Min-Time Coverage in Constricted Environments: Problem Formulations and Complexity Analysis

Spyros Reveliotis and Young-In Kim

Abstract—**This work deals with the task allocation and robot routing problems that arise in a class of networked mobile robot systems. The defining characteristics of these systems are (i) the constricting nature of their operational environment and (ii) the need to coordinate the robot motion in order to maintain full connectivity of a multi-hop wireless communication network connecting the robots with each other and a master controller that supervises the entire operation. These two elements give rise to resource allocation structures and traffic dynamics that transcend the state of art of the corresponding theory, and challenge our current understandings and insights for these dynamics and their effective management. We provide: (i) a systematic introduction of the considered problems and of the elements that differentiate them from similar task allocation and traffic scheduling problems already studied in the literature; (ii) complete analytical characterizations of these problems in the form of mathematical programming formulations; and (iii) a formal analysis of the worst-case computational complexity of these problems and of certain factors that determine this complexity. Furthermore, the presented results define a base for the development of solution approaches to the considered problems able to manage effectively and systematically the identified trade off between the operational efficiency of the derived solutions and their computational cost.**

Index Terms—**Networked mobile robotic systems; multi-robot coordination; combinatorial scheduling; coverage problems**

I. Introduction

The deployment and coordination of multiple mobile robot systems (MMRS) is a topic that has received extensive attention by many research communities. An excellent account of all this activity, its current achievements, but also the remaining challenges, is presented in [1].

In this work, we focus on a particular MMRS class that has been previously studied in [2], [3]. These two works consider a fleet of mobile robots that must visit a number of pre-specified locations. The robots can communicate wirelessly with each other and with a command and control center that manages the entire operation. But the target locations are accessible through a tunnel system that has a dendrite structure and constricts significantly, both, the robot advancement towards the target locations and the aforementioned communication. Hence, the robot motion must be carefully coordinated in a way that guarantees (i) their physical safety in terms of collision avoidance, and (ii) their ability to remain connected to the command and control center. In the considered operations, physical safety is ensured through the imposition of a "zoning" scheme on the links of the underlying tunnel system, that enforces the physical separation of the robots. On the other hand, the connectivity of each robot to the command and control center is ensured through the preservation of a multi-hop wireless communication network enabling the exchange of messages between each robot and the center via the intermediary robots that are located on the corresponding path in the underlying tunnel system. More specifically, the robot advancements through the available zones of the tunnel system must be further coordinated in order to ensure that every robot remains connected to the multi-hop communication network.[1]

MMRS that possess communication capabilities of the type that is described in the previous paragraph, are characterized as *networked* in the corresponding literature. According to [1], the main body of the research on networked MMRS concerns the pertinent and efficient acquisition, dissemination, and processing by the cooperating robots of all the information that is necessary for the effective execution of their mission. Some more specific overviews of these endeavors can be found in [4], [5]. The work of [4] considers the distributed control of a fleet of mobile robots in order to attain tasks like rendezvous, formation control and flocking. Hence, from a control-theoretic standpoint, the problems addressed in [4] are regulation and tracking problems defined in terms of the entire fleet, and the efficacy of the presented control schemes is based on the information exchange taking place between each robot and the other robots that are located in its vicinity. On the other hand, the work of [5] introduces a notion of efficiency and optimization in the control problems that are addressed in it. The fleet of the mobile robots is perceived as a reconfigurable sensing and communication network, with the performance of its sensing and communication functions being dependent on the spatial distribution of the robots over the covered area. The undertaken objective is to optimize the employed configuration with respect to the attained performance, while accounting for further (re-)configuration costs. Furthermore, a significant part of the developments in [5] concerns the characterization and the assessment of the quality of the wireless connectivity that is provided by each location in the covered area.

When it comes to the coverage tasks addressed in [2], [3], (re-)configurability issues for the maintained multi-hop communication network, like those addressed in [5], are predetermined by the topology of the underlying tunnel system and the imposed zoning scheme. On the other hand, an additional

S. Reveliotis and Y.-I. Kim are with the School of Industrial & Systems Engineering, Georgia Institute of Technology. Emails: spyros@isye.gatech.edu, ykim902@gatech.edu. This work was partially supported by NSF grant ECCS-1707695. The first author also wants to thank Dr. Savvas Loizou for introducing him to the class of OMRP^{2}G problems considered in this work.

[1] A detailed, formal description of the considered MMRS and the operational problems that are addressed in this work are provided in Section II.

important issue in the context of these operations, is the expediency of the execution of the corresponding coverage tasks. This issue can be effectively addressed through 1) a careful allocation of the various subtasks to the available robots, and 2) the coordination of the resulting traffic in a way that ensures a fast execution of these subtasks and of the overall mission of the entire robot team. In [2], [3], these two issues are addressed through some simple heuristics that complement, almost as "after-thoughts", the technical developments that are the primary focus of those works.[2] Yet, a more careful consideration of these issues gives rise to some hard but very interesting combinatorial optimization problems, involving resource allocation structures and traffic dynamics that transcend the state of art of the corresponding theory and our current understanding of this class of problems. This work seeks to provide a formal characterization of these problems, and a first analytical investigation of the involved dynamics and the underlying complexity.

In order to position the considered problems and the intended contribution of this work in the context of the existing literature on MMRS, next we overview the main developments of this literature regarding the problems of task allocation and robot routing, focusing primarily on those results that address time-based performance objectives and have tried to provide a rigorous analytical base for these problems.

One of the first works that has tried to provide a theoretical framework for the task allocation problems arising in MMRS is that of [8]. More specifically, the work of [8] (i) defined a taxonomy for these task allocation problems, (ii) mapped the various elements of this taxonomy to some more abstract problems that are studied by Operations Research (OR), and (iii) used these mappings in order to (a) draw some conclusions about the worst-case computational complexity of the task allocation problems identified by the developed taxonomy, and (b) propose some pertinent solution approaches and policies. The particular features utilized for the problem classification in the taxonomy of [8] are: 1) the ability of a robot to get simultaneously involved in more than one tasks; 2) the number of robots that must be assigned to a task (one vs. many); and 3) the type of the optimization that is pursued by the adopted criterion (myopic vs. longer-term). It was also observed that (i) while the existing OR theory suggested optimizing strategies for some of the identified problem classes, this was not possible for some of the more complex problem classes of the taxonomy, and (ii) there are additional task allocation problems which arise in MMRS and are not even captured by the proposed taxonomy of [8].

Since the publication of [8], an extensive amount of work has strengthened the theoretical base for the task allocation and robot routing problems that arise in the context of the MMRS applications, and has translated the more theoretical findings to effective and efficient operational policies. These endeavors have further pursued and exploited the connections of the considered task allocation and robot routing problems with a host of problems that have been investigated by the disciplines of Operations Research, Theoretical Computer Science, Control Theory, and Artificial Intelligence. At the same time, the identified connections have motivated and defined further research in all of the aforementioned disciplines.

At a first level, many of the task allocation and routing problems arising in MMRS can be referred to the Vehicle Routing (VR) problem [9], that has been studied extensively by the OR community. The basic VR problem consists of two primary subproblems that must be addressed simultaneously: (i) a set of standing service requests, dispersed over a geographical area, must be assigned to a set of vehicles that are available at a "depot"' location, and (ii) every vehicle must be routed through the locations of its assigned requests in an optimal manner. Some typical objectives pursued in the resolution of these two issues are: 1) The minimization of the completion time of the entire plan; this completion time is known as the plan "makespan", and its minimization is equivalent to maximizing a notion of "throughput". 2) The minimization of the mean waiting time for the serviced requests. 3) The minimization of the total distance travelled by all vehicles; this measure constitutes a surrogate for the energy expended during the considered operation. Strongly related to the VR problem is the Traveling Salesman Problem (TSP) [10], that concerns the optimized routing of the system vehicles, and places the VR problem in the class of NP-Hard problems [10].

Of particular interest in the MMRS applications is the Dynamic Vehicle Routing (DVR) problem [11], where the tasks and the corresponding locations that must be serviced by the system robots are not known a priori, but they appear dynamically during the system operation. One way to address this dynamic problem version is by applying to it the theory of on-line algorithms that has been developed by competitive analysis [12], [13]. The design and the evaluation of these algorithms is based on a "worst-case" analysis regarding the arrival process of the service requests. An alternative approach to the DVR problem and its manifestation in the MMRS is presented in [14]. The work of [14] assumes the knowledge of the stochastic processes that generate the service requests and their service times, and develops a queueing-theoretic framework that enables the performance evaluation of a set of policies under some limiting regimes of the offered load. The developed theory also supports the accommodation of priority schemes and of timing constraints for the service requests, and the decentralization of the employed control logic.

