Rollout with Reshuffling

The multiagent rollout framework offers many options for further improvement. Among them, we focus on the order of agents according to which the minimization similar to (7) is carried out. Ideas on optimizing orders for similar optimization problems have appeared in the context of, for example, path planning of robots, under the name of 'prioritized planning;' see (Latombe, 1991, Section 2.3.1) and (Čáp et al., 2015). In this work, however, we will rely on orders that are generated randomly. Before getting into the details of our scheme, we first note that the performance bound (9) remains valid even if the agent-by-agent optimization is performed in a different order other than the default one.

The basis of our algorithm is an assumption that we know a 'rule' that provides a verdict on whether a control $u \in U(x)$ at a given state x is 'good.' We introduce permutation functions, denoted by σ and τ , which are one-to-one mappings with both domain and image being $\{1, \dots, m\}$. Thus, their inverse images exist and are denoted as σ^{-1} and τ^{-1} respectively.

Given the current state x and permutation σ , we define $\tau_0 = \sigma$. In addition, we define ℓ_0^i as

$$\ell_0^i = \tau_0^{-1}(i), \quad i = 1, \dots, m. \quad (11)$$

Then with a slight abuse of notation, we rearrange the elements of control according to τ_0 , i.e.,

$$u = \left(u^{\ell_0^1}, u^{\ell_0^2}, \dots, u^{\ell_0^m} \right),$$

perform a sequence of minimization

$$\begin{aligned} \tilde{u}_0^{\ell_0^1} &\in \arg \min_{u^{\ell_0^1} \in U^{\ell_0^1}(x)} \left\{ g(x, u^{\ell_0^1}, \mu^{\ell_0^2}(x), \dots, \mu^{\ell_0^m}(x)) \right. \\ &\quad \left. + \alpha J_\mu \left(f(x, u^{\ell_0^1}, \mu^{\ell_0^2}(x), \dots, \mu^{\ell_0^m}(x)) \right) \right\}, \end{aligned}$$

$$\begin{aligned}
\tilde{u}_0^{\ell_0^2} &\in \arg \min_{u^{\ell_0^2} \in U^{\ell_0^2}(x)} \left\{ g(x, \tilde{u}_0^{\ell_0^1}, u^{\ell_0^2}, \mu^{\ell_0^3}(x), \dots, \mu^{\ell_0^m}(x)) \right. \\
&\quad \left. + \alpha J_\mu \left(f(x, \tilde{u}_0^{\ell_0^1}, u^{\ell_0^2}, \mu^{\ell_0^3}(x), \dots, \mu^{\ell_0^m}(x)) \right) \right\}, \\
&\dots \quad \dots \quad \dots \\
\tilde{u}_0^{\ell_0^m} &\in \arg \min_{u^{\ell_0^m} \in U^{\ell_0^m}(x)} \left\{ g(x, \tilde{u}_0^{\ell_0^1}, \dots, \tilde{u}_0^{\ell_0^{m-1}}, u^{\ell_0^m}) \right. \\
&\quad \left. + \alpha J_\mu \left(f(x, \tilde{u}_0^{\ell_0^1}, \dots, \tilde{u}_0^{\ell_0^{m-1}}, u^{\ell_0^m}) \right) \right\},
\end{aligned} \tag{12}$$

and define a control \tilde{u}_0 as

$$\tilde{u}_0 = \left(\tilde{u}_0^{\ell_0^1}, \tilde{u}_0^{\ell_0^2}, \dots, \tilde{u}_0^{\ell_0^m} \right).$$

