



MACHINE LEARNING RESEARCH GROUP  
ROBOTICS RESEARCH GROUP  
DEPARTMENT OF ENGINEERING SCIENCE  
UNIVERSITY OF OXFORD

---

Technical Report

# Lazy auctions for multi-robot collision avoidance and motion control under uncertainty

Authorship:

Jan-Peter Calliess\*, Daniel Lyons<sup>†</sup> and Uwe Hanebeck<sup>†</sup>

Oxford, July 2011

Reference No.:

**PARG-11-01**

\*Dept. of Engineering Science, University of Oxford, UK.

<sup>†</sup>Inst. f. Anthropomatik, Karlsruhe Institute of Technology, Germany.

TECHNICAL REPORT

Title: Lazy auctions for multi-robot collision avoidance and motion control under uncertainty

Author(s): Jan-Peter Calliess<sup>‡</sup>, Daniel Lyons<sup>†</sup> and Uwe Hanebeck<sup>§</sup>

This report is an extended version of a workshop paper to be published in Springer's Lecture Notes series. The original publication of the workshop version will be available at [www.springerlink.com](http://www.springerlink.com).

University of Oxford  
Department of Engineering Science  
Parks Road  
Oxford, OX13PJ  
United Kingdom

We propose a distributed, iterative bidding protocol that can be employed in multi-robot coordination where the individual robots' problems can be modeled as mixed-integer programs and interaction via binary constraints. While our approach is myopic during bid calculations to render the coordination problem computationally tractable, we provide first experimental evidence that our auction-based mechanism achieves coordinated solutions that outperform fixed priority methods and match the socially optimal solution on a large proportion of randomized graph planning problems. Furthermore, we discuss theoretical conditions under which guaranteed coordination in a finite number of iterations can be proved.

As a key contribution, we describe how our method can be utilized for multi-vehicle control by combining it with an established particle-based model-predictive single-UAV control method. In simulations, we illustrate the effectiveness and efficiency of this merger and derive a distribution-independent bound on the collision probability in the presence of uncertainty.

# 1. Introduction

Owing to its practical importance, multi-agent coordination has been subject to ever increasing research efforts over the past decades. One of its subfields, multi-robot coordination, focusses on problems that reflect the specific nature of robotic agents and their environment. In contrast to strategic settings, in multi-robot coordination problems, the mechanism designer can typically afford to assume obedient agents and hence does not need to burden herself with ensuring design goals such as incentive compatibility or strategyproofness. This freedom should be much welcomed considering that robots typically interact in a complex and uncertain physical world and often can choose from a continuum of control signals (actions).

Many important planning and control problems can be stated in terms of a solution to a mixed-integer linear program (MILP). An example can be found among particle methods which have become increasingly popular for stochastic model-predictive single-vehicle control and path planning under uncertainty [5]. The drawn particles can serve to approximately bound the probability of a collision with an obstacle via chance constraints that are added as binary constraints to the MILP formulation of the vehicle’s cost-optimizing control problem [5]. The resulting plans (sequences of control inputs) are shown to result in low-cost trajectories that avoid all obstacles with adjustably high certainty.

A simple method to extend single-robot problems to multi-robot problems is to combine their individual optimization problems into one large, centralized MILP (e.g. [26]). While delivering cost-optimal results, such approaches have the architectural disadvantages of centralized approaches, and scale poorly in the number of robots and interaction-constraints. Therefore, their application is typically restricted to coordination tasks of low complexity. As finding a socially optimal solution is known to be NP-complete, most practically applicable coordination methods constitute a compromise between tractability and optimality.

As one of the best established of such approximation methods we consider *fixed priority methods* [13]. Here, the robots plan in order of their assigned priority with the highest priority robot beginning. When it is Robot  $r$ ’s turn it is informed of the plans of all higher-priority robots whose trajectories become obstacles in the space-time planning domain that  $r$  needs to avoid. While these methods are computationally attractive and several extensions have been suggested [1] [22], the rigidity of the fixed priority scheme can lead to joint solutions whose summed (*social*) cost can be unattractively high.

By contrast, the main idea of our approach is to base the decision which robot may pass a conflicting point in the space-time domain (which multiple robots initially plan on passing) not based on a fixed priority alone but chiefly on a bid each robot computes based on local

information.

To achieve socially optimal coordination all robots could in principle be non-myopic and compute VCG-bids for combinatorial bundles of resources (space-time points) [25]. Unfortunately, this approach is once again intractable and generally belongs to a class of combinatorial allocation problems which are known to be NP-hard (cf. [10]).

Addressing the inevitable tradeoff between optimality and tractability, we propose a myopic, iterative bidding protocol where each robot bids for one conflicting resource at a time without taking other potentially ensuing conflicts at subsequent resources into account (hence the term myopic). Other applications include but are not limited to distributed reinforcement learning [3], constrained decentralized allocation of atomic resources or graph routing.

Furthermore, our coordination mechanism is distributed and lazy in the sense that, instead of asking for bids on all conceivable combinations of plans, all robots plan independently and bidding only takes place for resources that turn out to be overbooked (i.e. which two or more robots plan to use simultaneously). Thereby, their coordinated paths are guaranteed to be collision-free, while at the same time the exponential blow-up resulting from considering all combinations is avoided.

Our model assumption is that the individual robots' problems are MILPs with all interaction modeled via (hard) binary constraints. This is in contrast to another large body of works on coordination that focusses on agent interaction via objective functions (e.g. [14, 15, 6]). Since our model is based on MILPs we can employ our method in the context of particle-based multi-robot open-loop control [5].

The result of this application is a distributed coordination mechanism that (with adjustably high certainty) generates collision-free paths without prior space-discretization and which can take uncertainty into account (the latter may be desirable due to sensor noise and model-inaccuracies).

The remainder of this paper is structured as follows. After placing our work in the context of the literature, we discuss the model assumptions in greater detail and describe our bidding protocol in generality. We use the notion of a *virtual obstacle* as an intuition for constraints that are successively generated as a result of coordination iterations, designed to prohibit resource conflicts (i.e. violations of binary inter-robot constraints). We then elucidate several of our method's properties in the context of graph path planning as a didactic example application. For a mild restriction of the algorithm and problem domain it is possible to prove termination in a finite number of coordination iterations. However, the rather lengthy and technical discussion of this theoretical guarantee had to be deferred to an extended version of this work [7].

Before concluding, we propose how to link our approach to stochastic control and present experiments illustrating how it can be utilized for efficient, distributed multi-vehicle control under uncertainty. For such settings, we derive probabilistic collision bounds that can provide a guideline for choosing the size of the virtual obstacles, which is an important design parameter that can be expected to influence the trade-off between conservatism

and social cost.

## 2. Related Work

Multi-robot coordination is a broad topic with numerous strands of works. The approach we present to collision avoidance and control is germane to a number of these strands comprising both approaches designed to operate in both continuous and in discrete worlds. It is beyond the scope of this paper to present an exhaustive survey of the extensive body of previous work that ranges across various disciplines. For surveys focussing on market-based approaches refer to [11, 18].

As a rather coarse taxonomy, present methods can be divided into centralized and decentralized approaches. *Centralized approaches* (e.g. [26][24]) typically rely on combining the individual agents' plans into one large, joint plan and optimizing it in a central planner. Typically, they are guaranteed to find an optimal solution to the coordination problem (with respect to an optimality criterion, such as the sum of all costs). However, since optimal coordination is NP-hard it is not surprising that these methods scale poorly in the number of participating agents and the complexity of the planning environment. With worst-case computational effort growing exponentially with the number of robots, these methods do provide the best overall solutions, but are generally intractable except for small teams.

In contrast, decentralized methods distribute the computational load on multiple agents and, combined with approximation methods, can factor the optimal problem into more tractable chunks.

There are two classes of *decentralized coordination mechanisms*. The first class imposes local interaction rules designed to induce a global behavior that emerges with little or no communication overhead. For instance, based on a specific robot motion model, Pallottino et. al. [23] propose interaction policies that result in guaranteed collision avoidance and can accommodate new robots entering the system on-line. Furthermore, under the assumption that robots reaching their goals vanish from the system, the authors prove that eventually all robots will reach their respective destination locations. While in its present version uncertainty is not explicitly taken into account, it may be worthwhile endowing their method with an explicit error model and performing a similar analysis as we provide in Sec. 7.

The second class focusses on the development of mechanisms where coordination is achieved through information exchange succeeding the distributed computations.

*Distributed optimization* techniques have been successfully employed to substitute the solution of a centralized optimization problem by solving a sequence of smaller, decoupled

---

problems (e.g. [6], [20], [21], [3] and [17]).

For example, Bererton et. al. [3] employ Dantzig-Wolfe Decomposition [8] to decentralize a relaxed version of a Bellman BLP to compute an optimal policy. However, due to the relaxation of the collision constraints, collisions are only avoided in expectation. Many of these algorithms have a market interpretation due to passing Lagrangian multipliers among the subproblems.

Generally, *market-based approaches* have been heavily investigated for multi-robot coordination over the past years [27] [16] [11]. Among these, *auction mechanisms* allow to employ techniques drawn from Economics. They are attractive since the communication overhead they require is low bandwidth due to the fact that the messages often only consist of bids. However, as optimal bidding and winner determination for a large number of resources (as typically encountered in multi-robot problems) is typically NP-hard, all tractable auction coordination methods constitute approximations and few existing works provide any proof of the social performance of the resulting overall planning solution beyond experimental validation. An exception are *SSI auctions* [18, 19]. For instance, Lagoudakis et. al. [19] propose an auction-based coordination method for multi-robot routing. They discuss a variety of bidding rules for which they establish performance bounds with respect to an array of team objectives, including social cost. While multi-robot routing is quite different from the motion control problem, we consider some of their bid design to be related in spirit to ours. It may be worthwhile considering under which circumstances one could transfer their theoretical guarantees to our setting. One of the main obstacles here may be the fact that in SSI auctions, a single multi-round auction for all existing resources (or bundles) is held. This may be difficult to achieve, especially if we, as in Sec. 7, desire to avoid prior space discretization and take uncertainty into account.

Most frequently used in approximate Bayesian inference but recently applied to coordination are message passing methods such as *max-sum* [14]. In these algorithms, agent interaction is modeled to take place exclusively via the agents' cost functions and coordination is achieved by message passing in a factor graph that represents the mutual dependencies of the coordination problem. While dualization of our inter-robot resource constraints into the objective function could be leveraged to translate our setting into theirs, several problems remain. First, the resulting factor graph would be exceptionally loopy and hence, no performance or convergence guarantees of max-sum can be given. Second, the interconnecting edges would have high weights (cf. [15]) whose removal would correspond to a relaxation of the collision-avoidance constraints and hence, render pruning-based max-sum-based methods [15] inapplicable.

Among all multi-robot path planning approaches, *fixed priority methods* are perhaps the most established ones. In its most basic form introduced by Erdmann and Lozano-Perez [13], robots are prioritized according to a fixed ranking. Planning is done sequentially according to the fixed priority scheme where higher ranking robots plan before lower ranking robots. Once a higher ranking robot is done planning, his trajectories become *dynamic obstacles*<sup>1</sup> for all lower ranking robots, which the latter are required to avoid. If

---

<sup>1</sup>The notion *dynamic obstacle* loosely corresponds to our *virtual obstacles* (cf. Sec. 7). The difference is

independent planning under these conditions is always successful, coordination is achieved in  $A$  planning iterations that spawn the necessity to broadcast  $A - 1$  messages in total (plans of higher priority agents to lower priority ones) where  $A$  is the number of robots.