Decentralized approaches for the DVR problems that arise in the MMRS operations have also been based on auction mechanisms where the robots compete for the service of the emerging service requests through a bidding process. The employed bidding mechanisms are based on concepts and insights that underlie the modeling and the analysis of the basic VR problem. A theoretical characterization of these bidding mechanisms and a worst-case analysis of their relative performance to the performance of the optimal policy is provided in [15].

An important limitation of all the developments that were discussed in the previous paragraphs, is that they do not account for the presence of (i) obstacles in the operational

[2]Some further heuristic approaches to the task allocation and robot routing problems considered in this work, developed in the context of the broader communication-aware robotic problems that are discussed in [5], can be found in [6], [7].

environment, and (ii) the potential conflicts that these obstacles may imply for the determined routes. These problems become especially prominent in MMRS that possess large robot fleets or operate in highly constricted environments due to the presence of extensive cluttering, narrow aisles, etc. For familiar and structured operational environments, like those considered in this work, a robust approach to cope with these complications is by imposing a *zoning scheme* that partitions the entire operational environment in a number of zones and sets a limit to the number of robots that can be simultaneously present in any zone [16], [17], [18]. In this way, the buffering capacity of every zone becomes a resource that must be negotiated by the traveling robots and the system controller according to a zone allocation protocol [19], [20], [21].

A formal representation of the zoning scheme that was described in the previous paragraph, is a graph where the nodes represent the zones and the edges represent the neighboring structure among the zones. Each node of this graph is assigned a numerical value that defines the buffering capacity of the corresponding zone. Furthermore, the nodes and/or the edges of this graph can be assigned additional weights representing traversal times, distances, or other elements that define a metric structure of interest. Then, a typical problem investigated in the resulting representation is the transport of a set of robots from their original locations to a set of destinations in a way that (i) observes the zone buffering capacity and the corresponding zone allocation protocol, and (ii) optimizes an objective function that is defined by the robot routes and the weights of the nodes and edges belonging on these routes. In the literature, this problem has been characterized as the problem of Optimal Multi-Robot Path Planning on Graphs [22], and it will be referred to as the OMRP^{2}G problem in the following.

The complexity, and even the feasibility, of the OMRP^{2}G problem depends strongly on a number of factors pertaining to (i) the structure of the imposed zoning scheme, (ii) the zone allocation protocol, (iii) the navigational capabilities of the robots within the zones, and (iv) the substitutability of the robots in the pursued tasks.

A first formal analysis regarding the feasibility of the OMRP^{2}G problem appeared in [23], which used permutation group theory in order to address the (n^2-1)-puzzle problem, i.e., the problem of rearranging n^2-1 labeled tokens located on an $n \times n$ grid into a new arrangement through a sequence of "sliding moves" that use the single free node of the grid. It was shown that if the supporting grid is non-bipartite, the original token configuration can be re-arranged to any target configuration. On the other hand, if the supporting grid is bipartite, the set of the possible token distributions on this grid is partitioned into two equivalence classes in terms of the considered reachability requirement. The results of [23] were subsequently extended in works like those of [24], [25], which considered additional graph topologies and more relaxed conditions for the occupancy of the graph nodes by the circulating tokens. These works have also provided polynomial-time algorithms for testing the feasibility of the considered problem.

More recently, works originating from the robotics community, like those presented in [22], [26], [27], [28], [29], [30], have addressed the OMRP^{2}G problem under robot-coordinating schemes involving the simultaneous advancement of entire robot groups that occupy neighboring zones, and potential payload transfers among groups of robots. These features enhance the feasibility of the resulting versions of the OMRP^{2}G problem, and enable a more focused investigation of the optimization part. The vast majority of these optimization problems are shown to be NP-hard in the aforementioned works. However, in [26] it is shown that the OMRP^{2}G problem can admit a polynomial-time solution if the robots are indistinguishable and they can be matched with the targeted destinations in an arbitrary manner by the solution algorithm.

There is also a large amount of work that seeks to provide pertinent heuristic approaches for various instantiations of the OMRP^{2}G problem. Many of these approaches are based on local search schemes that seek to effect improvements to an incumbent routing schedule by identifying and rerouting robots that determine most drastically the cost of this schedule. We refer to [30] for a concrete example of such a heuristic scheme, and for a more systematic survey of the corresponding literature. Furthermore, in [28], [29] it is shown that the current commercial solvers can provide optimal solutions to the mathematical programming (MP) formulations of some large instantiations of the OMRP^{2}G problem versions that are studied in those works, in reasonable times.

Another factor that can complicate substantially the solution of the OMRP^{2}G problem, is a potential inability of the robots to reverse the direction of their motion within their current zone; the corresponding MMRS have been characterized as irreversible in the corresponding literature. Irreversible MMRS are susceptible to deadlock and livelock, which necessitates the deployment of additional control logic for ensuring the liveness of the generated traffic, i.e., the preservation of the ability of every robot to access every zone *ad infinitum*. A systematic treatment of the liveness preservation problem for a broad range of MMRS classes is provided in [31], [32].

Also, the works of [30], [33] consider the embedding of the OMRP^{2}G problem discussed in the previous paragraphs in a model predictive control (MPC) scheme that enables the underlying MMRS to address more complex and dynamically generated sequences of service requirements, and to respond to any slippages in the execution of the original schedules and other experienced contingencies through replanning.

Finally, there is a set of works, like those presented in [34], [35], that have tried to add arbitrary constraints of a more logical nature to the OMRP^{2}G problem, using representations and computational tools borrowed from the areas of formal methods [36] and Discrete Event System (DES) theory [37]. The corresponding theory is very elegant, but the eventual applicability of the results may be limited by the explosive size of the employed representations and/or the complexity of the MP formulations that must be solved for the derivation of the sought schedules.

In the rest of this document we shall show that the connectivity constraints that must be observed by the networked MMRS of [2], [3], give rise to task allocation and robot routing problems with a very different analytical structure from the task allocation and robot routing problems that have been investigated in the forecited literature. Our intention in this work is to (i) define these problems in a way that pronounces the most salient elements differentiating them from the past literature, (ii) provide analytical formulations for them, and (iii) investigate their worst-case computational complexity and the factors that determine this complexity. This analysis also suggests a research plan for developing pertinent solution methods for these problems. At the same time, the presented results constitute novel applications and extensions of the methodological tools of complexity theory and the combinatorial optimization and scheduling theory that are used in this work.

The rest of the paper is organized as follows: Section II provides a systematic introduction of the networked MMRS considered in this work, and the coverage problems that we address in the context of these systems. Section III provides an analytical characterization of these coverage problems in the form of MP formulations. Section IV presents the results regarding the computational complexity of the considered problems. Finally, Section V concludes the paper by discussing the implications of the presented developments for the further study of the OMRP^{2}G problems that are considered in this work, and it also suggests some variations and extensions of these OMRP^{2}G problems.

II. The considered MMRS

In this section we provide a systematic description of the MMRS considered in this work. This MMRS is a formal abstraction of the MMRS described in [3]. A brief description of the structure of that MMRS and its operation is as follows:

A set of underground locations must be inspected by a fleet of mobile robots. These locations are connected to the point where the robots are initially located through a network of tunnels that constitutes a tree. The initial location of the robots defines the root of this tree and the targeted locations are its leaves. Furthermore, the tunnels are narrow and the robots have limited sensing and maneuvering capability. Therefore, for safety reasons, the robots must be separated through the imposition of a zoning scheme with unit buffering capacity for each zone.

The robots possess wireless communication capability, but their communication range is drastically limited by their operational environment. Since these wireless communication links are the only way for each robot to communicate with its operational environment, the robot motion must be coordinated in a way that, at any time point, the active links among the robots define a multi-hop communication network connecting each robot to each other and to a command and control center that is located at the origin. A natural way to ensure this connectivity is by defining the imposed zones in a way that neighboring zones ensure the required connectivity between the robots that occupy these zones, and further stipulating that, at any time point, a zone cannot be occupied unless its parent zone in the underlying tree is also occupied.

Finally, following standard practice in the formal study of the traffic dynamics that are generated by zoning schemes similar to those considered in this work, we further assume that zones are defined in a way that they have uniform traversal time. Then, picking this traversal time as the time unit, we can study the resulting traffic dynamics in discrete time.[3]

In view of the above description, the considered MMRS can be formally represented by a tuple $\mathcal{M} = \langle \mathcal{R}, \mathcal{T} \rangle$, where $\mathcal{R}$ is the set of the robots and $\mathcal{T}$ is the rooted tree representing the tunnel system. The node set V of $\mathcal{T}$ represents the zones of the tunnel system, and the edge set E represents the neighboring relation among the zones.