If the control \tilde{u}_0 is deemed to be ‘good’ according to the rule, we set $\tilde{\mu}(x) = \tilde{u}_0$. Otherwise, we generate according to certain probability distribution a new permutation τ_1 , define ℓ_1^i similar to (11), and carry out agent-by-agent optimization similar to (12). This procedure is repeated until after generating, say j , new permutations, the corresponding \tilde{u}_j is considered as ‘good.’ Then we define $\tilde{\mu}(x) = \tilde{u}_j$, and the state x' and permutation σ' at the next stage is given by

$$x' = f(x, \tilde{\mu}(x)), \quad \sigma' = \tau_j.$$

As an example of the rule for selecting controls, let us assume that there exists a function $\hat{J} : X \mapsto \mathfrak{R}$ that sets a realistic bound for the multiagent rollout policy. As stated earlier, the performance bound (9) remains valid regardless of the order of agents being optimized. Thus, after obtaining control \tilde{u}_j , we may define a value \hat{V}_j as

$$\hat{V}_j = g(x, \tilde{u}_j) + \alpha J_\mu(f(x, \tilde{u}_j)),$$

and compare \hat{V}_j and $\hat{J}(x)$. If $\hat{V}_j \leq \hat{J}(x)$, then the control \tilde{u}_j is applied to the system by setting $\tilde{\mu}(x) = \tilde{u}_j$. In view of the construction of $\tilde{\mu}$, we have that

$$J_{\tilde{\mu}}(x) \leq \hat{J}(x), \quad \forall x \in X.$$

3. MODELLING FOR THE WAREHOUSE PATH PLANNING PROBLEM

In this section, we present details of the model for the warehouse path finding problem. We will first describe the components of the state x and the control u at each stage, and then introduce the system dynamics $f(\cdot)$ and stage cost $g(\cdot)$ used in our numerical studies presented in the next section. There is a considerable amount of flexibility in the modeling process, and judged by the numerical results, our choices here strike a good balance between implementation expediency and satisfactory performance.

Given that the warehouse environment is modeled as a grid world, we assign to each cell a unique scalar index. We assume that each one of the agents, the goods, as well as the delivery points occupy a single cell. Thus, for i th agent, its position at k th stage can be described by the index of the cell at which it is located, and is denoted as p_k^i . Its current target position is the index of the cell of the good or the delivery point it aims to reach, and is denoted as t_k^i . Therefore, information related to the i th agent is adequately captured by the pair of indices (p_k^i, t_k^i) . All the other information that is varying over stages and is

useful for the purpose of defining the dynamics is lumped in a vector denoted as b_k . Among others, it includes the positions of goods, which act as obstacles when present. Collecting indices pairs (p_k^i, t_k^i) and b_k together, we define the state at k th stage as

$$x_k = \left((p_k^1, t_k^1), (p_k^2, t_k^2), \dots, (p_k^m, t_k^m), b_k \right).$$

As for control of the i th agent, its control constraint set $U^i(x_k)$ at state x_k contains at most $C = 5$ elements. They represent the actions of moving to the four neighboring cells located to the left, right, up, and down of its current position p_k^i , as well as staying still. Depending on the values of (p_k^i, t_k^i) and b_k , some of those controls may be removed due to the boundary of the grid world, as well as other goods acting as obstacles. Therefore, the dependence of set $U^i(x_k)$ on x_k is only through (p_k^i, t_k^i) and b_k , but not the positions of other agents. This means that some control u_k that leads to a collision at $(k+1)$ th stage (i.e., several agents occupying the same grid at the same time) is considered feasible, namely $u(k) \in U(x_k)$. On the other hand, they are incentivized to avoid such situations by a large cost. Our choice for modeling collision avoidance through cost, rather than control constraints, is mainly motivated by the streamlined mathematical analysis in the preceding section. In practice, there is little difference in the implementation.

The system dynamics $f(\cdot)$ is composed of three parts: mechanisms for updating p_k^i, t_k^i , and b_k respectively. Given its current position p_k^i , i th agent applying control u_k^i would lead itself to the position p_{k+1}^i that is either a neighbor of p_k^i or equaling to p_k^i in the case of standstill. Note that we allow the possibility that $p_k^i = p_k^j$ for $i \neq j$, consistent with the feasible control discussed before. As for the target position, we define $t_{k+1}^i = t_k^i$ if $p_k^i \neq t_k^i$, meaning that a target is not updated until the target is reached by its corresponding agent. On the other hand, if $p_k^i = t_k^i$, a new target t_{k+1}^i is assigned based upon x_k and u_k , the mechanism of which is assumed to be deterministic. In addition, the vector b_k is updated accordingly depending on the remaining goods.