By contrast, in our mechanism,  $A$  such messages need to be sent *per coordination iteration*. Although our results indicate that the number of these iterations scale mildly in the number of robots and obstacles in typical obstacle avoidance settings, such an additional computation and communication overhead needs to be justified with better coordination performance. Our experiments in subsequent sections indeed illustrate the superior performance of our flexible bidding approach over fixed priorities.

Note, our mechanism also incorporates an (in-auction) prioritization (as expressed by the robots' indices) that becomes important for winner determination whenever there is a bidding tie.

In priority methods, the overall coordination performance depends on the choice of the ranking and a number of works have proposed methods for a priori ranking selection (e.g. [2]). Conceivably, it is possible to improve our method further by optimizing its in-auction prioritization (robot indexing) with such methods. Exploring how to connect our mechanism to extensions of priority methods, such as [22], could have the potential to improve the communication overhead. Investigating the feasibility of such extension will have to be done in the course of future research efforts.

### 3. Problem Formulation

While our approach could be applied to more general scenarios, in this paper, we restrict our focus to the following multi-robot path planning problem (**MRPPP**): *A team of robots  $\mathcal{A} = \{1, \dots, A\}$  desires to find individual plans  $p^1, \dots, p^A$ , respectively, that translate to collision-free paths in free-space such that each robot  $r$ 's path leads from its start  $S(r)$  to its destination  $D(r)$ .*

Since our approach is motivated by multi-robot path planning, we interpret a plan as being in a one-to-one relationship with a path in free space. For instance, a plan could be a sequence of control inputs that linearly relates to a trajectory of locations (resources) in an environment. For simplicity of exposition, we will from now on assume that plan  $p^r$  is a time-indexed sequence  $(p_t^r)_{t \in \mathbb{N}}$  where  $p_t^r$  corresponds to a decision specifying which resource to consume at time  $t$ . However, we will lift this assumption again in Sec. 7 where the plans are indeed control inputs that linearly relate to locations.

---

that our virtual obstacles are only present at a particular time step whereas the dynamic obstacles span the whole range of all time steps. Furthermore, we described how to adjust the box-sizes to control the collision probability in the presence of uncertainty.



---

Obviously, the robots need to make sure that plans are *legal*, that is they adhere to the laws of the environment. We call the set of all legal plans the *global feasible set*  $G$ .

For example, consider a routing scenario in a graph with edges  $E$  and vertex set  $V$ . A plan could be to find a path through the network represented as a sequence of vertices that respects the graph's topology. To enforce this, we could specify a global feasible set as a subset of  $\{(p_t)_t | \forall t : (p_t, p_{t+1}) \in E\}$ . The global feasible set is global in the sense that the constraints it enforces apply to all robots in the system.

By contrast, each robot  $r$  may desire to enforce individual constraints upon the plans it generates. We can represent them as a *local feasible set*  $L^r$ . For instance, in the routing example, robot  $r$  may wish to ensure that he finds a path that leads from its start location to its destination:  $p^r \in L^r \subset \{p^r = (p_t^r)_t | p_0^r = S(r), \exists k \forall t \geq k : p_t^r = D(r)\}$ .

Depending on the environment, there might be many (possibly infinitely many) plans that are both legal and locally feasible. In most applications however, robots may have a preference over different plans implied by a local cost function  $c^r : G \rightarrow \mathbb{R}$  that assigns a cost to different plans. (For instance,  $c^r(p^r)$  may quantify the path length.) So, if robot  $r$  could plan independently, he would like to execute the solution to optimization problem:

$$\min_{p^r \in G \cap L^r} c^r(p^r).$$

Unfortunately, this is not possible in environments with multiple robots as they need to avoid collisions (i.e. plans where two robots simultaneously use the same non-divisible resource). Let  $p^{-r} = (p^r)_{r \in \mathcal{A} - \{r\}}$  denote the collection of plans of all robots except  $r$ . If  $r$  knew fixed  $p^{-r}$ , he could react to it by solving

$$\min_{p^r \in G \cap L^r \cap R(p^{-r})} c^r(p^r) \quad (3.1)$$

where  $R(p^{-r})$  is the set of all paths that are not in conflict with the paths generated by  $p^{-r}$ . If  $p^{-r}$  is a collection of tentative plans, we can interpret  $R(p^{-r})$  as the set of all plans that do not use any resource that are already used by any robot in  $\mathcal{A} - \{r\}$  based on the current belief that all other resources will be available. Notice, that  $R(p^{-r})$  would typically be specified by a set of binary ( or integer) constraints. Therefore, the individual optimization problem would be a (*binary*) *mixed-integer linear program* (MILP) which could be solved by the robot employing either standard mixed-integer-solvers or a problem domain specific algorithm of the robot's own choice. Unfortunately, due to the mutual interdependence of the constraint sets, for all  $r$ ,  $R(p^{-r})$  is unknown a priori and hence, the individual optimization problems are unknown (since the feasible sets are interdependent). This is where the necessity for coordination arises.

We can now restate the overall task description (comprising (MRPPP) as a special case) in general terms:

**TASK:** Assume each robot  $r$  ( $r = 1, \dots, A$ ) can choose a plan  $p^r \in G \cap L^r$ . Coordinate the planning process such that the overall outcome  $(p^1, \dots, p^A)$  of plans is conflict free (i.e.  $\forall t \forall a, r \in \mathcal{A}, a \neq r : p_t^r \neq p_t^a$ ) and such that the social cost  $\sum_{a \in \mathcal{A}} c^a(p^a)$  is small.

The socially optimal solution can be stated quite easily as the solution of the centralized optimization problem

$$\min_{(p^1; \dots; p^A) \in G^A \cap \times_r L^r \cap I} \sum_{r=1}^A c^r(p^r) \quad (3.2)$$

where  $I$  is a set defined by inter-robot constraints that prohibit collisions (conflicts). In other words,  $I$  is the set of all overall plans  $p = (p^1; \dots; p^A)$  such that all plans  $p^a, p^r$  use distinct resources (for  $a, r \in \mathcal{A}, r \neq a$ ). Typically we will have to specify  $I$  via binary constraints, rendering the overall optimization problem a binary mixed-integer linear program (MILP) that could in principle be solved by a centralized planning agent.

Unfortunately, such centralized approaches are known to scale poorly in the number of robots, even in expectation. They are NP-hard in the worst case and are limited by the typical architectural down-sides of multi-robot systems that rely on centralized planners. For example, central planners constitute computational and communication choke-points and a single points of failure ( cf. e.g. [9]).

Since the centralized optimization problem acc. to (3.2) scales poorly, we will seek to replace it by iteratively solving a sequence of individual, tractable problems similar to (3.1). Due to the hardness of the original problem we will have to be satisfied if the ensuing overall solution is not always socially optimal.

## 4. Mechanism

We propose an iterative mechanism that proceeds as follows:

*In each iteration, agents plan independently based on their current beliefs of available resources. Initially each agent assumes all resources are available. The planning process in each agent  $r$  is done solving an opt. problem of the form (3.1).*

*Whenever a conflict is detected, the conflicting agents participate in an auction for the contested resource. The winner is allowed to proceed as if no conflict had occurred while the losers add new constraints preventing them from using the lost resource at the specific time  $t$  where the conflict occurred in future iterations (i.e. they update their beliefs about the available resources as encoded by  $R$ ). Conflicts are resolved in time step order. That is, a conflict that would lead to a collision at time  $t$  is resolved before a detected conflict that would lead to a collision at time step  $t' > t$ . If we define the auction horizon to be the largest time step  $t$  where a conflict has been resolved then this horizon increases monotonically from coordination iteration to iteration until no more conflicts arise.*

*Whenever an agent has won a resource for a certain time step  $t$  in past iterations that*

she does not need anymore in her current plan, she releases it for  $t$  and informs the other agents of this event. Once all conflicts are resolved, the agents can execute their final plans.

Winner determination of an auction proceeds as follows: *All agents who simultaneously (at the same coordination iteration  $i \in \mathbb{N}_0$ ) plan to use a resource at the same time step  $t$  submit a bid. The bid  $b^r(i)$  that each contestant  $r$  submits equals  $l^r(i) - s^r(i)$ . Here,  $l^r(i)$  is the cost  $r$  expects to experience (given its current belief in  $i$  of the available resources) if it would lose the resource. And,  $s^r(i)$  is the cost  $r$  expects (given its current belief of the available resources) to incur if it can keep using the contested resource. The winner is determined to be the agent who submits the highest bid. If multiple agents have greater or equal high bids than all the other ones ( $|\arg \max_{a \in \mathcal{A}} b^a(i)| \geq 2$ ), the robot with the highest index wins.*

To gain an intuitive motivation for the bidding rule, notice the bid quantifies the regret an agent expects to have for losing the auction (given its current belief of the availability of resources). Acknowledging that  $s^{winner}(i) + \sum_{a \in losers} l^a(i)$  is the estimated social cost (based on current beliefs of available resources) after the auction, we see that the winner determination rule greedily attempts to minimize social cost:  $\forall r : b^w(i) \geq b^r(i) \Leftrightarrow \forall r : s^r(i) + \sum_{a \neq r} l^a(i) \geq s^w(i) + \sum_{a \neq w} l^a(i)$ .

Notice, there are several degrees of freedom regarding the architectural implementation of the mechanism. For instance, to detect a conflict, all agents communicate their current plans to all other agents. With broadcast messages the communication effort per coordination iteration is hence in  $O(A)$  where  $A$  is the number of agents. Then each agent would be responsible to detect the next conflict and arrange an auction with the other agents. Alternatively, the mechanism designer could set up a number of additional dedicated conflict detectors and auctioneers (e.g. one for a set of time steps or a set of resources).

Before applying our mechanism to continuous distributed control under uncertainty in Sec. 7 we devote the next section illustrating its behavior in deterministic graph routing.

## 5. Coordinated Path Planning in Graphs

In this section, we will discuss our mechanism's properties in the context of path planning in finite graphs.

Graphs are mathematical abstractions that are simple but reflect the essence of many real-world planning scenarios. A graph  $G = (V, E)$  is a pair consisting of a set  $V$  of *vertices* or *nodes* and a set  $E \subset V^2$  of edges. The edges impose a relation structure on the vertices. In a robot path planning scenario the vertices could correspond to locations. Assuming

## 5. Coordinated Path Planning in Graphs

discretized time we could construct  $G$  such that  $(v, v') \in E$  iff a robot can travel from location  $v$  to  $v'$  in one time step. Finally, we assume the robot incurs a cost  $c^e > 0$  for traversing an edge  $e \in E$ . Depending on the objective, such a cost can model the time delay (e.g.  $\gamma(e) = 1$  [sec]) for moving from  $v$  to  $v'$  (where the vertices are chosen such that  $e = (v, v')$ ).

As an illustration, consider a simple graph routing example. Two agents desire to find low-cost paths in a graph with transition costs as depicted in Fig. 5.1(a). Agent 1 desires to find a path from Node 1 to 5, Agent 2 from Node 2 to 6.