The root node $o \in V$ is the initial location of all robots and the point of command and control for the entire system. The set of the leaf nodes of $\mathcal{T}$ is denoted by L. Each zone $v \in L$ must be visited by some robot for inspection purposes. The inspection of a leaf zone can be carried out by the visiting robot in the time interval corresponding to a discrete period.

The set of neighbors of a zone $v \in V$ is denoted by $\mathcal{N}(v)$, and for any zone $v \neq o$, $p(v)$ denotes the parent of v in $\mathcal{T}$. Let $z(r,t)$ denote the zone $v \in V$ occupied by robot r at period t. Then, $z(r,t+1) \in \{z(r,t)\} \cup \mathcal{N}(z(r,t))$; i.e., robot r can either remain in the same zone at period $t+1$, or advance to a neighboring zone $v' \in \mathcal{N}(v)$. Furthermore, at any period t, zones $v \neq o$ cannot contain more than one robot. On the other hand, a group of robots can coordinate their advancement over a path of neighboring zones; i.e., for a group of robots $r_1, r_2, \ldots, r_n$ with $z(r_i,t) \in \mathcal{N}(z(r_{i-1},t))$, for $i = 2,\ldots,n$, we allow $z(r_i,t+1) = z(r_{i+1},t)$, $i = 1,\ldots,n-1$, provided that robot r_n moves itself to a free zone or to the root zone o at period $t+1$.

Robots can reverse the direction of their motion within their zone. This assumption is reasonable in the context of the considered applications, and furthermore, it is necessary due to the tree structure of the underlying tunnel system.

Finally, as observed in the opening part of this section, the communication connectivity among the robots and the system controller is established by stipulating that, for every zone $v \neq o$ and every period t,

$$\exists\ r \in \mathcal{R} \text{ with } z(r,t) = v \implies \exists\ r' \in \mathcal{R} \text{ with } z(r',t) = p(v) \tag{1}$$

The above requirement implies that for every zone v occupied by a robot in period t, all the zones in the path connecting zone v to the root zone o in tree $\mathcal{T}$ are also occupied by a robot in period t. Furthermore, the root zone o is always occupied by at least one robot.

We want to determine a plan that will advance the robots $r \in \mathcal{R}$ in a way that is consistent with the above assumptions regarding the robot capabilities and the zone allocation protocol, and at the end of its execution, each leaf zone $v \in L$

[3]Nonuniform traversal times for the system zones can be easily introduced into our model. But this feature would overload the employed notation and the pursued analysis without adding anything substantial to this analysis.

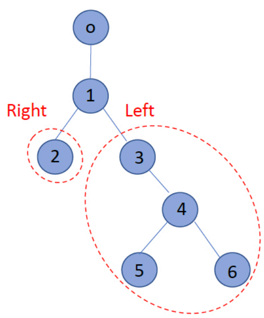

First / Problem	Left	Right
M	5	6
TVT	14	13

Fig. 1: An MMRS where the M-problem and the TVT-problem have different optimal plans. The tree $\mathcal{T}$ for this MMRS is depicted in the left part of the figure, and it is further assumed that $|\mathcal{R}| \geq 6$. The table in the figure provides the objective values for the resulting M- and TVT-problems under different priorities for the two subtrees that emanate from node 1. We can see that the optimal plan for the M-problem gives priority to the left subtree, while the optimal plan for the TVT-problem gives priority to the right subtree.

will have been visited by some robot.[4] Let $\mathcal{P}$ denote the set of feasible plans, and for every $P \in \mathcal{P}$ and $v \in L$, let $C(v; P)$ denote the first period that plan P places a robot in zone v. We are especially interested in plans P^* such that

$$P^* = \arg\min_{P \in \mathcal{P}} \max_{v \in L} C(v; P) \tag{2}$$

or

$$P^* = \arg\min_{P \in \mathcal{P}} \sum_{v \in L} C(v; P) \tag{3}$$

Each of Eqs 2 and 3 defines a combinatorial optimization – or scheduling – problem. In the following, we shall refer to the scheduling problem defined by Eq. 2 as the *Makespan-minimization* problem, or the *M-problem*. The scheduling problem defined by Eq. 3 will be characterized as the *Total Visitation Time-minimization* problem, or the *TVT-problem*. Figure 1 shows that these two problems are in a Pareto optimal relationship [38], [22]; i.e., there are MMRS where the sets of optimal plans for these two problems have no common element. Hence, these two problems require separate treatments.

In the rest of this paper we provide complete formulations for the combinatorial optimization problems that are defined by Eqs 2 and 3, and we analyze the computational complexity of these problems. However, before concluding this section, we provide some further justification for the tree topology presumed for the tunnel system of the considered MMRS. For this, first we notice that, according to [2], [3], sewage networks and other similar public infrastructure facilities, and also the tunnel systems deployed in underground operations like mines and archeological excavation sites, *do* possess a tree structure. But trees are also prominent in the studies pertaining to the connectivity analysis of communication networks and its enforcement, through the notion of "spanning trees" [39], [1], [4], [6]. Hence, the restriction of this first investigation of the considered MMRS to guidepath networks with a tree topology is in line with the realities of the targeted applications and it does not compromise the applicability of the results.

III. MP formulation of the considered problems

In this section we provide analytical characterizations for the M- and TVT-problems in the form of MP formulations. Besides providing succinct, unambiguous characterizations for the corresponding problems, the presented formulations also have practical value, since, as remarked in the introductory section, the works of [28], [29] have demonstrated the capability of modern solvers to cope effectively with the MP formulations of sizable instances of other OMRP^{2}G problems.

In the subsequent discussion, we employ the notation introduced in Section II. We also let $\bar{T}$ denote an upper bound for the completion time of an optimal plan P^* for each problem. One way to obtain such an upper bound is by considering the completion time of the plan P that tries to reach one leaf zone at a time, while scanning the tree $\mathcal{T}$ in a depth-first sense.

The decision variables employed by this formulation consist of: (i) a set of "state" variables that trace the distribution of the robots to the system zones over time; (ii) a second set of "control" variables that determine the evolution of the robot distribution to the system zones; and (iii) some auxiliary variables that facilitate the testing and/or the enforcement of certain conditions on the system behavior, and the formulation of the objective function. The complete list of the employed decision variables is as follows:

- State variables
  - $x_{v,t}$, $v \in V$, $t \in \{0, 1, \ldots, \bar{T}\}$: a nonnegative integer variable indicating the number of robots in zone v at period t.
- Control variables
  - $u_{v,v',t}$, $v \in V$, $v' \in \mathcal{N}(v)$, $t \in \{1, \ldots, \bar{T}\}$: a nonnegative integer variable representing the number of robots moving from zone v to neighboring zone v' at period t.
- Auxiliary variables
  - $y_{v,t}$, $v \in L$, $t \in \{1, \ldots, \bar{T}\}$: a binary variable for testing whether leaf zone v has been visited by period t.[5]
  - s_t, $t \in \{1, \ldots, \bar{T}\}$: a binary variable for testing whether the entire leaf-node visitation task has been completed by period t.

The technological constraints employed by the MP formulations of the M- and the TVT-problems are as follows:

[4] In more technical terms, a plan is a sequence of distributions, $\mathcal{D}_t, t = 0, 1 \ldots$, of the system robots to the various zones of the underlying tunnel system. The distribution $\mathcal{D}_{t+1}$, for period $t+1$, is obtained from the distribution $\mathcal{D}_t$ by relocating a number of robots from their zones at period t to some neighboring zone, while abiding to the introduced assumptions about the maneuvering capabilities of the robots and the zone allocation protocol. This characterization is specified further through the MP formulations of Section III.