The stage cost $g(\cdot)$ consists mainly of two parts: a large penalty for collisions, and a negative cost (equivalently, a ‘reward’) for reaching the goal. In particular, the stage cost of applying u_k at x_k is defined as

$$g(x_k, u_k) = g_1(x_k, u_k) + \sum_{i=1}^m g_2(p_k^i, t_k^i),$$

where $g_1(\cdot)$ encodes the collision cost, which is defined as

$$g_1(x_k, u_k) = c_1 \cdot n(f(x_k, u_k)),$$

with c_1 being a large positive constant, and n mapping a state to the number of collisions that occurred in the state. The function $g_2(p_k^i, t_k^i)$ takes a negative value c_2 if $p_k^i = t_k^i$ and 0 otherwise.

At this point, together with a suitably defined discount factor α , we obtain a multiagent deterministic optimal control problem with details filled in. The problem is then addressed by the multiagent rollout with reshuffling algorithm introduced in Section 2. We apply two base policies μ_1 and μ_2 , with μ_1 being the shortest paths for each agent computed by A^* , while the μ_2 being the same

as μ_1 except that it freezes robots for different numbers of steps depending on their indices, i.e., staying still for a few consecutive stages. The rule used to define a ‘good’ control is that a control $u_k \in U(x_k)$ at x_k is considered ‘good’ if $g_1(x_k, u_k) = 0$, or equivalently, no immediate collision is caused by applying u_k at x_k . The testing results with our algorithm are presented next.

4. NUMERICAL STUDIES

This section presents the results of our numerical studies. The environment used here is composed of 47×115 grids, whose layout resembles that of (Wurman et al., 2008, Fig. 3). There are m agents which are tasked to collaboratively deliver 1183 goods to their designated delivery points. At any given time, if the number of remaining goods is larger than m , each robot is assigned randomly a good, whose delivery point is also chosen randomly.¹ If the scheduling mechanism is available, this random procedure can be duly replaced. A new good among the remaining ones is assigned to a robot once its current task is completed. Note that our approach applies to the life-long delivery mission, by which we mean that there are infinite amounts of goods arriving sequentially for delivery. For convenience, we restrict our attention to a finite goods case. A test where all the goods have been delivered successfully or where a collision has occurred is named an ‘episode.’ The problem data of the optimal control problem described in Section 3 are $\alpha = 0.999$, $c_1 = 10^{20}$, and $c_2 = -10^4$.

We use cooperative A* introduced in (Silver, 2005) as a benchmark for comparison (‘CA*’ for short in the figures). It computes the shortest paths of agents according to a fixed order via A* in a space-time grid, while treating the paths of preceding agents as obstacles. As the scheme is a one-short algorithm (see the discussion in Section 1), it is executed every time a new assignment is given (or equivalently, a delivery is completed, provided that there are remaining goods to be assigned). Since simulation in our scheme deals with a static environment, we can compute all the shortest paths to all goods and delivery points offline and store them for online simulation. This procedure is called pre-computation. Our scheme with and without pre-computation is labeled as ‘MA-rollout’ and ‘MA-rollout w/o pre-comp.’ in the figures respectively. For those approaches, we test cases with agent number m ranging from 100 to 200 in the same environment, and the listed results are averaged over successful tests among 200 episodes for each agent count. An Apple M1 processor was used for all tests, which were performed on a single thread. The detailed implementation of our scheme as well as demonstrations of a few test cases can be found in <https://github.com/will-em/multi-agent-rollout>.

4.1 Success rate, computation time, and memory usage

The rates of successful episodes with our scheme and cooperative A* are shown in Fig. 1. Our approach is able to complete all delivery missions without any collision occurring in all tests with on average less than 2 reshuffling

¹ Throughout our implementation, all random events are generated according to a uniform distribution.