Figure 5.1.: Two examples. Numbers next to the edges denote the transition costs. Ex.2 : S(a)/D(a): start/destination of agent a. Coordinated plans depicted in blue (Agent 1) and cream (Agent 2) which happens to be socially optimal.

In the first iteration ( $i = 1$ ), Agent 1 and Agent 2 both assume they can freely use all resources (nodes). Solving a binary linear program they generate their shortest paths as  $p^1 = (1\ 3\ 4\ 5\ 5\dots)$  and  $p^2 = (2\ 3\ 4\ 6\ 6\dots)$ , respectively. Detecting a conflict at time step 2 and 3, the agents enter an auction for contested Node 3. Agent 1’s estimated “detour cost” for not winning Node 3 (assuming he will be allowed to use all other nodes in consecutive time steps) is 2 which he places as a bid  $b^1(i) = 2$ . On the other hand, Agent 2’s detour cost is  $b^2(i) = l^2(i) - s^2(i) = 12 - 4 = 8$  and hence, she wins the auction. Having lost, Agent 1 adds a constraint to his description of his feasible set (more precisely to  $R$ ) that from now on prevents it from using Node 3 in time step 2. Replanning results in updated plans  $p^1 = (1\ 4\ 5\ 5\dots)$  and  $p^2 = (2\ 3\ 4\ 6\ 6\dots)$ . Being conflict-free now, these plans can be executed by both agents.

Notice how the laziness of our method protected us from unnecessary computational effort: the initial conflict at time 3 (Node 4) was implicitly resolved by the first auction without the need to set up an explicit auction for Node 4 or bidding on all combinations of availability of Nodes 3 and 4.

Of, course, this positive effect of laziness may not always bear fruit - in several situations resolving a collision at one node may not prevent collisions from happening (or, trigger new ones) at other nodes. As an example consider Ex. 2 in Fig. 5.1(b) and assume Agent 2’s initial plan visits Vertex I11 - after this conflict is resolved there will be a second at




Figure 5.2.: Results of comparison between different methods over 2000 randomized problem instances. The bars represent the percentages of the trials where... 1:  $\Gamma(AUC) \leq \Gamma(FPM)$  2:  $\Gamma(AUC) < \Gamma(FPM)$  3:  $\Gamma(AUC) \leq \Gamma(BEST - FPM)$  4:  $\Gamma(AUC) < \Gamma(BEST - FPM)$  5:  $\Gamma(AUC) = \Gamma(OPT)$  6:  $\Gamma(FPM) = \Gamma(OPT)$ .

Vertex I21.

Nonetheless, Ex. 1 was designed to provide an intuition that it often can lead to favorable coordination outcomes. In Sec. 7, we provide an experimental investigation of the number of collisions triggered in a typical multi-robot path planning scenario.

**Comparison of social cost on randomized graphs.** As explained above the myopic and lazy nature of our method may save computational effort during coordination possibly at the price of higher social cost. On the other hand, its coordination effort may at times be higher than that of fixed priority methods, so the overhead only seems justifiable if resulting in lower social cost.

To obtain a first assessment of our method’s (AUC) performance we compared it against the fixed priority method (FPM). The priorities were the same as the internal priorities in our methods, i.e. equivalent to the robots’s indices. As an absolute performance benchmark, we compared both methods to the optimal solutions computed by a centralized BLP solver (CS).

The comparisons were conducted on 2000 randomized graph planning problems. In each randomized trial, the planning environment was a forward directed graph similar in structure to the one in Fig. 5.1(b) (b). Each graph had a random number of vertices ( $L \times N$  - graphs where number of layers  $L \text{ Unif}(\{3, \dots, 11\})$ , number of nodes per layer  $N \text{ Unif}(\{3, \dots, 11\})$  ) and randomized vertex-transition costs drawn from  $\text{Unif}([1, \dots, 200])$ . The coordination task was to have each robot find a cost optimal path through the randomized graph where the robots had a randomized start location in the first layer and a destination vertex in the last layer.

For each trial we compared the social costs of the plans generated with the different coordination methods. The results are depicted in Fig. 5.2.

For a given problem instance, let  $\Gamma(AUC)$ ,  $\Gamma(FPM)$ ,  $\Gamma(OPT)$  denote the social cost of the coordinated plan generated by our method, the fixed priority methods FPM, and the optimal social cost, respectively. Finally, let  $\Gamma(BEST - FPM)$  be the social cost the plans of the fixed priority method with the best choice of priorities in hindsight would have incurred.

The data show that our auction method performed optimally on 93 % of the problems (5) while the fixed priority method did so on only 62.2 %. Conversely, the fixed priority method outperformed our method only on 2.1 % (see bar (1)) while it was strictly outperformed on 35.1 % of the randomized trials (2).

## 6. Modified Bids and Termination Guarantee

Many iterative coordination methods suffer from the limitation that they cannot be guaranteed to terminate in a finite number of coordination steps which can for instance happen due to running into deadlocks or cycles.

In this section, we will present a slight modification of our mechanism and discuss conditions under which it is guaranteed to terminate with a valid coordination solution when applied in environments with a finite number of states (resources) but where we do allow an infinite planning horizon.

In lieu of the graph planning problems introduced in the previous section, we assume that the planning environment is a graph  $G = (V, E)$ . However, in order to prove termination of our mechanism we make some additional assumptions. **In particular, we assume**  $G$  is connected, finite, edge relation  $E$  is reflexive, the edge costs are uniformly 1,  $\forall v \in V : deg(v) \geq |\mathcal{A}|$  each vertex has at least as many children as there are robots. Finally, we assume that a robot that reaches its goal node is effectively removed from the system in subsequent time steps, i.e. seizes claiming any further resources (nodes) other robots might want to visit after it has reached its goal.

In the modified version of our mechanism bidding for resources (nodes) will usually proceed as described in Sec. 4. As a new exception to the normal rule, whenever a robot  $r$  bids for the right to visit a node  $v \in V$  in time step  $t$  but has lost the right to use some node for some later time step  $t + \tau$  ( $\tau \geq 1$ ) or has lost  $v \in V$  for time step  $t$  to some other robot in a previous coordination iteration (who may have released it again in the mean time) then  $r$  will submit the modified bid:  $b^r(i) = 1 - \delta(l^r(i) - s^r(i))$  where  $\delta : \mathbb{R} \rightarrow \{0, 1\}$  denotes the Kronecker delta function and  $l, s$  are defined as usual (cf. Sec. 4).

**Theorem 6.0.1 (Convergence Guarantee )** *Under the assumptions listed above, our modified auction mechanism generates a joint plan in a finite number of coordination iterations that is legal (i.e. two consecutive nodes in the plan sequence of nodes are always*

connected by an edge in  $E$ ), collision-free and allows all robots to reach their destinations, regardless of the (mutually exclusive) start and goal nodes.

The formal proof is technical and rather lengthy and is omitted at this point. However, we will provide a detailed discussion and derivation of a generalized version of the modified algorithm in Sec. A.

## 7. Distributed control in a spatially continuous world and under uncertainty

### 7.1. Preliminaries- Sampling-based control and obstacle avoidance

Multi-robot motion planning and control problems in continuous maps have been addressed with mixed-integer linear programming (MILP) techniques [26]. Typically they rely on time-discretization only, without prior space-discretization. However, they are commonly solved with a centralized planner and typically do not take uncertainty into account. Recently, stochastic control methods have been suggested for single-robot path planning that accommodate for uncertainty in the effect of control signals. For instance, Blackmore et. al. [5] discuss a particle-based method that can be used to generate a low-cost trajectory for a vehicle that avoids obstacles with adjustably high confidence. In their model, the plans  $p^a$  are time-discrete sequences of control inputs. The spatial location  $x_t^a$  of Robot  $a$  at time  $t$  is assumed to be a linear function of all previous control inputs plus some iid random perturbations  $\nu_0, \dots, \nu_{t-1} \sim \mathcal{D}$ . So, given plan  $p^a$ , drawing  $n$  samples of perturbations for all time steps generates  $N$  possible sequences of locations (*particles*)  $(x_t^{a,(j)})_t$  ( $j = 1, \dots, N$ ) Robot  $a$  could end up in when executing his plan.

Formally,  $x_t^{a,(j)} = f_t(x_0^{a,(j)}, u_0^a, \dots, u_{t-1}^a, \nu_0^{(j)}, \dots, \nu_{t-1}^{(j)})$  ( $j = 1, \dots, N$ ) where  $f_t$  is a linear function and  $u_0^a, \dots, u_{t-1}^a$  is a sequence of control inputs as specified by Robot  $a$ 's plan. Due to this functional relationship we can constrain Robot  $a$ 's MILP's search for optimal control inputs by adding constraints on the particles.

Let  $T$  be the number of time steps given by the time horizon and temporal resolution. That is,  $t \in \{1, \dots, T\}$ . Furthermore, let  $F$  be the free-space, i.e. the set of all locations that do not belong to an obstacle. Obstacle avoidance is realized by specifying a chance constraint  $\Pr((x_t^a)_{t \in T} \notin F) \leq \delta$  on the actual location of the robot. For practical purposes,  $\Pr((x_t^a)_{t \in T} \notin F)$  is estimated by Monte-Carlo approximation leading to the approximated chance constraint  $\frac{1}{N} |(x_t^{a,(j)})_{t \in T} \notin F, i = 1, \dots, N| \leq \delta$  which we add to Robot  $a$ 's individual MILP [5].

If  $\mathcal{D}$  is a unimodal and light-tailed distribution (e.g. a Gaussian), the particles  $x_t^{a,(1)}, \dots, x_t^{a,(N)}$  for  $a$  at time step  $t$  typically form a cluster mostly centered around the mean.

Note that the uncertainties due to the random perturbations accumulate over time. Hence, the standard error of the particle clusters along a robot’s trajectory can be expected to increase with  $t$ .

## 7.2. Multi-Robot motion control under uncertainty

As collision-free plans are found by solving a MILP we could combine both approaches to a multi-robot stochastic control mechanism: Integrating the individual MILP’s into one large central MILP (cf. to Eq. 3.2 in Sec. 3) we could then add an appropriate inter-robot constraint for each combination of particles in order to avoid collisions. Unfortunately, the number of integer constraints would grow superlinearly in the number of particles and even exponentially in the number of robots, rendering this approach computationally intractable.

Instead, we propose to apply our mechanism as follows: Each robot solves its local MILP to find a plan that corresponds to sequences of  $n$  particle trajectories. When two (or more) robots  $a, r, \dots$  detect their particle clusters  $\{x_t^{a,(1)}, \dots, x_t^{a,(N)}\}, \{x_t^{r,(1)}, \dots, x_t^{r,(N)}\}, \dots$  ‘get too close’, they suspect a conflict and participate in an auction. The winner gets to use the contested region, while the losers receive constraints that correspond to a *virtual obstacle* (that is valid for time step  $t$ ) and replan. Notice, for notational convenience, we omit the explicit mention of the coordination iteration  $i$  in our notation throughout the rest of the section.