[5] Actually, the pricing of the variables $y_{v,t}$ is part of the "informational state" that drives the underlying decision making process at period t. We have included these variables in the set of the auxiliary variables since their values are determined by the values of the corresponding variable sets $\{x_{v,q}, q = 1, \ldots, t\}$.

$$x_{o,0} = |\mathcal{R}| \tag{4}$$

$$\forall v \in V \setminus \{o\}, \quad x_{v,0} = 0 \tag{5}$$

$$\begin{gathered} \forall v \in V, \ \forall t \in \{1, \ldots, \bar{T}\}, \\ x_{v,t} = x_{v,t-1} + \sum_{v' \in \mathcal{N}(v)} \big(u_{v',v,t} - u_{v,v',t}\big) \end{gathered} \tag{6}$$

$$\forall v \in V, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad \sum_{v' \in \mathcal{N}(v)} u_{v,v',t} \leq x_{v,t-1} \tag{7}$$

$$\forall v \in V \setminus \{o\}, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad x_{v,t} \leq 1 \tag{8}$$

$$\forall v \in V \setminus \{o\}, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad x_{v,t} \leq x_{p(v),t} \tag{9}$$

$$\forall v \in L, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad y_{v,t} \leq \sum_{q \in \{1, \ldots, t\}} x_{v,q} \tag{10}$$

$$\forall v \in L, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad s_t \leq y_{v,t} \tag{11}$$

Constraints 4 and 5 define the initial distribution of the robots by means of the state variables $x_{v,0}, \ v \in V$. Constraint 6 expresses the evolution of the robot distribution to the system zones at period t based on the control decisions that are expressed by the variables $u_{v,v',t}$. Constraint 7 stipulates that the control decisions at period t must be feasible with respect to the robot distribution over the system zones at period $t-1$. Constraint 8 enforces the buffering capacity of the zones $v \neq o$. Constraint 9 enforces the condition of Eq. 1. Constraint 10 forces the binary variable $y_{v,t}$ to zero if leaf zone v has not been visited by period t. Finally, Constraint 11 forces the binary variable s_t to zero if there is a leaf zone v that has not been visited by period t.

The M-problem can be expressed by the following formulation:

$$\max \sum_{t \in \{1, \ldots, \bar{T}\}} s_t \tag{12}$$

s.t. Constraints 4 – 11 plus the sign restrictions for the problem variables specified during the introduction of these variables.

The TVT-problem can be expressed by the following formulation:

$$\max \sum_{v \in L} \sum_{t \in \{1, \ldots, \bar{T}\}} y_{v,t} \tag{13}$$

s.t. Constraints 4 – 10 plus the sign restrictions for the problem variables specified during the introduction of these variables.

An optimal solution of each of these two formulations determines an optimal plan P^* for the corresponding scheduling problem through the quantities $[u_{v,v',t} - u_{v',v,t}]^+$ for every pair (v, v') of neighboring zones and period t.[6] Also, these formulations can be easily adjusted to enforce an arbitrary buffering capacity $C_v \geq 1$ for each zone $v \in V \setminus \{o\}$; we leave the relevant details to the reader.

Finally, an additional very important remark from a computational standpoint, is that we can replace the original sign restrictions of the variables $x_{v,t}$, $y_{v,t}$ and s_t with the following constraints that relax the integrality requirements for these variables:

$$\forall v \in V, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad x_{v,t} \geq 0 \tag{14}$$

$$\forall v \in L, \ \forall t \in \{1, \ldots, \bar{T}\}, \quad 0 \leq y_{v,t} \leq 1 \tag{15}$$

$$\forall t \in \{1, \ldots, \bar{T}\}, \quad 0 \leq s_t \leq 1 \tag{16}$$

Indeed, it can be easily checked that, as long as we retain the integrality requirement for the variables $u_{v,v',t}$, Constraints 4–6 will ensure the integrality of the variables $x_{v,t}$, and this fact subsequently preserves the mechanism that establishes the correct pricing of the variables $y_{v,t}$ and s_t in any optimal solution of the resulting formulation.

IV. Complexity Analysis

This section establishes that the M- and TVT-problems are NP-hard [10]. However, the last part of the section also shows that for specific structures of $\mathcal{T}$, these problems may admit polynomial solutions that even take the convenient form of priority rules similar to those used in other applications of combinatorial scheduling theory [40]. Besides defining a boundary between hard and easy cases for the considered class of OMRP^{2}G problems, results of the last category are useful for developing pertinent suboptimal solutions for the harder cases.

We start with establishing the NP-hardness of the M-problem.

Theorem 1: The M-problem defined in Section II is strongly NP-hard.

Proof: We prove the result of Theorem 1 in three stages. First, we consider a modified version of the M-problem that does not require the presence of a robot at the root zone o, and we show that the *decision version* of the modified M-problem can provide a representation of the Bin Packing problem [41]. Next, we use this result in order to obtain a strong NP-hardness proof for the modified M-problem from the NP-hardness proof of [41] for the Bin Packing problem. Finally, we adapt the NP-hardness proof for the modified M-problem to the original M-problem that stipulates the presence of at least one robot at zone o.

We start by specifying the decision version of the considered M-problem. This version does not seek a plan P^* that is optimal according to the criterion of Eq. 2, but for any given instance $\mathcal{M} = \langle \mathcal{R}, \mathcal{T} \rangle$ and some constant K, it asks whether there exists a plan $P \in \mathcal{P}$ with $\max_{v \in L} C(v, P) \leq K$.

Also, the Bin Packing problem can be stated as follows [41]: Given a set $\mathcal{A} = \{a_1, a_2, \ldots, a_N\}$ of N positive integers (the items), and two more integers, C (the capacity) and B (the number of bins), determine whether set $\mathcal{A}$ can be partitioned into B subsets, each of which has a total sum of at most C.

In the semantics of the modified M-problem, the given set of items $\{a_1, a_2, \ldots, a_N\}$ is represented by the tree that is

[6]We remind the reader that $[x]^+ = \max\{x, 0\}$.

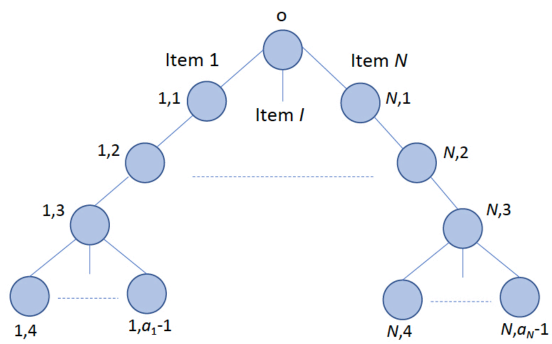

Fig. 2: The tree representing the N items of the Bin Packing problem used in the proof of Theorem 1.

TABLE I: The item sizes employed in the proof of [41].

Item	Size
Marked item cor. to b_i	$10M^4 + iM + 1$
Other items cor. to b_i	$11M^4 + iM + 1$
Marked cor. to g_j	$10M^4 + jM^2 + 2$
Other items cor. to g_j	$11M^4 + jM^2 + 2$
Marked cor. to h_k	$10M^4 + kM^3 + 4$
Other items cor. to h_k	$8M^4 + kM^3 + 4$
Triplet $(b_i, g_j, h_k) \in T$	$10M^4 + 8 - iM - jM^2 - kM^3$

depicted in Figure 2. The leaf nodes in the subtree representing item I can be visited by a team of four robots in a_I periods in a way that, at period a_I, one robot will have returned at the root zone o and the other three robots will be located at zones $v_{I,1}$, $v_{I,2}$ and $v_{I,3}$. Furthermore, the arrangement of the robots at period a_I enables them to start the visitation of the subtree corresponding to another item J at the next period. Hence, considering a team of four robots as a bin, and assuming that this bin is allocated the item set $\mathcal{I} \subseteq \mathcal{A}$, we can see that the corresponding robot team is able to visit all the leaf nodes of the subtrees representing these items in $\left(\sum_{I \in \mathcal{I}} a_I\right) - 1$ periods.[7]

Next we use the correspondences that were established in the previous paragraph, in order to adapt the strong NP-hardness proof for the Bin Packing problem that is provided in [41], to a strong NP-hardness proof for the decision version of the modified M-problem. The proof of [41] is based on a polynomial reduction of the 3D-Matching problem (also known as the Tripartite Matching problem) to the Bin Packing problem. Next, we introduce the 3D-Matching problem, and we overview the main points of the proof of [41].

In the 3D-Matching problem, we are given three nonintersecting sets $B = \{b_1, b_2, \ldots, b_n\}$, $G = \{g_1, g_2, \ldots, g_n\}$ and $H = \{h_1, h_2, \ldots, h_n\}$, of common cardinality n, and also a set of triplets $T = \{t_1, t_2, \ldots, t_m\} \subseteq B \times G \times H$, and we are asked whether there is a subset of n triplets in T such that each element from each set B, G and H is contained in one of the n triplets.

[7]For the subtree to be processed last, there is no need to return a robot to the root node. We also notice that the item representation of Figure 2 cannot represent items I with $a_I < 5$. But this representational capability is not necessary in the context of this proof.