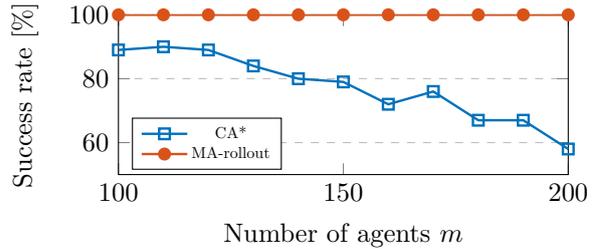


Fig. 1. The success rate (in percent) at which the methods are able to complete the task.

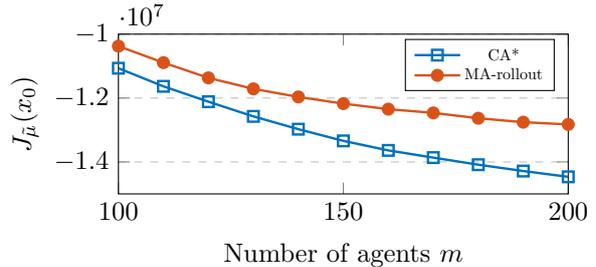


Fig. 2. The cumulative cost $J_{\bar{\mu}}$, averaged over 200 episodes. This is far smaller than \bar{J} , which has an average value of around 10^{20} , verifying inequality (9).

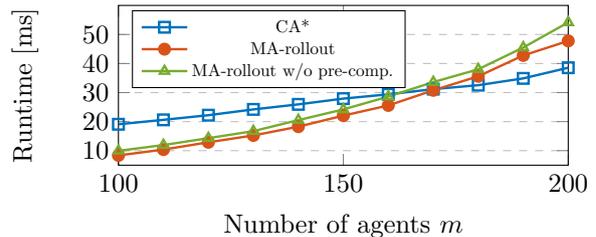


Fig. 3. The average computation time required to obtain feasible controls for all agents.

procedures, while the success rate drops to about 60% when $m = 200$ in the case of cooperative A*. Note that for $m = 100$, our result of cooperative A* is consistent with the result reported in the original work (Silver, 2005, Fig. 3). Among those where cooperative A* completes the task, its average cost is lower than that of our approach, as shown in Fig. 2. Still, both performance bounds (9) and (10) for our scheme are verified in those tests.

The computation times for different approaches are averaged over their respective successful episodes and shown in Fig. 3. Although cooperative A* has a slight edge over our method when $m > 160$, its successful episodes are less than 80% of the total tests. In addition, similar to the generic rollout method, our scheme is well-suited for parallel computation. Therefore, the computation time for our scheme would be around $1/C$ of what is shown in Fig. 3 had parallel computation been fully brought to bear, see (Bertsekas, 2021a) for further details. The pre-computation of shortest paths reduces about 20% computation time in our tests.

Besides, the scale of the graphs where the shortest-path searches are performed differs between the methods. The graph needed by cooperative A* is expected to be N

times larger than that of our scheme, where N denotes the largest number of stages needed for all agents to reach their current respective targets. This may therefore pose limitations on the maximum path length of cooperative A* from a memory standpoint. For the pre-computation used in our scheme, the total size of the data is the dimension of the environment times the sum of the number of the spots where the goods are placed and that of the delivery points, which in our example is around 7 Mb.

4.2 Adaptivity

Since the real-world operation of a warehouse is prone to unforeseeable events, it is of great importance for a warehouse path-finding algorithm to be able to adapt to these circumstances. The proposed algorithm can adapt to a changing environment by means of online replanning, and we demonstrate this ability by addressing randomly occurring robot malfunction, which is modeled by using a failure mode for the robots. In particular, up to 20% of robots are fixed at their position indefinitely starting from a random stage and are acting as obstacles ever since. Their assigned tasks are later assigned to the functioning robots. The indices of malfunctioning robots are assumed known. To cope with this situation, our proposed scheme simply regards the malfunctioned robots as the ones with modified control constraint sets $\hat{U}^i(x_k)$, containing nothing but one control that corresponds to ‘staying still.’ Then our scheme can be applied without any further modification. We carried out 50 tests with 100 agents. Our scheme adapted to this kind of scenario, and the remaining agents continued their operation by circumnavigating these malfunctioned robots.