Next, we will explain the application of our mechanism to the continuous path planning problem in greater detail. Every robot employs the path planning algorithm as described in [5] to generate a particle-trajectory that is optimal for him. As explained in Sec. 4 the mechanism requires the robots to exchange their plans in every coordination iteration. However, they do not need to exchange all particles constituting their trajectories – it suffices only to exchange the optimal control inputs that lead to the particle trajectories (alongside the state or seed of their own pseudo-random-generator with which they drew their disturbance parameters).

With this knowledge all the other robots are able to exactly reconstruct each others’ particle trajectories. Now each robot locally carries out a test for collision by calculating the probability of a collision for each plan of every other robot.

Let  $\{x_t^{a,(1)}, \dots, x_t^{a,(N)}\}$  be the particle cluster that probabilistically describes the desired position of Robot  $a$  at time step  $t$ . Furthermore, let  $\{x_t^{r,(1)}, \dots, x_t^{r,(N)}\}$  be the particle cluster of Robot  $r$ . Let  $\epsilon$  be a predetermined parameter representing the minimum distance allowed between two robots. For instance, we could set  $\epsilon = 2d$  where  $d$  is the diameter of the robots which is a reasonable choice when defining a robot’s location as the cartesian coordinates of his center point.



The probability of a collision of Robot  $a$  and Robot  $r$  at time step  $t$  is

$$\Pr(\|x_t^a - x_t^r\| < \epsilon) = \mathbb{E}_{x_t^a, x_t^r} \{\chi_C\} = \int \int \chi_C(x_t^a, x_t^r) f(x_t^a) f(x_t^r) dx_t^a dx_t^r \quad (7.1)$$

$$\approx \frac{1}{N^2} \sum_{k=1}^N \sum_{j=1}^N \chi_C(x_t^{a,k}, x_t^{r,j}) \quad (7.2)$$

where  $f(x_t^a)$  and  $f(x_t^r)$  are the densities representing the uncertainty regarding Robot  $a$ 's and Robot  $r$ 's locations, respectively, given the histories of their control inputs and where

$$\chi_C(x_t^a, x_t^r) := \begin{cases} 1 & , \text{ for } \|x_t^a - x_t^r\| < \epsilon \\ 0 & , \text{ otherwise.} \end{cases}$$

Therefore, the probability of a collision of Robot  $a$  and Robot  $r$  at time step  $t$  is approximated by their respective particle representations. If this approximated probability is above a predefined threshold  $\delta$ , the robots engage in an auction for the contested spatial resource, as described in previous sections. The resource in this case corresponds to the right to pass through. We propose its denial to be embodied by a new *virtual obstacle* the loser of the auction, say Robot  $r$ , will have to avoid (but only at time  $t$ ). By placing the virtual obstacle around the winner's location estimate at time step  $t$ , we will reduce the chance of a collision. We represent the new obstacle by a square (if planning takes place in higher dimensions a hypercube)  $B_{\alpha+\epsilon}(\bar{x}_t^a)$  with side length  $\alpha + \epsilon$  and centered at the sample mean  $\bar{x}_t^a$  of Robot  $a$  at time step  $t$ . The choice of this representation is motivated by the fact that the chance constraints for a square-obstacle can be encoded by merely four linear and a few additional integer constraints [4, 5].

Obviously, the larger the virtual obstacle, the lower the probability of a collision between the robots. On the other hand, an overly large additional obstacle shrinks the free-space and may unsuitably increase path costs or even lead to deadlocks. Next, we will derive coarse mathematical guidelines for how to set the size of the virtual obstacle in order to avoid a collision with a predefined probability.

Let  $t$  be a fixed time step. Let  $C := \{(x_t^a, x_t^r) | (\|x_t^a - x_t^r\| < \epsilon)\}$  be the event of a collision and  $E := \{(x_t^a, x_t^r) | \|x_t^a - \bar{x}_t^a\|_2 \leq \alpha\}$  the event that the true position of Robot  $a$  at time step  $t$  deviates no more than  $\alpha$  from the mean of its position estimate given by sample mean  $\bar{x}_t^a$ . By introducing a chance constraint with threshold  $\frac{\delta}{2}$ ,

$$\Pr[x_t^r \in B_{\epsilon+\alpha}(\bar{x}_t^a)] < \frac{\delta}{2} \quad (7.3)$$

we enforce a bound on the collision probability. Introduction of the virtual obstacle to Robot  $r$ 's constraints induces his planner to adjust the control inputs such that the fraction of particles  $(x_t^{r,(j)})_{j=1,\dots,N}$  that are inside the square box  $B_{\epsilon+\alpha}(\bar{x}_t^a)$  with edge length  $\alpha + \epsilon$  around sample mean  $\bar{x}_t^a$  is bounded (and by particle approximation of the chance constraint, hence also the (approximated) probability that Robot  $r$  is inside the box).

Parameter  $\alpha$  needs to be specified after the desired  $\delta$  is defined and we will now discuss a proposal how this can be done.

Let  $K$  be the event  $\{(x_t^a, x_t^r) | x_t^r \in B_{\epsilon+\alpha}(\bar{x}^a)\}$ .

We have  $\Pr(C) = \Pr(C \cap E) + \Pr(C \cap \neg E) = \Pr(C \cap E \cap K) + \Pr(C \cap E \cap \neg K) + \Pr(C \cap \neg E) = \Pr(C \cap E \cap K) + \Pr(C \cap \neg E)$  where the last equality holds since  $\Pr(C \cap E \cap \neg K) = 0$ . Furthermore,  $\Pr(C \cap E \cap K) \leq \Pr(K)$  and  $\Pr(C \cap \neg E) \leq \Pr(\neg E)$ . Hence,

$$\Pr(C) \leq \Pr(K) + \Pr(\neg E) \quad (7.4)$$

Due to chance constraint (7.3) we know that control inputs are found that (for sufficiently large  $N$ ) ensure that  $\Pr(K) < \frac{\delta}{2}$ . Hence, all we are left to do is to determine box parameter  $\alpha$  such that  $\Pr(\neg E) \leq \frac{\delta}{2}$ .

### 7.2.1. Collision bounds assuming isotropic Gaussian noise

For now, let the distributions of Robot  $a$  be an isotropic Gaussian with covariance matrix  $\Sigma = \sigma^2 I$  where  $I$  is the identity matrix. We can then control  $\Pr(\neg E)$  by computing the  $\sigma$ -bounds of the normal distribution (considering the masses of its tails). For instance, an upper bound  $\frac{\delta}{2} = 0.05$  on the collision probability can be achieved by setting  $\alpha := 2\sigma$  and a bound of 10 percent by setting  $\alpha := 1.64\sigma$ . Of course such an ad hoc method would not work if we can sample from the noise distributions but are oblivious of their analytical nature and it would be desirable to have collision bounds that are distribution-independent. Deriving such bounds is our next objective.

### 7.2.2. Distribution-independent collision bounds

Since  $\|\cdot\|_2 \leq \frac{1}{\sqrt{2}} \|\cdot\|_\infty$ , we have  $\Pr(\neg E) = \Pr(\|x_t^a - \bar{x}_t^a\|_2 \geq \alpha) \leq \Pr(\frac{1}{\sqrt{2}} \|x_t^a - \bar{x}_t^a\|_\infty \geq \alpha) = 1 - \Pr(\|x_t^a - \bar{x}_t^a\|_\infty \leq \sqrt{2}\alpha) =: P_t^a(\alpha)$ . Utilizing Whittle's generalization of Chebyshev's inequality [28] yields an upper bound  $\beta(x_t^a, \alpha)$  on  $P_t^a(\alpha)$ . For the two-dimensional case we have  $\Pr(\neg E) \leq P_t^a(\alpha) \leq \beta(x_t^a, \alpha)$

where

$$\beta(x_t^a, \alpha) = \frac{1}{\alpha^2} (c_{t,11}^a + c_{t,22}^a) + \frac{1}{\alpha^2} \sqrt{(c_{t,11}^a + c_{t,22}^a)^2 - 4(c_{t,12}^a)^2}$$

and  $c_{t,ij}^a$  denotes the covariance of  $x_t^a$  between dimensions  $i$  and  $j$ .

For a larger number of particles (so we can expect  $\bar{x}_t^a$  to coincide with the true mean) this bound constitutes a formula, describing how to set box-size parameter  $\alpha$  to ensure  $\Pr(\neg E) < \frac{\delta}{2}$ .

Finally, by referring to Eq. 7.4 and Eq. 7.3 we see that we have

$$\Pr(C) \leq \frac{\delta}{2} + \Pr(\neg E) \leq \frac{\delta}{2} + \beta(x_t^a, \alpha) \quad (7.5)$$

which provides a recipe that allows us to bound collision probability  $\Pr(C)$  below free parameter  $\delta$  by adjusting virtual obstacle parameter  $\alpha$ .

Note, since approximate bound  $\beta$  is distribution-independent the bound holds for any noise distribution that governs our uncertainty.

### 7.3. Experiments

We consider three different path planning

scenarios, all with planning horizon of length ten, in our simulations:

- A simple example with only two robots to illustrate the very basic functionality of the mechanism.
- A quantitative evaluation of the average runtime behaviour for an increasing number of robots in an environment with a fixed number of obstacles.
- A quantitative evaluation of the average number of conflicts to be resolved by the mechanism in an increasingly complex environment for a fixed number of robots.

In all simulations the sample distribution for the robots was chosen as isotropic zero-mean white Gaussian noise with standard deviation  $\sigma = 0.001$ .

For an illustration, consider the simulations of a two-robot planning scenario depicted in Fig 7.1. Here two robots 1 and 2 started at locations at the bottom of a map. When

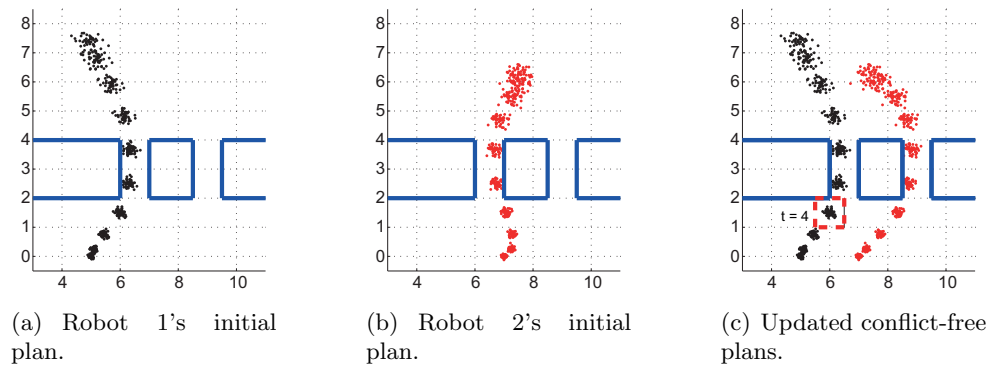


Figure 7.1.: Simple example. Blue box: obstacle. Dashed box: virtual obstacle for Robot 2 for time step 4 (after he lost an auction against Robot 1).

generating paths to destinations at the far side of the map, they desired to avoid the obstacles (blue rectangles). Their control inputs were accelerations and their state space consisted of locations and velocities. Each robot's cost function quantified the sum of expected  $\ell_1$  distances to the robot's destination of the generated trajectory.