For a given instance of the 3D-Matching problem, the proof of [41] constructs an instance of the Bin Packing problem that contains $N = 4m$ items: one for each triplet and one for each occurrence of the elements b_i, g_i and h_i, $i = 1, \ldots, n$, in T. Furthermore, among the set of items corresponding to the occurrences of b_i (resp., g_i or h_i) one is singled out as "marked".

The N items are assigned the sizes indicated in Table I. The parameter M that appears in Table I is a sufficiently large number that is polynomially related to the dimensions n, m of the 3D-Matching problem instance.[8] Hence, the unitary representation of the item sizes quoted in Table I through the tree of Figure 2 results in a tree size that is polynomially related to n, m. The specification of the induced Bin Packing problem is completed by setting $C = 40M^4 + 15$ and $B = m$.

In [41] it is shown that there exists a solution for the induced Bin Packing problem if and only if there is a 3D-matching for the original 3D-Matching problem. In that case, the solution of the Bin Packing problem will employ m full bins, with each bin containing exactly four items. These four items will be (i) an item corresponding to a triplet (b_i, g_j, h_k) of T, (ii) an item corresponding to one of the occurrences of b_i, (iii) an item corresponding to one of the occurrences of g_j, and (iv) an item corresponding to one of the occurrences of h_k. Furthermore, the last three items in this list will either be all marked or none of them will be marked. Finally, the triplets containing the marked items define a solution for the 3D-Matching problem.

In view of the correspondence between the Bin Packing problem and the modified M-problem that was established in the earlier parts of this proof, the developments of [41] can provide a polynomial reduction of the 3D-Matching problem the modified M-problem as follows:

For each item defined in the proof of [41], the new reduction employs a subtree with the structure depicted in Figure 2. Also, the reduction employs $4m$ robots working in 4-robot teams for the satisfaction of the visitation requirements that are defined by any item-representing subtree. Taking into consideration Footnote 7, the bound K for the induced instance of the modified M-problem is set equal to $K = C - 1 = 40M^4 + 14$. Then, clearly, a feasible instance of the 3D-Matching problem induces a feasible instance of the modified M-problem. To see that an infeasible instance of the 3D-Matching problem induces an infeasible instance of the modified M-problem, it suffices to show that there is no feasible plan for the modified M-problem instance that visits the leaf nodes of an item-representing subtree in an intermittent manner. But this claim is true because a plan possessing such an intermittent visitation pattern would visit the corresponding nodes $v_{I,1}, v_{I,2}, v_{I,3}$ more than once, and this is not affordable under the time budget defined by K (remember the correspondence between the item sizes and the representation of these sizes by the visitation time of the corresponding subtrees by a robot team, that was discussed in the first part of this proof).

Next we modify the above NP-hardness proof for the modified M-problem in order to establish the strong NP-hardness of the original version of the M-problem that was

[8]Quoting [41], we can think of M as $100n$, for a more concrete example.

defined in Section II. In this case, the M-problem induced by the input 3D-Matching problem is specified similarly to the modified M-problem that was considered in the previous part of the proof, but the new MMRS has $4m+1$ robots. Robots work in teams of four as in the previous case, and an extra robot is kept at the root zone o.

Suppose that a 4-robot team is currently occupied with the visitation of the subtree representing some item I. Then, thanks to the availability of the additional robot at the root zone, the system can initiate the visitation of the subtree representing another item J as soon node v_{I,a_I-1} has been visited. This remark further implies that the total (discrete) time required for the visitation of the subtrees corresponding to the four items included in a single bin in the proof of [41], can be reduced by up to three periods.[9]

Clearly, the potential reductions of the bin processing times that are discussed in the previous paragraph, preserve the previously established fact that the induced M-problem instance is feasible if the original 3D-Matching problem instance is feasible. Next, we argue that the potential reductions of the "bin processing times" are not able to turn an infeasible instance of the modified M-problem into a feasible instance for the original M-problem. Hence, the provided construction still resolves effectively the feasibility of the input 3DM-matching problem instance, and therefore, the original M-problem is strongly NP-hard.

In order to establish the necessary result, first we notice that from the item sizes that are reported in Table I and the employed C and K values, it is clear that any plan for the M-problem instance having a chance to be feasible with respect to employed bound K, cannot assign more than 4 item-representing subtrees to any robot team. And since there are $N = 4m$ such subtrees and m robot teams, it follows that any viable assignment will assign exactly four item-representing subtrees to each team. Hence, even in the case of infeasible 3D-Matching problem instances, meaningful reductions of the total time required for a team assignment cannot be more than 3 periods. This potential time gain cannot turn an infeasible modified M-problem instance into a feasible instance for the original M-problem version, due to the following two reasons: (a) The addition of any extra robots to a four-robot team cannot expedite the processing of an item-representing subtree, due to the "bottleneck" role of the corresponding zone $v_{I,3}$ in this process. (b) Furthermore, it takes four periods to reach a leaf node in an item-representing subtree that is not currently processed. Hence, the early completion of a robot team assignment by 3 periods, cannot increase the processing capacity of another team within the time-horizon that is defined by K. □

The modified M-problem used in the proof of Theorem 1 has a strong conceptual affinity to the classical scheduling problem of minimizing the makespan of the processing of a set of tasks by a set of identical parallel machines. The employed teams of robots play the role of the machines, and the subtrees corresponding to the different items define the tasks to be processed by these machines. This conceptual analogy can be useful when contemplating various properties of the considered MMRS. On the other hand, this analogy cannot support an NP-hardness proof for the M-problem (e.g., through a polynomial reduction from the Partitioning problem [10]), because in the tree of Figure 2, the task durations have a *unitary* representation by the number of nodes in the corresponding subtree.

The proof of Theorem 1 also highlights the fact that the concurrent advancement of the robots towards the leaf nodes of $\mathcal{T}$ is constrained by two different problem elements: (i) the size $|\mathcal{R}|$ of the robot fleet, and (ii) the unit buffering capacity of the various zones, which constrains the flow of the available robots towards the various parts of $\mathcal{T}$. The next result shows that the M-problem remains strongly NP-hard even if we avail of an arbitrarily large fleet of robots.

Theorem 2: The M-problem defined in Section II remains strongly NP-hard even if $|\mathcal{R}| \geq |V|$, i.e., even if we have more robots than the number of zones in the underlying guidepath network.

Proof: The proof of this theorem is based on a polynomial reduction from the 3D-Matching problem that is built upon the reduction used in the proof of Theorem 1. Furthermore, the notation and the terminology that is used in this proof are similar to the notation and the terminology that were introduced in that previous proof.

The tree $\mathcal{T}$ employed in the new reduction is depicted in Figure 3. The subtree of $\mathcal{T}$ that is rooted at node o' and consists of the N subtrees labeled 'Item 1', ..., 'Item N' plays a role similar to the role of the tree that was used in the proof of Theorem 1. This subtree of $\mathcal{T}$ is constructed in exactly the same manner with the corresponding tree of Figure 2, except for the fact that the item-representing subtrees have a depth of 5 instead of 4; this modification is explained in the following.

The new tree $\mathcal{T}$ has $(m-1)$ additional item-representing subtrees that are rooted at node o'. The specific structure and

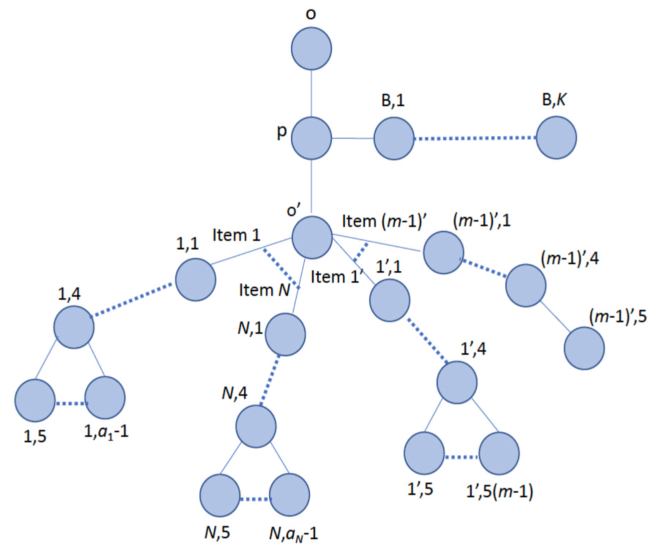

Fig. 3: The tree $\mathcal{T}$ used in the proof of Theorem 2.

[9]The incurred reduction for certain bins might be less than three periods if there is a simultaneous completion of the processing of two or more item-representing subtrees by the robot teams working on them. In this case, only one new subtree will be started by the robot currently available at the origin.

the role of these subtrees are also explained in the following.