5. CONCLUSION

We proposed the multiagent rollout with reshuffling for solving the warehouse robot path planning problem. Based upon the framework introduced in (Bertsekas, 2021b), our scheme can produce paths for a large number of robots in real time while achieving collision avoidance. Unlike existing schemes that are searching in environments associated with both space and time, our method relies on the use of simulation in a smaller static environment. It also inherits the generic features of rollout and can adapt to a changing environment by means of online replanning. Our numerical studies demonstrated that the scheme scaled well with a large number of robots and could adapt to situations where robot malfunctions occurred.

ACKNOWLEDGEMENTS

The authors are grateful to Prof. Dimitri P. Bertsekas for valuable suggestions and comments in various stages of the project. They also appreciate the contributions of Laura Briffa in implementing early versions of the scheme.

REFERENCES

Bellman, R. (1958). On a routing problem. *Quarterly of Applied Mathematics*, 16(1), 87–90.
 Bertsekas, D. (2019). *Reinforcement learning and optimal control*. Athena Scientific.

Bertsekas, D. (2021a). Multiagent reinforcement learning: Rollout and policy iteration. *IEEE/CAA Journal of Automatica Sinica*, 8(2), 249–272.
 Bertsekas, D. (2021b). *Rollout, policy iteration, and distributed reinforcement learning*. Athena Scientific.
 Bertsekas, D. and Tsitsiklis, J.N. (1996). *Neuro-dynamic programming*. Athena Scientific.
 Bertsekas, D.P. (1991). An auction algorithm for shortest paths. *SIAM Journal on Optimization*, 1(4), 425–447.
 Bhattacharya, S., Kailas, S., Badyal, S., Gil, S., and Bertsekas, D. (2021). Multiagent rollout and policy iteration for pomdp with application to multi-robot repair problems. In *Conference on Robot Learning*, 1814–1828. PMLR.
 Čáp, M., Novák, P., Kleiner, A., and Selecký, M. (2015). Prioritized planning algorithms for trajectory coordination of multiple mobile robots. *IEEE transactions on automation science and engineering*, 12(3), 835–849.
 Dijkstra, E.W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1), 269–271.
 Emanuelsson, W., Riveiros, A.P., Li, Y., Johansson, K.H., and Mårtensson, J. (2022). Multiagent rollout with reshuffling for warehouse robots path planning. *arXiv preprint arXiv:2211.08201*.
 Hart, P., Nilsson, N., and Raphael, B. (1968). A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2), 100–107.
 Latombe, J.C. (1991). *Robot motion planning*, volume 124. Springer Science & Business Media.
 Li, J., Tinka, A., Kiesel, S., Durham, J.W., Kumar, T.K.S., and Koenig, S. (2021). Lifelong Multi-Agent Path Finding in Large-Scale Warehouses.
 Ma, H., Li, J., Kumar, T.S., and Koenig, S. (2017). Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’17*, 837–845. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC.
 Sharon, G., Stern, R., Felner, A., and Sturtevant, N.R. (2015). Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219, 40–66. doi: 10.1016/j.artint.2014.11.006.
 Silver, D. (2005). Cooperative Pathfinding. In *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment*, 117–122.
 Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419), 1140–1144.
 Sutton, R.S. and Barto, A.G. (2018). *Reinforcement learning: An introduction*. MIT press.
 Tesauro, G. and Galperin, G. (1996). On-line policy improvement using monte-carlo search. In *NeurIPS*.
 Wagner, G. and Choset, H. (2011). M*: A complete multirobot path planning algorithm with performance bounds. In *2011 IEEE/RSJ IROS*, 3260–3267. doi: 10.1109/IROS.2011.6095022.
 Wurman, P.R., D’Andrea, R., and Mountz, M. (2008). Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine*, 29(1), 9–9.