Planning independently with the particle-control method, the robots found their individually cost-optimal trajectories as depicted in Figs. 7.1(a) and 7.1(b). Note, how the spread

of their particle clusters increases as the uncertainties accumulate over time. Getting too close to each other at time step four (i.e. causing our coll. probability estimate to exceed our threshold  $\delta$ ) and auction was invoked where Robot 1 was determined to be the winner. Hence, Robot 2 got a constraint corresponding to a virtual obstacle (dashed box) for time step 4 denying access through the left gap for  $t = 4$  and inducing him to instead take the (originally costlier) way around through the right gap (Fig. 7.1(c)).

It should be expected that the number of iterations of our mechanism depends on the number of collisions during coordination, which in turn, should increase with the number (and size) of obstacles (or decrease with available free-space) and the number of robots in the system. To develop an intuition for the dependence of run-time on these factors we conducted randomized experiments (with varying robot destinations and obstacle placements) in which run-time and number of collisions were recorded. The results for ten robots with varying starts, destinations and obstacles are depicted in the left part of Fig. 7.2.

In a third round of simulations, the obstacles were placed at fixed positions together with fixed, equally spaced, starting positions for the robots. In order to provoke potential conflicts, the robots' goals were drawn at random from a uniform distribution. We iteratively added more robots to the planning scenario and set up the mechanism to calculate conflict-free plans for varying numbers of robots. The results are depicted in the right plot of Fig. 7.2.

The simulations were implemented in MATLAB, with no particular emphasis on run-time optimization and all experiments were executed on a standard desktop computer. In summary, Fig. 7.2 illustrates that both the number of coordination iterations (collisions) and run-time increased moderately with increasing problem complexity.




Figure 7.2.: Left: Number of arising conflicts vs. varying number of obstacles. Right: Runtime in seconds vs. number of robots. Plots show averages and standard deviations over 50 Monte-Carlo runs of randomized problems.

---

## 8. Conclusions

In this paper, we presented a distributed, auction-flavoured multi-robot coordination mechanism. It is lazy in the sense that the agents only coordinate when necessary (due to conflicting resource usage), generating conflict free plans. We showed how such a myopic coordination mechanism can be applied to the MILP formulation of a multi-robot path planning problem, taking uncertainty about the robots' positions into account. Using our distributed mechanism in this scenario can be expected to be computationally more attractive, as we replace the intractable centralized problem by a succession of significantly more tractable single-robot MILPs.

Our simulations suggest that the overall coordination and communication overhead scales well in the number of agents and the complexity of the environment. For planning problems in continuous, uncertain environments we have linked our approach to a particle-based open-loop control method and have given probabilistic bounds on the collision probability as a function of the size of the virtual obstacles we introduced to achieve coordination.

Furthermore, for a slight modification of the bidding rules, we showed it is possible to give theoretical termination guarantees for a large class of graph planning problems (refer to [7]). While our mechanism will not always achieve socially optimal coordination, our initial experiments on graph planning problems suggest that it can on most problems and that our approach seems often advantageous over simple fixed priority methods that are still widely used in practice. A theoretical elucidation of the nature of problem instances where this is guaranteed to be (or not to be) the case, alongside more extended experimental validations, will be conducted in the context of future work.

## Acknowledgements

The authors would like to thank Achim Hekler for useful discussions during early stages as well as Stephen Roberts and Mark Ebdon for comments on related work and the final draft during late stages of this project. Jan Calliess is grateful for funds via the UK EPSRC "Orchid" project EP/I011587/1. Daniel Lyons gratefully acknowledges partial support by the German Research Foundation (DFG) within the Research Training Group GRK 1194 "Self-organizing Sensor-Actuator-Networks".

## A. Coordination in discrete graphs

If a coordination mechanism is to be applied in a time-critical scenario it can be important to consider the termination properties. That is, is the coordination method guaranteed to reach a coordination solution in a finite number of iterations? For a slight modification of our mechanism we can give an affirmative answer in environments with a finite number of states (resources) but a (countably) infinite planning horizon. Such planning problems can be formalized as path planning problems in graphs.

After formalizing the problem statement we provide a formal discussion of the circumstances under which we are indeed able to prove that our modified mechanism is guaranteed to achieve coordination in a finite number of iterations. The section concludes with a proof of this termination guarantee.

### A.1. Problem statement and conventions

We assume the following model: Let  $R = \{R_1, \dots, R_M\}$  be the set of resources (states/locations) in the system that robots  $1, \dots, A$  can claim in a discrete sequence of time steps  $t = 1, 2, \dots, \infty$ . Remember, safety dictates that resource claims are mutually exclusive (i.e. no two robots can simultaneously use the same resource). Therefore, we assume that the number of resources is at least as high as the number of robots,  $0 < A \leq M$ .

Assuming not all states are reachable from all other states in one time step, we can express the reachability (within one time step) as a sequence of relations  $E_t = (R \times R)_t$  ( $t \in \mathbb{N}_0$ ) where  $(R_u, R_v) \in E_t$  iff a robot can reach  $R_v$  at time step  $t + 1$  when it is in  $R_u$  at time step  $t$ .<sup>1</sup> Let  $Z_t^a$  denote the state robot  $a$  plans to occupy at time  $t$ . In this environment, a legal open-loop plan corresponds to a sequence of states  $(Z_t^a)_{t \in \mathbb{N}_0}$  such that  $\forall t \in \mathbb{N}_0 : (Z_t^a, Z_{t+1}^a) \in E_t$ . Furthermore, the *overall plan* now corresponds to the joint sequence  $(Z_t)_{t \in \mathbb{N}_0}$  where  $Z_t = Z_t^1 \times \dots \times Z_t^A$ .

We can express this model as a graph planning problem. Spatio-temporal graph  $\mathcal{G} = (V, E)$  consists of vertices  $V = \{v_{t,m} | t \in \mathbb{N}_0, m \in \{1, \dots, M\}\}$  where vertex  $v_{t,m}$  corresponds to state  $R_m$  at time  $t$ . So the graph consists of an infinite number of time layers (“rows”) indexed by  $t$  and each of the  $M$  “columns” corresponds to one state each. The set of edges reflects the reachability conditions defined above. That is, we assume  $(v_{t,l}, v_{t+1,m}) \in E$  iff  $(R_l, R_m) \in E_t$ . For an illustration refer to Fig. A.1.

---

<sup>1</sup>For simplicity, we tacitly assume homogenous robot capabilities that are known in advance.

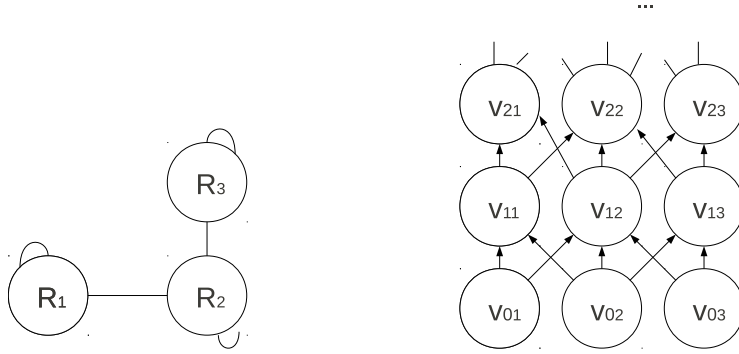


Figure A.1.: Left: Spatial graph with  $M = 3$  resources. Right: First three time layers (rows) of corresponding spatio-temporal graph.

With these notions in mind, robot  $a$ 's plan in coordination iteration  $i$  can be represented as a sequence of vertices in spatio-temporal graph  $\mathcal{G}$ . We denote this sequence as  $(z_t^a(i))$  where  $z_t^a(i) \in \{v_{t,1}, \dots, v_{t,M}\}$  is the vertex robot  $a$  intends to visit in time layer  $t$ . Notice the relationship  $z_t^a(i) = v_{t,m} \Leftrightarrow Z_t^a(i) = R_m$ .

When executing a plan that involves a state transition corresponding to using an edge in  $E$  the robot receives a cost specified by edge cost function  $\gamma^a : E \rightarrow \mathbb{R}_+$ . (Note, above-mentioned cost function  $c^a$  on the set of feasible plans can now be defined as the sum of all costs of edges (acc. to  $\gamma^a$ ) the plan would have to use to generate the implied sequence of nodes.)

At time  $t = 0$ , robots  $1, \dots, A$  are in their start states  $S(1) \in R, \dots, S(A) \in R$ , respectively ( $S(a) \neq S(r), \forall a \neq r$ ). That is they are all at distinct nodes in (time-) layer 0. The planning problem can now be expressed as finding a collection of low-social cost open-loop plans corresponding to legal sequences of states (nodes)  $(z_t^a)_{t \in \mathbb{N}_0}$  ( $a \in \mathcal{A}$ ) such that the local constraints are met, they are collision-free (i.e.  $z_t^a \neq z_t^r, \forall t \in \mathbb{N}_0 \forall a \neq r$ ) and the robots reach their destinations from their start states. Note, for convenience we will not always explicitly distinguish between a column  $(v_{0,m}, v_{1,m}, \dots)$  in the graph and the corresponding resource  $R_m$ .

## A.2. Further Assumptions

To show that our mechanism is guaranteed to achieve coordination, we will consider the system dynamics. Notice, there are different facets we can consider: (I) the *coordination dynamics* which we define to govern the evolution of plans or global state assignments  $(z_t)_{t \in \mathbb{N}_0}(i)$  as a function of iteration index  $i$ , and (II) the *execution dynamics* which we define to dictate the evolution of states  $z_t(i)$  as a function of time index  $t$  for a given, fixed iteration  $i$ . Below, we will consider an intermediate level of dynamics. We will elucidate

the execution dynamics until a finite time horizon for iterations that have established the absence of conflicts until that horizon.

Our mechanism can be demonstrated to empirically perform well in real-world settings. However, in order to be able to deliver a formal proof that establishes a guaranteed coordination success in a finite number of coordination iterations we introduce a number of additional assumptions:

- (A1) There exists a path from every vertex (location) to every other vertex in later time layers of the spatio-temporal graph.
- (A2) Every vertex has at least  $|\mathcal{A}|$  children, where  $\mathcal{A}$  is the set of robots competing for resources.
- (A3) State transition costs are uniformly positive. That is, all agents incur constant costs for using an edge:  $\gamma^a(e) = k \in \mathbb{R}_+, \forall a \in \mathcal{A}, e \in E$
- (A4) Once a robot reaches its goal, it disappears from the system/planning problem (i.e. discontinues to claim any resources in subsequent time steps).
- (A5) The robots' individual planning algorithms are suitable and produce only legal plans. Furthermore, given the number of available resources each robot's planner is always capable of finding a cost optimal path between start state and destination state given the robots knowledge of available states.
- (A6) Robots can opt to remain where they are. That is, the transition relation is reflexive:  $\forall t \in \mathbb{N}_0, m \in \{1, \dots, M\} : (v_{t,m}, v_{t,m}) \in E$ .
- (A7) The environment is static, i.e.  $\forall t \in \mathbb{N}_0 : E_t = E_{t+1}$ .
- (A8) The reachability relation is delayed symmetric, i.e.  $\forall j, m \in \{1, \dots, M\} \forall t : (R_m, R_j) \in E_t \Rightarrow (R_j, R_m) \in E_{t+1}$ .