Finally, tree $\mathcal{T}$ has an additional part that consists of the path $\langle o, p, o' \rangle$ and the path $\langle (B, 1), \ldots, (B, K) \rangle$ that should be perceived as a subtree rooted at node p. This subtree contains $K = 40M^4 + 14$ nodes.

The number of robots, $|\mathcal{R}|$, is picked as any number that satisfies the condition of Theorem 2.

Similar to the case of the reduction of Theorem 1, we assign 5m+1 robots to work on the visitation of the leaf nodes of the subtrees corresponding to items $1, \ldots, N$. These robots are advanced to the corresponding subtree rooted in node o' in the first part of the employed plan, and they are expected to work as m teams of 5 robots per team for the visitation of the leaf nodes in this subtree, while an extra robot will be located at node o', providing the necessary connectivity between the various teams and the root zone o.

However, in the current case, the first team of robots will start working on its assigned items with a lag of two periods, that corresponds to the traversal of the path $\langle p, o' \rangle$. The second team will experience an additional lag of 5 periods in the initiation of its assignment, that results from the advancement of the first team in the subtree rooted at o'. Similarly, every other robot team will start working on its assignment with a lag of 5 periods compared to the previous robot team. Hence, the last team will start its assignment with a lag of $2+5(m-1)$ periods. For this reason, we set the target makespan equal to $K' = K+5m-3$, where K is the target makespan in the proof of Theorem 1; i.e., eventually we set $K' = 40M^4 + 5m + 11$.

Furthermore, the insertion in the new tree $\mathcal{T}$ of the subtrees labeled 'item $1'$' to 'item $(m-1)'$' intends to occupy the robot teams that start early and therefore will finish their assignment earlier than some other teams. In particular, the completion of the task labeled 'item $1'$' accounts for the lag of $5(m-1)$ periods between the start by the first robotic team that is advanced to the subtree that is rooted at o' and the start by the last robotic team to reach this subtree. Similarly, the task labeled 'item $(m-1)'$' accounts for the lag of 5 periods between the team starting next to last and the team starting last. The remaining $(m-3)$ subtrees are defined in a similar manner. Then, it is clear that no team can interfere with the work of the other teams due to the experienced lags. On the other hand, it is also true that, in the new regime, every team except for the last will process 5 item-representing subtrees, instead of 4 that was the case in the proof of Theorem 1. This fact necessitates the increase of the depth of the item-representing subtrees by one node, in order to avoid possible interference among the robot teams due to presence of the extra robot at node o'; c.f. the relevant discussion in the last part of the proof of Theorem 1.

The above discussion implies that when provided with $5m+1$ robots, the subtree that is rooted at node o' can resolve the feasibility of the input 3D-Matching problem instance in a way similar to the corresponding mechanism that was established in the proof of Theorem 1. The fact that this mechanism will not be jeopardized by the provision of additional robots to this subtree is ensured by the presence of the path $\langle (B, 1), \ldots, (B, K) \rangle$, that is rooted at node p. More specifically, after the $5m+1$ robots have been provided to the subtree rooted at node o', and the last of the m teams has started working on its assignment, the unit buffering capacity of node p must be used during the remaining time until the end of the provided time horizon of K' periods, for the advancement of the necessary robots for the visit of the leaf node (B, K).

Hence, the constructed M-problem instance satisfies the assumptions of Theorem 2, and it is equivalent to the input 3D-Matching problem instance in terms of their feasibility. The proof concludes by further noticing that the size of the constructed M-problem instance is polynomially related to the dimensions of the 3D-Matching problem instance. □

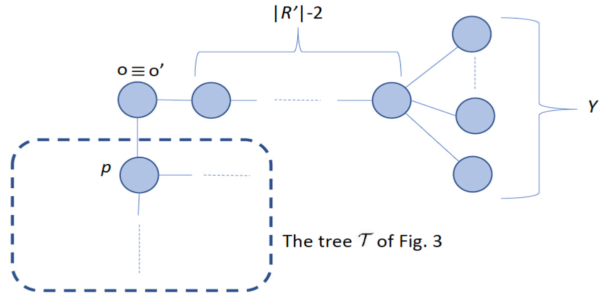

Fig. 4: The tree $\mathcal{T}'$ used in the proof of Theorem 3.

The next theorem establishes the strong NP-hardness of the TVT-problem.

Theorem 3: The TVT-problem defined in Section II is strongly NP-hard.

Proof: We provide a polynomial reduction of the 3D-Matching problem to the TVT-problem, by showing that the M-problem instance constructed in the proof of Theorem 2 can be polynomially reduced to a TVT-problem instance. In the following, we shall respectively denote these two problem instances by $(\mathcal{M}, K)$ and $(\mathcal{M}', K')$.[10] Also, we shall use the notation $\mathcal{M} = \langle \mathcal{R}, \mathcal{T} = (V, E) \rangle$ and $\mathcal{M}' = \langle \mathcal{R}', \mathcal{T}' \rangle$, and we shall denote the set of the leaf nodes of tree $\mathcal{T}$ by L, and the entire set of leaf nodes in $\mathcal{T}'$ by L'.

The tree $\mathcal{T}'$ of the induced TVT-problem instance is depicted in Figure 4. Tree $\mathcal{T}'$ is obtained from the tree $\mathcal{T}$ depicted in Figure 3, by adding to that tree the new subtree that is depicted in Figure 4. It is important to notice that by setting the length of the internal path of the new subtree equal to $|\mathcal{R}'| - 2$, we ensure the accessibility of the leaf nodes of the new subtree, but we also ensure that the visits of the leaf nodes in the original tree $\mathcal{T}$ and in the new subtree cannot occur simultaneously.

Next we discuss the specification of $|\mathcal{R}'|$. The key requirement that drives the specification of this quantity is that an optimized plan for the resulting TVT-problem will always visit the leaf nodes of the original tree $\mathcal{T}$ before visiting any leaf node of the new subtree. This requirement can be satisfied by setting

$$|\mathcal{R}'| \;\geq\; \bar{K}Y + 2 \;\geq\; |\mathcal{R}| \tag{17}$$

[10]The reader should notice that the value of K in the current proof is the value of K' in the proof of Theorem 2.

where $\bar{K}$ is a polynomial upper bound to the makespan of an optimal plan for the input M-problem.[11]

Indeed, under the left inequality of Eq. 17, the visit of a leaf node of the new subtree before the completion of the processing of tree $\mathcal{T}$ will increase the visitation time of at least one leaf node in $\mathcal{T}$ by $|\mathcal{R}'|-1$ periods. On the other hand, the increase in the total visitation time of the nodes of the new subtree that will result by the earlier processing of tree $\mathcal{T}$ is upper-bounded by $\bar{K}Y$.

The right inequality of Eq. 17 ensures that $|\mathcal{R}'| \geq |\mathcal{R}|$, and since $|\mathcal{R}| \geq |V|$ (c.f. Theorem 2), the potential increase of the number of robots in the induced TVT-problem that is implied by Eq. 17, does not alter the optimal makespan for the processing of tree $\mathcal{T}$ (which should take place before the processing of the new subtree, for the reasons that were discussed in the previous paragraph).

The proof of Theorem 2 has established that any competitive plan for the corresponding M-problem instance will keep pumping robots from the root node o to the rest of the tree $\mathcal{T}$ for the entire time interval $\{1,\ldots,K\}$. Hence, when operating under any of these plans for the visitation of the leaf nodes $v \in L$ of the tree $\mathcal{T}'$, at period K, there will be K robots located in the subtree of $\mathcal{T}'$ that is rooted at node p. If it also holds that

$$\bar{K}Y + 2 \;\geq\; 2K+1 \tag{18}$$

then, any competitive plan for the considered TVT problem instance will also have deployed K robots on the path of the $|\mathcal{R}'|$-2 nodes in the new subtree of $\mathcal{T}'$ by period K, and therefore, the number of the additional robots needed to reach any of the leaf nodes of this subtree is $|\mathcal{R}'|-1-K$.

The rest of the proof exploits all the above remarks in order to price the parameters Y and K' of the induced TVT-problem instance in a way that establishes the equivalence of this TVT-problem instance to the input M-problem instance.

Hence, suppose first that the input M-problem instance $(\mathcal{M},K)$ is feasible. Then, there exists a plan P for this problem instance with $\max_{v\in L}\{C(v;P)\} \leq K$. Consider the plan P' for the TVT-problem instance that executes plan P initially, and subsequently it visits the leaf nodes of the new subtree, one node at a time. Then, under the assumptions of Eqs 17 and 18,

$$\begin{aligned}\sum_{v\in L'} C(v;P') &\leq |L|K + \sum_{j=1}^{Y}\big(K + (|\mathcal{R}'|-2-K)+j\big) \\ &= |L|K + Y(|\mathcal{R}'|-2) + \frac{Y(Y+1)}{2}\end{aligned} \tag{19}$$

We set

$$K' = |L|K + (|\mathcal{R}'|-2)Y + \frac{Y(Y+1)}{2} \tag{20}$$

and in the following we shall set Y so that there is no plan P' with $\sum_{v\in L'} C(v;P') \leq K'$ for infeasible M-problem instances.

Consider an input M-problem instance $(\mathcal{M},K)$ that is infeasible, and recall that any optimized plan P' for the induced TVT-problem instance will process the original tree $\mathcal{T}$ before visiting any leaf nodes of the new subtree. For any such plan P', it will hold

$$\begin{aligned}\sum_{v\in L'} C(v;P') &> K + \sum_{j=1}^{Y}\big(K+1+(|\mathcal{R}'|-2-K)+j\big) \\ &= K + Y(|\mathcal{R}'|-1) + \frac{Y(Y+1)}{2}\end{aligned} \tag{21}$$

The first term in the right-hand-side of the first row of Eq. 21 acknowledges the fact that, due to the infeasibility of the input M-problem instance, at least one leaf node $v \in L$ in tree $\mathcal{T}$ will have a visitation time higher than K (and this also justifies the strict inequality in the first part of the equation). The summation appearing in the same part of the equation is a lower bound for the total visitation time of the leaf nodes in the new subtree. This lower bound results from the facts that (i) these nodes will be visited after the completion of the processing of tree $\mathcal{T}$, and (ii) due to the infeasibility of the input M-process instance, the completion of $\mathcal{T}$ will not take place before period $K+1$.

In view of Eqs 19–21, in order to establish the sought equivalence between the input M-problem instance and the induced TVT-problem instance, Y must satisfy the following inequality:

$$\begin{aligned}K + Y(|\mathcal{R}'|-1) \;\geq\; |L|K + (|\mathcal{R}'|-2)Y \;\Longrightarrow \\ Y \;\geq\; K(|L|-1)\end{aligned} \tag{22}$$

Picking Y as specified by Eq. 22 also meets the requirement of Eq. 18. The proof is completed by setting

$$Y \geq \max\big\{K(|L|-1), \big(|\mathcal{R}|-2\big)/\bar{K}\big\} \tag{23}$$

so that it satisfies the right inequality in Eq. 17, as well. □

Currently, we do not have a result similar to that of Theorem 2 that would characterize the worst-case computational complexity of the TVT-problem in the case that we have a very large fleet of robots. But in the rest of this section, we use this particular problem version to demonstrate (i) how the presence of special structure in tree $\mathcal{T}$ can enable the development of efficient solution algorithms for the considered OMRP^{2}G problems, and also (ii) how the relevant analysis can be further facilitated by referring these problems to some classical problems of combinatorial scheduling theory.

The special structure for the tree $\mathcal{T}$ considered in the following discussion is depicted in Figure 5. It consists of a path of length P_0 leading from the root zone o to zone p, and there are k subtrees rooted at node p possessing what we shall call a *"claw"* structure. More specifically, each of these subtrees consists of an internal path of some length $P_i,\ i=1,\ldots,k$, and the last node of this path, q_i, has Q_i additional paths emanating from it and having equal length R_i. We want to come up with a visitation plan P for the leaf zones of tree $\mathcal{T}$ that minimizes the total visitation time for these zones. Furthermore, we assume that $|\mathcal{R}| \geq |V|$, where $|V|$ is the number of zones of the considered tree $\mathcal{T}$.

The assumption $|\mathcal{R}| \geq |V|$ implies that there is no need to reuse the robots directed to any claw for the visitation of the leaf nodes of some other claw. Also, the unit capacity of each zone q_i implies that the robots that must be directed to the i-th

[11] E.g., $\bar{K}$ can be set to $2\cdot|L|\cdot$ (the depth of $\mathcal{T}$).

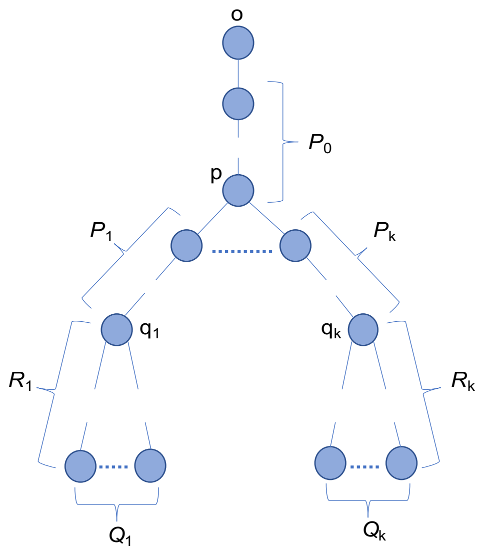

Fig. 5: A tree $\mathcal{T}$ with a "claw" structure.

claw for an expedient visitation of its leaf nodes, is exactly $(P_i + R_i)$. On the other hand, the path from o to p implies that, at each period, we can advance a robot from the origin to only one claw; i.e., this path serializes the service of the claws. Hence, eventually the considered OMRP^{2}G problem reduces to the determination of an optimal sequence for forwarding to each claw the required robots.

In order to determine such an optimal sequence, let $[i]$, $i = 1, \ldots, k$, denote the claw that is in the i-th position of this sequence. Then, it is easy to see that the total visitation time for the leaves of this claw is equal to

$$\sum_{j=1}^{Q_{[i]}} \left\{ P_0 + \sum_{l=1}^{i-1} \left(P_{[l]} + R_{[l]} \right) + P_{[i]} + (j \times R_{[i]}) \right\} =$$
$$\sum_{j=0}^{Q_{[i]}-1} \left\{ P_0 + \sum_{l=1}^{i} \left(P_{[l]} + R_{[l]} \right) + (j \times R_{[i]}) \right\} =$$
$$Q_{[i]} P_0 + Q_{[i]} \sum_{l=1}^{i} \left(P_{[l]} + R_{[l]} \right) + \frac{R_{[i]} Q_{[i]} (Q_{[i]} - 1)}{2} \qquad (24)$$

The breakdown of the claw TVT according to the right-hand-side of Eq. 24 implies that, in order to minimize the total visitation time over all claws, it suffices to minimize the quantity

$$\sum_{i=1}^{k} Q_{[i]} \sum_{l=1}^{i} \left(P_{[l]} + R_{[l]} \right)$$

This quantity is reminiscent of the problem of minimizing the weighted flow time of k jobs on a single server. The quantities $(P_i + R_i)$, $i = 1, \ldots, k$, correspond to the job processing times, and the quantities Q_i, $i = 1, \ldots, k$, correspond to the job weights. It is well known that for this single-machine scheduling problem, an optimal sequence is provided by the *"Weighted Shortest Processing Time (WSPT)"* rule: Jobs are processed in increasing ratios of their processing times over their weights [40]. When translated in the context of the considered TVT-problem, this result implies that the k claws must be processed in increasing order of the ratios $r_i \equiv (P_i + R_i)/Q_i$. Hence, the considered version of the TVT-problem is resolved very efficiently.

V. Conclusions

This paper introduced a new class of OMRP^{2}G problems that are characterized by the fact that the underlying robot fleet must observe some "connectivity" requirements while trying to reach a set of target locations. Also, due to the particular attributes of the applications that have motivated these problems and the nature of the imposed "connectivity" requirements, the guidepath network that supports the robot traffic is a tree. We have provided formal characterizations and complete mathematical programming formulations for these problems, and we have analyzed their worst-case computational complexity. In their general positioning, the considered OMRP^{2}G problems are NP-hard. But we have also demonstrated that it is possible to construct computationally efficient solutions for some instantiations of these problems, by identifying and exploiting special structure in the tree $\mathcal{T}$ that defines the guidepath network.

The presented developments have also revealed that the combinatorial structure that underlies the considered OMRP^{2}G problems, and defines their computational complexity and the prospective structure of their solution algorithms, is substantially different from the corresponding structure that underlies the OMRP^{2}G problems studied in the past literature. The new OMRP^{2}G problems present a stronger affinity to various machine scheduling problems than to the vehicle routing and the traveling-salesman problems that have been associated with the earlier OMRP^{2}G problems. At the same time, the interaction of the resources involved in this new class of problems – i.e., the robots and the zone buffering capacities – together with the imposed "connectivity" requirements generate a richer and much more complex set of resource allocation patterns and dynamics than those encountered in the more classical scheduling theory.