Assumption A1 is necessary to ensure all mutually exclusive configurations of start and destination locations correspond to a feasible plan resulting in a path that connects them. Positivity of costs required in A3 is reasonable in virtually all physical systems. If we lifted it (see below) this can result in simplifications guaranteeing optimal solutions of our mechanism for certain cost maps and in robots never having an incentive to reaching their destinations (also see the discussion below). The requirement that edge cost is a constant implies that a cost-optimal path is also a path with a minimal number of hops (vertex visitations). Assumption A2 will prove important to ensure that any robot that loses a resource actually can be guaranteed to be able to execute an alternative plan not involving this resource under all circumstances. Finally, A4 is an assumption mostly found in multi-robot coordination papers that provide theoretical guarantees. It ensures that no robot's destination is occluded by other (idle) robots that already reached their goals.

While assumption A4 obviously does restrict the applicability to some settings. Still, there many scenarios where A4 is met or which can be modified accordingly (eg by an external robot pick-up mechanism). For formal reasons, we model A4 by defining that if



Robot  $r$  reaches its destination in iteration  $T_r$  along a conflict-free path then it plans on using a virtual dummy resource  $R_{m+r}$  (which no other robot ever claims) from then on:  $Z_{T_r}^r(i) = D(r) \Rightarrow Z_{T_r+t}^r = R_{m+r}, \forall t \in \mathbb{N} - \{0\}$ . Notice, it would be equally sufficient to lift A4 and replace it by another assumption stating that the destinations of the robots are never blocking the pathways of other robots to their destinations. This could be assured if the robots' destinations are dead ends (in which case we would adapt A2 to exclude destinations). For instance, in a multi-vehicle scenario the destinations could be distinct car ports.

A5 simply states that the robots are capable of making individually optimal decisions. In the graph planning scenario they could simply employ prominent shortest-path algorithms such as  $A^*$  or Dijkstra's algorithm [12]. It is obvious that this requirement is important since our mechanism relies on the cost estimations of the local planning algorithms to compute the bids and since all robots need to have working planning algorithms to compute the paths that lead them to their destinations.

### A.3. Modified mechanism and termination guarantee

**Definition A.3.1 ((Individual) Loss Horizon (ilh))** *For coordination iteration  $i$ , Robot  $r$ 's current individual loss horizon  $ilh^r(i)$  is defined as the largest time step where Robot  $r$  has lost a resource to a competing agent in iteration  $i$  or in previous coordination iterations. We say that the vertices  $v_{t,m}$  where  $t \in ilh^r(i)$  are on the individual loss horizon and define  $ILH^r(i) := \{v_{t,m} \in V | t \in ilh^r(i)\}$  to be the set of these spatio-temporal vertices.*

Note, that in iteration  $i$ , every robot  $r$  believes that all spatio-temporal vertices  $v_{t,m}$  with  $t > ilh^r(i)$  are available (of course, it may turn out that  $r$  loses  $v_{t,m}$  in a subsequent iteration  $i' > i$  in which case  $ilh^r(i) < ilh^r(i')$ ).

#### A.3.1. Modified bidding behaviour

As mentioned above we will modify our mechanism in order to establish provable convergence. In the modified mechanism each robot  $r$  maintains a list data structure  $L^r$  containing spatio-temporal vertices which he has lost in past iterations. We will call this list *loss horizon list* as it keeps track of  $r$ 's individual loss horizon. Furthermore, it serves as a memory of the iteration order in which spatio-temporal vertices have been initially lost. The bidding mechanism will be as described in Sec. 4 for the most recently lost vertex and vertices beyond a robot's loss horizon but will ensure that bids for all other vertices are restricted to  $\{0, k\}$ . While this property is a restriction in bidding flexibility (and possibly may result in higher social cost for certain problem instances) we leverage this restriction in our proof of the convergence theorem given below.

At the beginning of the coordination process each robot starts out with an empty list  $L^r$ . Whenever  $r$  loses a spatio-temporal vertex, say  $v_{t,m}$ , which is not already in the list, he

inserts it to the first position of the list such that there is no vertex with larger time index  $t' > t$  behind  $v_{t,m}$  in the list. Hence, if  $L_u^r$  denotes the  $u$ th element of the list then we have  $\forall u < u' \in \mathbb{N} : ((L_u^r = v_{t,m} \wedge L_{u'}^r = v_{t',j} \Rightarrow t \geq t')$  and (if  $t = t'$  then  $v_{t,m}$  was lost in a later coordination iteration than  $v_{t',j}$ ). By construction, the first entry  $\text{top}(L^r)$  of list  $L^r$  is the vertex of the highest time layer that  $r$  lost (for the first time) in the most recent iteration.

For reasons that are important for the proof of Thm. A.3.2 below, the loss horizon list is constructed to never shrink— although it is permissible to remove all nodes of lower time steps than  $\text{top}(L^r(i))$  to save storage space.

The key difference between the generic mechanism described in Sec. 4 comes to bear when we consider bidding. Assume Robot  $r$  bids for vertex  $v_{t,j}$  in iteration  $i$  and let  $v_{t',m} = \text{top}(L^r)$ . If  $t > t'$  or  $v_{t,j} = v_{t',m}$  then  $r$  submits the usual bid

$$b^r(i) = l^r(i) - s^r(i)$$

as described in Sec 4. Otherwise however,  $r$  submits the bid

$$b^r(i) = k - k \delta(l^r(i) - s^r(i)) \tag{A.1}$$

where  $k$  is the constant transition cost acc. to A3 and  $\delta : \mathbb{R} \rightarrow \{0,1\}$  denotes the Kronecker delta function. Therefore, bids below the individual loss horizon are always in  $\{0, k\}$  and can only qualitatively express whether losing a node is a detour ( $b^r(i) = k$ ) or not ( $b^r(i) = 0$ ). By contrast, bids for vertices beyond  $ilh^r(i)$  express a quantification of the magnitude of the detour cost losing the auction is anticipated to entail. We will use this bid construction to show that the bids consequently always are in  $\{0, k\}$  which in turn is a stepping stone in proving convergence Thm. A.3.2 stated below.

### A.3.2. Termination guarantee

**Theorem A.3.2 (Termination Guarantee)** *Provided the definitions and assumptions above hold, our modified mechanism generates a joint plan in a finite number of iterations that is legal, collision-free and allows all robots to reach their destinations, regardless of the (mutual exclusive) start states. That is:*

$$\forall \text{ injective } S, D : \mathcal{A} \rightarrow R \exists T^1, \dots, T^A, i \in \mathbb{N}_0 \forall a, r \in \mathcal{A}, a \neq r, t : z_t^a(i) \neq z_t^r(i) \wedge Z_{T^a}^a(i) = D(a) \wedge Z_0^a(i) = S(a).$$

Before attempting a formal proof we will give an informal outline as follows: The motivation behind augmenting the bidding rule by the one in Eq. A.1 was to ensure that we can show that all bids for a vertex whose loss would result in an increase in the path length of the bidders all equal  $k$  (see Lemma A.3.3 given below). Therefore, the robot  $A$  with the highest priority only loses bids for resources when there is an alternative route to his destination of equally short path length (avoiding the lost spatio-temporal vertex).

| Notation               | Meaning   |
|------------------------|---|
| $\mathcal{A}$          | set of robots   |
| $A$                    | number of robots competing for resources ( $A \leq  \mathcal{A} $ )                               |
| $a, r$                 | indices of robots $r$ and $a$   |
| $t$                    | ordinal time index  |
| $i$                    | coordination iteration index  |
| $R_1, \dots, R_M$      | resources / locations / nodes in spatial graph  |
| $\mathcal{G} = (V, E)$ | spatio-temporal graph   |
| $E_t$                  | edges connecting temporal layers $t$ and $t + 1$ .  |
| $v_{t,m}$              | spatio-temp. vertex in time layer $t$ corresp. to resource $R_m$ in spatial graph                 |
| $z_t^r(i)$             | spatio-temp. vertex in time layer $t$ $r$ plans (in coordination iteration $i$ ) to use.          |
| $Z_t^r(i)$             | spatial vertex (resource) used by $r$ at time $t$ in iteration $i$ .                              |
| $\delta$               | Kronecker delta   |
| $S(r), D(r)$           | start, goal vertex of robot $r$ .   |
| $k$                    | positive, uniform edge cost   |
| $ilh^r(i)$             | indiv. loss horizon of robot $r$ in coord. iteration $i$ (ref. Def. A.3.1)                        |
| $b^r(i)$               | bid robot $r$ submits in coord. iteration $i$   |
| $H(v_{t,m}, R_s)$      | min-hop-distance/ geodesic distance between $R_m$ and $R_s$                                       |
| $h_i^r(v_{t,m}, R_s)$  | $r$ 's hop-dist. between nodes in spatio-temp. graph given the history of vertices lost up to $i$ |
| $V_t^r(i)$             | set of reachable nodes in time layer $t$ $r$ has not lost up to coord. iteration $i$              |

Table A.1.: Reference table of frequently used notation.

Hence, this robot will always reach his destination in a finite number  $T^A$  of time steps, where  $T^A$  is the number of state transitions along a shortest path (cf. Lemma A.3.4 given below). Due to A4, Robot  $A$  is effectively removed from the path-planning process beyond time step  $T^A$  and we can show convergence of the remaining  $A - 1$  robots' coordination problem with a recursive argument.

### A.3.3. Derivation of Theorem A.3.2

Before commencing with our derivations we will need to establish a bit more notation and terminology that builds on the notions already defined above. For convenience a summary of the introduced notation is depicted in Tab.

We define the *min-hop-distance*  $H(v_{t,s}, R_d)$  between a spatio-temporal vertex  $v_{t,s}$  and a resource  $R_d$  as the minimum length of a path leading from  $v_{t,m}$  into the column representing  $R_d$ . The min-hop-distance is a property of the graph. It could also be defined as the geodesic distance or shortest path length between  $R_s$  and  $R_d$  in the spatial graph.

By contrast, for robot  $r$ , we define his *conceived hop-distance*  $h_i^r(v_{t,s}, R_d)$  between a spatio-temporal vertex  $v_{t,s}$  and a resource  $R_d$  as the minimum length of a path leading from  $v_{t,m}$  into the column representing  $R_d$ , given  $r$ 's current belief about available vertices for time steps  $\tau \geq t$ , given the auction history of won and lost resources up to iteration  $i$ . The

term *hop – distance* is motivated by the fact that the *length* of the path measures the number of vertex transitions (“hops”). A *shortest* path between two vertices or resources is one with minimum length.

Before proving our termination guarantee, we will need to establish the following two lemmata.

**Lemma A.3.3** *Let  $k$  be the constant transition cost mentioned in Assumption A3. Assume, in iteration  $i$ , Robot  $r$  bids for a spatio-temporal vertex  $v_{t,m}$  where  $t \in \mathbb{N}, m \in \{1, \dots, M\}$ . Furthermore, assume A1-A5 hold and A6- A8 hold for all time layers of the spatio-temporal graph beyond the robot’s loss horizon (i.e.  $\forall E_\tau$  where  $\tau \geq ilh^r(i)$ ).*

Then we have

$$b^r(i) = l^r(i) - s^r(i) = \begin{cases} k, & l^r(i) > s^r(i) \\ 0, & \text{otherwise} \end{cases}$$

where  $b^r(i)$  is Robot  $r$ ’s bid in iteration  $i$  and  $l^r(i), s^r(i)$  are the path costs estimates under the assumptions that  $v_{t,m}$  is not or is available, respectively (cf. Sec. 4).