Our future work will seek to further analyze and understand these dynamics and their association to the existing scheduling theory. It will also employ the results and the insights that will be provided by this analysis, towards the development of solutions to the considered OMRP^{2}G problems able to manage systematically and effectively the need for a trade-off between operational and computational efficiency that was established by the results of Section IV of this paper. A third research task concerns the extension of the aforementioned developments to MMRS with guidepath networks of a more general topology.[12] Similarly, one can consider the extension of the methodological developments and the insights to be obtained from the aforementioned investigations, in order to support the expedient execution of coverage tasks that might evolve in less familiar and/or less structured environments, establishing, thus, a bridge between the research program

[12]Of course, all the complexity results of Section IV carry over immediately to this broader class of problems. But the corresponding MP formulations must account for the additional choice that is defined by the more general topology of the underlying guidepath network.

that has been specified in this work and some other research lines on networked robotic systems that were discussed in the introductory section. Another possible extension concerns the detailed specification and investigation of similar min-time coverage problems where, however, each robot has a distinct identity and role in the executed task. Finally, the aforementioned endeavors will not only provide a pertinent and powerful theory for the considered class of OMRP^{2}G problems, but they will also extend the boundaries of the current combinatorial scheduling theory.

REFERENCES

[1] L. E. Parker, D. Rus, and G. S. Sukhatme, "Multiple mobile robot systems," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer, 2016, pp. 1336–1379.

[2] D. Tardioli, A. R. Mosteo, L. Riazuelo, J. L. Villarroel, and L. Montano, "Enforcing network connectivity in robot team missions," *IJRR*, vol. 29, pp. 460–480, 2010.

[3] S. G. Loizou and C. C. Constantinou, "Multi-robot coverage on dendritic topologies under communication constraints," in *Proceedings of IEEE CDC 2016*. IEEE, 2016, pp. –.

[4] M. M. Zavlanos, M. B. Egerstedt, and G. J. Pappas, "Graph theoretic connectivity control of mobile robot systems," *Proceeding of the IEEE*, vol. 99, pp. 1525–1540, 2011.

[5] A. Muralidharan and Y. Mostofi, "Communication-aware robotics: Exploiting motion for communication," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 115–139, 2021.

[6] Y. Kantaros and M. M. Zavlanos, "Global planning for multi-robot communication networks in complex environments," *IEEE Trans. on Robotic Systems*, vol. 32, pp. 1045–1061, 2016.

[7] J. Stephan, J. Fink, V. Kumar, and A. Ribeiro, "Concurrent control of mobility and communication in multirobot systems," *IEEE Trans. on Robotic Systems*, vol. 33, pp. 1248–1254, 2017.

[8] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *IJRR*, vol. 23, pp. 939–954, 2004.

[9] P. Toth and D. Vigo (eds), *The Vehicle Routing Problem*. Philadelphia, PA: SIAM, 2001.

[10] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York, NY: W. H. Freeman and Co., 1979.

[11] H. N. Psaraftis, "Dynamic vehicle routing problems," *Vehicle Routing: Methods and Studies*, vol. 16, pp. 223–248, 1988.

[12] S. Irani, X. Lu, and A. Regan, "On-line algorithms for the dynamic traveling repair problem," *Journal of Scheduling*, vol. 7, pp. 243–258, 2004.

[13] P. Jaillet and M. R. Wagner, "Online routing problems: Value of advanced information and improved competitive ratios," *Transportation Science*, vol. 40, pp. 200–210, 2006.

[14] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith, "Dynamic vehicle routing for robotic systems," *Proceedings of the IEEE*, vol. 99, pp. 1482–1504, 2011.

[15] M. G. Lagoudakis, E. Markakis, D. Kempe, P. Keskinocak, A. Kleywegt, S. Koenig, C. Tovey, A. Meyerson, and S. Jain, "Auction-based multi-robot routing," in *Proceedings of the Robotics Science and Systems Conference*, 2005, pp. –.

[16] T. Ganesharajah, N. G. Hall, and C. Sriskandarajah, "Design and operational issues in AGV-served manufacturing systems," *Annals of OR*, vol. 76, pp. 109–154, 1998.

[17] I. F. A. Vis, "Survey of research in the design and control of automated guided vehicle systems," *European Journal of Operational Research*, vol. 170, pp. 677–709, 2006.

[18] S. S. Heragu, *Facilities Design (3rd ed.)*. CRC Press, 2008.

[19] S. A. Reveliotis, "Conflict resolution in AGV systems," *IIE Trans.*, vol. 32(7), pp. 647–659, 2000.

[20] M. P. Fanti, "Event-based controller to avoid deadlock and collisions in zone-controlled AGVS," *Intl. Jrnl Prod. Res.*, vol. 40, pp. 1453–1478, 2002.

[21] S. Reveliotis and E. Roszkowska, "Conflict resolution in free-ranging multi-vehicle systems: A resource allocation paradigm," *IEEE Trans. on Robotics*, vol. 27, pp. 283–296, 2011.

[22] J. Yu and S. M. LaValle, "Structure and intractability of optimal multi-robot path planning on graphs," in *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, 2013.

[23] R. M. Wilson, "Graph puzzles, homotopy, and the alternating group," *Journal of Combinatorial Theory, B*, vol. 16, pp. 86–96, 1974.

[24] D. Kornhauser, G. Miller, and P. Spirakis, "Coordinating pebble motion on graphs, the diameter of permutation groups, and applications," in *Proc. IEEE Symp. Found. Comput. Sci.* IEEE, 1984, pp. 241–250.

[25] V. Auletta, A. Monti, M. Parente, and P. Persiano, "A linear-time algorithm for the feasibility of pebble motion on trees," *Algorithmica*, vol. 23, pp. 223–245, 1999.

[26] J. Yu and S. M. LaValle, "Multi-agent path planning and network flow," in *Algorithmic Foundations of Robotics X*. Springer, 2013, pp. 157–173.

[27] J. Yu and D. Rus, "Pebble motion on graphs with rotations: Efficient feasibility tests and planning algorithms," in *Algorithmic Foundations of Robotics XI*, 2015.

[28] J. Yu and S. M. LaValle, "Optimal multirobot path planning on graphs: Complete algorithms and effective heuristics," *IEEE Trans. on Robotics*, vol. 32, pp. 1163–1177, 2016.

[29] H. Ma, C. Tovey, G. Sharon, S. Kumar, and S. Koenig, "Multi-agent path finding with payload transfers and the package-exchange robot-routing problem," in *AAAI 2016*, 2016, pp. 3166–3173.

[30] G. Daugherty, S. Reveliotis, and G. Mohler, "Optimized multi-agent routing for a class of guidepath-based transport systems," *IEEE Trans. on Automation Science and Engineering*, vol. 16, pp. 363–381, 2019.

[31] S. Reveliotis, "On the state liveness of some classes of guidepath-based transport systems and its computational complexity," *Automatica*, vol. 113 – 108777, 2020.

[32] S. Reveliotis, T. Masopust, and M. Ibrahim, "Polynomial-time optimal liveness enforcement for guidepath-based transport systems," *Nonlinear Analysis: Hybrid Systems*, vol. 41 – 101058, 2021.

[33] S. Reveliotis, "An MPC scheme for traffic coordination in open and irreversible, zone-controlled, guidepath-based transport systems," *IEEE Trans. on Automation Science and Engineering*, vol. 17, pp. 1528–1542, 2020.

[34] A. Ulusoy, S. L. Smith, X. C. Ding, C. Belta, and D. Rus, "Optimal multi-robot path planning with temporal logic constraints," in *Proc. of the 2011 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2011, pp. 3087–3092.

[35] C. Mahulea and M. Kloetzer, "Robot planning based on Boolean specifications using Petri net models," *IEEE Trans. on Automatic Control*, vol. 63, pp. 2218–2225, 2017.

[36] C. Baier and J. P. Katoen, *Principles of Model Checking*. London, England: The MIT press, 2008.

[37] C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems (2nd ed.)*. NY,NY: Springer, 2008.

[38] R. L. Rardin, *Optimization in Operations Research*. Prentice Hall, 1998.

[39] D. B. West, *Introduction to Graph Theory (2nd ed.)*. Upper Saddle River, N.J.: Prentice Hall, 2001.

[40] M. Pinedo, *Scheduling*. Upper Saddle River, NJ: Prentice Hall, 2002.

[41] C. H. Papadimitriou, *Computational Complexity*. Reading, MA: Addison-Wesley, 1995.