**PROOF.** If contested node  $v_{t,m}$  is below the robot’s individual loss horizon, by construction of the loss list  $L^r$ , the bid is  $b^r(i) = k - k\delta(l^r(i) - s^r(i)) \in \{0, k\}$  (cf. Sec. A, Eq. A.1). So, throughout the rest of the proof, we assume  $v_{t,m}$  is on or beyond  $r$ ’s individual loss horizon, i.e.  $t \geq ilh^r(i)$ .

Due to Assumption A3, there is a directly proportional one-to-one correspondence between length of a path (number of time steps until destination) and path cost. That is, any path of hop-length  $\lambda$  incurs a cost of  $k * \lambda$  where  $k$  is the uniform transition cost. Therefore, any min-cost path is also a shortest path (in terms of the number of hops).

Remember,  $l^r(i)$  is the cost of the min-cost path connecting  $S(r)$  and  $D(r)$  given robot  $r$ ’s belief (in iteration  $i$ ) about the available resources, under the condition that  $r$  will not be allowed to use  $v_{t,m}$ , i.e.  $l^r(i) = k(t-1) + k \min_{v \in V_t^r(i) - \{v_{t,m}\}} h_i^r(v, D(r))$  where  $V_t^r(i)$  is the set of all reachable (given the auction history) nodes in time layer  $t$  available to  $r$  (ie not lost). Similarly,  $s^r(i)$  is the cost of  $r$ ’s min-cost path under the condition that he assumes  $v_{t,m}$  to be available. That is,  $s^r(i) = k(t-1) + k h_i^r(v_{t,m}, D(r))$ . Notice, we can show with A2 that we have  $V_t^r(i) - \{v_{t,m}\} \neq \emptyset$ . Hence,  $b^r(i) = l^r(i) - s^r(i) = [\min_{v \in V_t^r(i) - \{v_{t,m}\}} k h_i^r(v, D(r))] - k h_i^r(v_{t,m}, D(r)) \in k\mathbb{N}_0$ . That is,

$$b^r(i) \in k\mathbb{N}_0.$$

Let  $v^* \in \arg \min_{v \in V_t^r(i) - \{v_{t,m}\}} k h_i^r(v, D(r))$  such that  $v^*$  would be chosen by  $r$ ’s individual planner if  $v_{t,m}$  were lost. Assumption A5 implies  $v_{t,m} \in \arg \min_{v \in V_t^r(i)} k h_i^r(v, D(r))$ , yielding

$$b^r(i) = [k h_i^r(v^*, D(r))] - k h_i^r(v_{t,m}, D(r)) \begin{cases} = 0, & h_i^r(v_{t,m}, D(r)) = h_i^r(v^*, D(r)) \\ \geq k, & \text{otherwise.} \end{cases}$$

Notice, for all vertices  $\{v_{\tau,j} | \tau \geq ilh^r(i) = t, j \in \{1, \dots, M\}\}$  on or beyond the individual loss horizon the min-hop-distance to the destination equals the conceived hop-distance. That is, we have  $H(v_{\tau,j}, D(r)) = h_i^r(v_{\tau,j}, D(r)), \forall \tau \geq ilh^r(i) \forall j$ . Hence,

$$b^r(i) = [k H(v^*, D(r))] - k H(v_{t,m}, D(r)) \begin{cases} = 0, & H(v_{t,m}, D(r)) = H(v^*, D(r)) \\ \geq k, & \text{otherwise.} \end{cases} \quad (\text{A.2})$$

Remember, we assume Robot  $r$  enters an auction in iteration  $i$  to compete for resource (vertex)  $v_{t,m}$  which is on or beyond its indiv. loss horizon. Compared to the resource  $Z_{t-1}^r(i)$  occupied in the previous time step, visiting contested vertex  $v_{t,m}$  would either result in

- (I) the same min-hop-distance ( $H(v_{t,m}, D(r)) = H(z_{t-1}^r(i), D(r))$ ),
- (II)  $r$  having moved closer to  $D(r)$  ( $H(v_{t,m}, D(r)) < H(z_{t-1}^r(i), D(r))$ ) or,
- (III)  $r$  having increased the min-hop-distance ( $H(v_{t,m}, D(r)) > H(z_{t-1}^r(i), D(r))$ ).

In the remainder of the proof we show  $b^r(i) \in \{0, k\}$  for each of the three cases.

*CASE I:* We have two subcases. (I.i) There is an alternative node  $v^* \in V_t(i)$  such that  $H(v^*, D(r)) = H(v_{t,m}, D(r))$  and hence,  $b^r(i) = 0$  (cf. Eq. A.2).

(I.ii) Due to A5 and since the robot had chosen  $v_{t,m}$  over  $v^* \in V_t(i)$ , the only alternative to (I.i) is  $H(v^*, D(r)) > H(v_{t,m}, D(r))$ . Let  $R_j := Z_{t-1}^r(i)$  be the resource occupied in time step  $t-1$  according to the current plan.  $(t+1) > ilh^r(i)$  implies that Robot  $r$  has not lost  $v_{t+1,j}$  in an auction and that Assumptions A7 and A8 hold in time layer  $t+1$ . A7 and A8 yield that  $v_{t+1,j}$  can be reached from  $v^*$  (incurring cost  $k$ ). Now,  $l^r(i) = k(t-1) + k + k H(v^*, D(r)) = k(t-1) + 2k + k H(v_{t+1,j}, D(r)) = k(t-1) + 2k + k H(v_{t,m}, D(r))$  where the last equality is implied by the fact that  $H(v_{t+1,j}, D(r)) = H(v_{t,j}, D(r)) = H(v_{t-1,j}, D(r)) = H(v_{t,m}, D(r))$  and the previous one both from  $H(v_{t+1,j}, D(r)) = H(v_{t,m}, D(r))$  and  $H(v^*, D(r)) = k + H(v_{t,m}, D(r))$ . On the other hand,  $s^r(i) = k(t-1) + k + k H(v_{t,m}, D(r))$ . Hence,  $b^r(i) = l^r(i) - s^r(i) = k$ .

*CASE II:* Obviously,  $r$ 's previous column in the spatio-temporal graph changes (i.e. he uses a different resource), since otherwise, the min-hop-distance could not decrease between time steps  $t-1$  and  $t$ . That is,  $R_m \neq R_j$  where  $z_{t-1}^r(i) = v_{t-1,j}$ . Losing  $v_{t,m}$  (ie  $R_m$  in time step  $t$ ) can either mean that (II.i) there is an alternative available vertex in time layer  $t$  (Robot  $r$  can route a path through) that reduces the min-hop-distance to the destination or, (II.ii) there is not. As always, let  $v^*$  denote the vertex  $r$ 's planning algorithm determines to be the second best to choose in time step  $t$  if  $v_{t,m}$  is unavailable.

First, assume (II.i). Then  $H(v_{t,m}, D(r)) = H(v^*, D(r))$  and thus (cf. Eq. A.2),  $b^r(i) = 0$ .

Conversely, assume (II.ii). Hence,  $l^r(i) > s^r(r)$ . Let resource index  $j$  be chosen such that

$R_j = Z_{t-1}^r(i)$ . We have  $t \geq ilh^r(i)$ .

Assume  $t = ilh^r(i)$ . Once again, we have multiple subcases: (a) In case  $v_{t,j}$  has been lost before, it is listed in loss horizon list  $L^r$  (cf. Sec. A.3.1). In that case, due to A5, the robot must have lost  $v_{t,m}$  in an even earlier iteration  $v_{t,j}$  (since it is the “better node”). Hence,  $v_{t,m} \in L^r, v_{t,m} \neq top(L^r)$  yielding  $b^r(i) = k$  acc. to Eq. A.1, as explained in Sec. A.3.1. (b) Otherwise, if  $v_{t,j}$  has not been lost it is available. That is, intuitively speaking, the robot can opt to keep using resource  $j$  also in time step  $t$  if beneficial (according to A6).

Likewise, if  $t > ilh^r(i)$ ,  $v_{t,j}$  has not been lost in an auction yet in previous iterations.

Since  $t+1 > ilh^r(i)$ ,  $v_{t+1,m}$  is also considered available by the robot. Since  $v_{t,m}$  was considered the best node to be taken right after  $v_{t-1,j}$  by  $r$ 's planning algorithm and since we assumed our planner to produce legal plans only, we have  $(v_{t-1,j}, v_{t,m}) \in E_{t-1}$ . Thus, by A8, we have  $(v_{t,j}, v_{t+1,m}) \in E_t$ . Hence (cf. Eq. A.2),  $b^r(i) = k H(v^*, D(r)) - k H(v_{t,m}, D(r)) = k H(v_{t,j}, D(r)) - k H(v_{t,m}, D(r)) = k H(v_{t,m}, D(r)) + k - k H(v_{t,m}, D(r))$  where the last equality is a direct consequence of A8.

*CASE III:* In conjunction with Assumption A6,  $H(v_{t,m}, D(r)) > H(z_{t-1}^r(i), D(r))$  implies  $ilh^r(i) = t$  and that  $r$  has lost all nodes  $v_t \in V_t(i)$  such that  $H(v_t, D(r)) < H(z_{t-1}^r(i), D(r)) \vee H(v_t, D(r)) = H(z_{t-1}^r(i), D(r))$ . Hence,  $H(v^*, D(r)) > H(z_{t-1}^r(i), D(r))$  where, as always,  $v^*$  is the next best node to  $v_{t,m}$  in terms of the number of remaining hops to the destination. Let  $v_{t-1,j} := z_{t-1}^r(i)$ . Since all nodes  $v_{\tau,j}$  such that  $\tau > t$  have not been lost (since  $t \geq ilh^r(i)$ ),  $v_{t+1,j}$  is available. Since  $(v_{t-1,j}, v^*) \in E_{t-1} \wedge (v_{t-1,j}, v_{t,m}) \in E_{t-1}$ , Assumptions A8 and A7 yield  $(v^*, v_{t+1,j}) \in E_t \wedge (v_{t,m}, v_{t+1,j}) \in E_t$ . Hence,  $H(v^*, D(r)) \leq k + H(v_{t-1,j}, D(r))$  and  $H(v_{t,m}, D(r)) \leq k + H(v_{t-1,j}, D(r))$ . In conjunction with  $H(v^*, D(r)) > H(v_{t-1,j}, D(r)) \wedge H(v_{t,m}, D(r)) > H(v_{t-1,j}, D(r))$  this yields  $H(v^*, D(r)) = H(v_{t-1,j}, D(r)) + k = H(v_{t,m}, D(r))$ . In conjunction with Eq. A.2 this yields  $b^r(i) = 0$ .

**q.e.d.**

**Lemma A.3.4** *Assume a graph planning problem with some prior auction history (comprising  $i' \in \mathbb{N}_0$  coordination iterations) consistent with our mechanism rules and where A1-A5 hold. Let  $A$  be the robot with the highest priority out of all resource-claiming robots  $\mathcal{A} = \{1, \dots, A\}$  currently in the system. Furthermore, assume A6- A8 are met for all time steps  $t \geq \min_{a \in \mathcal{A}} ilh^a(i')$ . Finally, for any  $i \in \mathbb{N}_0$  let  $T^A(i) \in \mathbb{N}_0$  such that  $Z_{T^A}^A(i) = D(A)$ . We have  $\forall i > i' : T^A(i) \leq \tau + H(z_\tau^A(i'), D(A)) \in \mathbb{N}_0$  where  $\tau := ilh^A(i')$ .*

**PROOF.**

Let  $i > i'$ . We show that, regardless of the spatial-temporal node  $z_t^A(i)$  which  $A$  plans to occupy in time layer  $t > \tau$ , he will always be able to decrease his min-hop-distance to his destination in subsequent time layers. That is,  $\forall t \geq \tau : H(z_t^A(i), D(a)) > H(z_{t+1}^A(i), D(a))$ .

Due to Lemma A.3.3 we know that, in all iterations  $i > i'$ , all robots submit a bid of  $k$  when losing a contested node would entail a detour and a bid of 0 when not. Hence, Robot  $A$  is among the highest bidders in an auction for a contested vertex  $v_{t,m}$  where losing  $v_{t,m}$  would increase the path cost to his goal. In such cases he is guaranteed to win any such critical resources due to the fact that he has the highest priority over competitors with equally high bids. Due to the equivalence of cost and number of hops (A3), this means that for all iterations  $i > i'$ ,  $A$  is able to make sure to gain any node beyond  $\tau = ilh^A(i')$  and due to planner optimality (A5), we have  $\forall t \geq \tau : H(z_t^A(i), D(a)) > H(z_{t+1}^A(i), D(a))$ . Remember, all reachable nodes beyond  $ilh^A(i')$  are available to Robot  $A$ . Since  $A$  is already in possession of spatio-temporal node  $z_\tau^A(i')$  in iteration  $i'$  and since he only loses auctions when not increasing the number of hops on his path between his start and destination, his total path length (in hops) will remain below the number of hops to  $z_\tau^A(i')$  (which equals  $\tau$ ) plus  $H(z_\tau^A(i'), D(A))$  during all subsequent coordination iterations. **q.e.d.**

Based on these lemmata and our assumptions, we are prepared to show Thm. A.3.2.

**Proof of Theorem A.3.2.** For Convenience, we repeat the theorem before proving it:

**Theorem A.3.2** *Provided the definitions and assumptions above hold, our modified mechanism generates a joint plan in a finite number of iterations that is legal, collision-free and allows all robots to reach their destinations, regardless of the (mutual exclusive) start states. That is:*

$$\forall \text{ injective } S, D : \mathcal{A} \rightarrow R \exists T^1, \dots, T^A, i \in \mathbb{N}_0 \forall a, r \in \mathcal{A}, a \neq r, t : z_t^a(i) \neq z_t^r(i) \wedge Z_{T^a}^a(i) = D(a) \wedge Z_0^a(i) = S(a).$$

**PROOF.**

Legality is enforced by Assumption A5. Since our coordination algorithm does not terminate until all plans are collision-free we only need to focus on termination in a finite number of coordination iterations. We prove this by induction on the number  $A$  of robots.

**Base case.**  $|\mathcal{A}| = 1$ : Is the single-robot case and here, the theorem holds due to Assumptions A1, A3 and A5.

**Induction hypothesis:**  $\exists n \in \mathbb{N} \forall$  agent sets  $\mathcal{A}', |\mathcal{A}'| \leq n$  : the theorem holds regardless of

- (a) the start and destination states of the participating agents and
- (b) for any prior auction history (of  $i' \in \mathbb{N}_0$  coordination iterations) on that graph that is consistent with the rules of the mechanism and

(c) for any spatio-temporal graph  $\mathcal{G}' = (V', E')$  for which our Assumptions A1-A5 hold and such that we have  $\forall t \geq \min_{a \in \mathcal{A}} ilh^a(i') : E'_t$  is consistent with Assumptions A6-A8.

**Induction step** ( $n \rightarrow n + 1$ ): Let  $i$  be any current coordination iteration where Robot  $A := n + 1 \in \mathcal{A}$  is the highest priority robot in the system potentially competing for resources. Lemma A.3.4 implies the existence of  $T^A \in \mathbb{N}_0$  such that  $\forall \iota \geq i \exists t \leq T^A : Z_t^A(\iota) = D(A)$ . By Assumption A4, Robot  $A$  hence discontinues claiming any vertices

in current spatio-temporal graph  $\mathcal{G}$  for time steps  $\geq T^A$  and thus, is effectively removed from the system beyond time horizon  $T^A$ . Notice, that Lemma A.3.4 also tells us that even if collisions in later time steps among the remaining robots cause collisions in time steps  $\leq T^A$ ,  $A$  is guaranteed to reach his destination before  $T^A$  and will be absent from claiming resources in later time steps.

Furthermore, the number of collisions (and hence, auctions) that can occur with Robot  $A$  are finite since (a) the length of his path is not exceeding  $T^A$  spatio-temporal vertices each corresponding to a representant chosen from a finite number of resources (b) the number of possible paths of length  $T^A$  is finite. Therefore, the other robots can collide with  $A$  only at a finite number of resources. Thus, there exists an iteration  $\hat{i}$  from when on no more collisions occur with  $A$ .

Let  $F = \{a \in \mathcal{A} | ilh^a(\iota) \leq T^A, \forall \iota \geq \hat{i}\}$  be the set of robots that reach their goals along collision free paths until at most time  $T^A$ . If  $F = \mathcal{A}$  we are done since this implies that all robots manage to find collision-free paths in a finite number of iterations.

Otherwise, we define  $\mathcal{A}' = \mathcal{A} - F$  to be the set of robots that require traveling durations to their destinations longer than  $T^A$ .

We choose  $i' \geq \hat{i}$  such that  $\min_{a \in \mathcal{A}'} ilh^a(i') > T^A$ . Note, such an iteration exists since by construction, the individual loss horizons are monotonically increasing. Furthermore, notice that  $A \in F$  and hence,  $|\mathcal{A}'| < |\mathcal{A}| = n + 1$ .

Coordination of the remaining robots in  $\mathcal{A}'$  reduces to a graph planning problem in graph  $\mathcal{G}'$  which we construct from spatio-temporal graph  $\mathcal{G}$  by pruning  $z_0^A(i'), \dots, z_{T^A}^A(i')$  from the latter.

Note,  $\mathcal{G}'$  is consistent with our assumptions A1-A5. Also, since only nodes before  $T^A$  were pruned and  $\min_{a \in \mathcal{A}'} ilh^a(i') > T^A$ , A6- A8 are met for all time steps  $t \geq \min_{a \in \mathcal{A}'} ilh^a(i')$  as well as required by the induction hypothesis and Lemma A.3.4. Since  $|\mathcal{A}'| < n + 1$ , the induction hypothesis entails that this remaining coordination problem is solved in a finite number of coordination iterations. **q.e.d.**

## Bibliography

- [1] M. Bennewitz, W. Burgard, and S. Thrun. Exploiting constraints during prioritized path planning for teams of mobile robots. In *IROIS*, 2001.
- [2] M. Bennewitz, W. Burgard, and S. Thrun. Finding and optimizing solvable priority schemes for decoupled path planning techniques for teams of mobile robots. *Robotics and Autonomous Systems*, 2002.
- [3] C. Bererton, G. Gordon, S. Thrun, and P. Khosla. Auction mechanism design for



- multi-robot coordination. In *NIPS*, 2003.
- [4] C. A. Bererton. *Multi-Robot Coordination and Competition Using Mixed Integer and Linear Programs*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 2006.
- [5] L. Blackmore, M. Ono, A. Bektassov, and B.C. Williams. A probabilistic particle approach to optimal, robust predictive control. *IEEE Trans. on Robotics*, 2010.
- [6] J. Calliess and G. J. Gordon. No-regret learning and a mechanism for distributed multiagent planning. In *Proc. of AAMAS*, 2008.
- [7] J. Calliess, D. Lyons, and U. Hanebeck. Lazy auctions for multi-robot collision avoidance and motion control under uncertainty. Technical Report PARG-01-11, Dept. of Engineering Science, University of Oxford, 2011.
- [8] G. B. Dantzig and P. Wolfe. Decomposition principle for linear programs. *Oper. Res.*, 8:101–111, 1960.
- [9] R.K. Dash, N.R. Jennings, and D. C. Parkes. Computational mechanism design: A call to arms. *IEEE Int. Syst.*, 2003.
- [10] S. de Vries and R. Vohra. Combinatorial auctions: A survey. *INFORMS J. Computing*, 2003.
- [11] M. B. Dias, R. M. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: a survey and analysis. *Proceedings of the IEEE*, 2006.
- [12] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1959.
- [13] M. A. Erdmann and T. Lozano-Perez. On multiple moving objects. *Algorithmica*, 1987.
- [14] A. Farinelli, A. Rogers, and N. Jennings. Decentralised coordination of low-power embedded devices using the max-sum algorithm. In *AAMAS*, 2008.
- [15] A. Farinelli, A. Rogers, and N. Jennings. Coordination using the max-sum algorithm. In *IJCAI-09 Workshop on Distributed Constraint Reasoning (DCR)*, 2009.
- [16] B. Gerkey and M. Mataric. Sold!: Auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation*, 19(5):758–768, 2002.
- [17] C. Guestrin and G. Gordon. Distributed planning in hierarchical factored mdps. In *UAI*, 2002.
- [18] Sven Koenig, Pinar Keskinocak, and Craig A. Tovey. Progress on agent coordination with cooperative auctions. In *AAAI*, 2010.
- [19] M. Lagoudakis, V. Markakis, D. Kempe, P. Keskinocak, S. Koenig, A. Kleywegt, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Int. Conf. on Robotics: Science and Systems*, 2005.

- [20] T. Nishi, M. Ando, and Masami Konishi. Distributed route planning for multiple robots using an augmented lagrangian decomposition and coordination technique. *IEEE Trans. on Robotics*, 2005.
- [21] T. Nishi, M. Ando, and Masami Konishi. Experimental studies on a local rescheduling procedure for dynamic routing of autonomous decentralized agv systems. *Robotics and Computer-Integr. Manuf.*, 2006.
- [22] K. Sycara P. Velagapudi and P. Scerri. Decentralized prioritized planning in large multirobot teams. In *IROS'10*, 2010.
- [23] L. Pallottino, V. G. Scordio, E. Frazzoli, and A. Bicchi. Decentralized cooperative policy for conflict resolution in multi-vehicle systems. *IEEE Trans. on Robotics*, 23(6):1170–1183, 2007.
- [24] D. Parsons and J. Canny. A motion planner for multiple mobile robots. In *ICRA*, 1990.
- [25] T. Sandholm. Algorithm for optimal winner determination on combinatorial auctions. *Artif. Int.*, 2002.
- [26] T. Schouwenaars, B. De Moor, E. Feron, and J. How. Mixed integer programming for multi-vehicle path planning. In *European Control Conference*, 2001.
- [27] A. Stentz and M. B. Dias. A free market architecture for coordinating multiple robots. Technical Report CMU-RI-TR-99-42, Carnegie Mellon, 1999.
- [28] P. Whittle. A multivariate generalization of tchebichev's inequality. *Quarterly Journal of Mathem.*, 1958.