Preservation of Traffic Liveness in MPC Schemes for Guidepath-based Transport Systems

Spyros Reveliotis
School of Industrial & Systems Engineering
Georgia Institute of Technology
spyros@isye.gatech.edu

***Abstract*— Guidepath-based transport systems is a pertinent abstraction for the traffic that is generated in many contemporary applications, ranging from industrial material handling and robotics, to computer game animations and the qubit transport systems that are employed in quantum computing. In some recent works of ours, we have proposed the traffic coordination in this class of systems according to a model predictive control (MPC) scheme that seeks to maximize the traffic throughput while retaining computational tractability for the corresponding scheduling problem. In this work we perform a more systematic investigation of the conditions that must be observed by the adopted MPC scheme in order to ensure the liveness of the resulting traffic. The presented results span a number of possible configurations of the underlying guidepath-based transport systems, and integrate and extend a variety of past results concerning the liveness-enforcing supervision of AGV and other complex resource allocation systems.**[1]

I. INTRODUCTION

This paper concerns the traffic that is generated by a set of agents circulating on a connected graph which is known as the "(supporting) guidepath network". The "mission" trips of these agents on the guidepath network are specified by node sequences that must be visited by the agents in the indicated order. Furthermore, during their trips to these destinations, the agents must observe certain regulations that are dictated by safety considerations. In particular, the agents must be sufficiently separated during their traveling, a requirement that is practically enforced by stipulating that these agents cannot cohabitate on the same edge of the guidepath network during their trips. This last stipulation is enforced by a traffic coordinator, and it turns the agent traveling towards their various destinations into a *resource allocation process*, with the negotiated resources being the edges of the guidepath network.

From an application standpoint, the traffic problems outlined in the previous paragraph arise naturally in the real-time operations of various automated unit-load material handling (MH) systems, like the AGV, overhead monorail and the complex crane and gantry systems used in many production and distribution facilities [1], but also in the physical medium that implements the various elementary operations in the context of quantum computing [2]; the reader is referred to [3], [4], [5] for an elaboration on these connections. In addition, similar guidepath-based traffic models have drawn recently the attention of the robotics community (e.g., [6], [7], [8], [9]), while, in the past, they have been studied even by the broader CS community in the context of some classical games like the, so called, "15-puzzle" where 15 uniquely numbered "pebbles" located on a 4×4 grid have to be re-arranged in the row-major order by "pebble sliding" through the single unoccupied vertex of the grid [10], [11].

A primary concern for the resource allocation process that takes place in the aforementioned transport systems is the establishment of a high throughput, through the facilitation of expedient traveling of the running agents to their various destinations. This objective is attained through (i) a pertinent coordination of the agent traversal of the various contested edges, and (ii) the effective utilization of the routing flexibility that is defined by the topology of the underlying guidepath network. At the same time, an additional important concern for the traffic coordinator and the corresponding resource allocation process is to maintain the *"liveness"* of the generated traffic, by ensuring the agents' ability to reach their successive target nodes while avoiding potential deadlocking and livelocking situations.

Under some further operational assumptions for the underlying transport system,[2] the aforementioned concerns give rise to a combinatorial scheduling problem that can be effectively formulated as a mixed integer program (MIP); some characteristic examples of such MIP formulations can be found in [12], [13], [14], [8], [3]. However, the practical applicability of these MIP formulations to many industrial environments, and their solvability under the "real-time" constraints that are posed by these environments, is rather limited. Hence, in a recent research program of ours on the considered traffic management problem, we have suggested a decomposing approach that is inspired by the Model Predictive Control (MPC) theory [15] for more traditional optimal control problems.

Generally speaking, this MPC-based approach to the considered traffic scheduling problem decomposes the original problem to a series of simpler scheduling problems that seek to route efficiently the traveling agents to their most immediate destinations. Furthermore, the active schedules are updated every time that one of these agents reaches its current destination and it is ready to set out for the next one. But as in the more traditional MPC theory, an

[1]This work was partially supported by NSF grant ECCS-1707695.

[2]We shall detail all these assumptions in the subsequent, more technical parts of this paper.

additional important concern in this decomposition process is the establishment of the feasibility of the subproblem sequence that is generated by it. *In the context of the traffic scheduling problem that is addressed in this work, this last requirement is tantamount to the preservation of "traffic liveness" by the considered MPC scheme.*

This paper intends to take a more systematic look into the traffic liveness-preservation requirement that was outlined in the previous paragraph, and in this way define a more solid base for the implementation of the proposed MPC framework in the considered transport systems. As it will be revealed by the subsequent developments, the corresponding results are contingent upon certain structural and operational characteristics of the underlying transport system. Among these characteristics, some of the most prominent are (i) the ability of an agent to freely reverse its motion on any given edge of the guidepath network, and (ii) the availability of a "depot" where the agents retire upon the completion of their mission trips.[3] Perhaps not surprisingly, our results also rely pretty heavily upon past results concerning the liveness-enforcing supervision of AGV systems ([16], [17], [18], [19], [20]), a prominent and thriving research area within the Discrete Event Systems (DES) community. When viewed from this standpoint, the primary contribution of this work is (a) to "catalog" those past results that are currently scattered in various parts of the past literature, (b) complement them with some new developments and insights, and, most importantly, (c) organize and present them in a way that will be practically useful in the MPC schemes that are promoted in this work.

With the above basic positioning of the paper content and contributions, the rest of the paper is organized as follows: Section II introduces the main structural and operational assumptions for the transport systems that are considered in this work. It also defines the traffic scheduling problem addressed in this paper, and outlines the MPC scheme for this problem that has been advocated in our past developments. Section III constitutes the main part of the paper, presenting a series of results that culminate into a set of practical conditions that must be observed by the applied MPC scheme in order to ensure the liveness of the overall generated traffic. As remarked in the earlier parts of this section, the corresponding developments are organized according to certain attributes of the considered transport systems that facilitate these developments and define the scope of their applicability. Section IV discusses the integration of the results of Section III in the considered MPC framework, and, at the same time, identifies remaining open issues that arise from the presented developments. Finally, Section V concludes the paper, summarizing its contributions and highlighting directions for future research.

[3]A third dimension for extending the presented taxonomy of guidepath-based transport systems can be based on the particular restrictions that are imposed by the adopted zone allocation protocol. Due to space considerations, this work does not expand in this direction, but we provide some further comments on this issue in later parts of the manuscript.

II. THE CONSIDERED GUIDEPATH-BASED TRANSPORT SYSTEMS, THE CORRESPONDING TRAFFIC SCHEDULING PROBLEM, AND THE PROPOSED MPC FRAMEWORK

A formal characterization of the considered guidepath-based transport systems and the generated traffic. The guidepath-based transport systems considered in this work are formally abstracted as follows: The system consists of a guidepath graph $G = (V, E)$ that is traversed by a set of agents, $\mathcal{A}$. Graph G is assumed to be connected and undirected. The edges $e \in E$ of G model the "zones" of the underlying quidepath network. These edges can be traversed by a traveling agent $a \in \mathcal{A}$ in either direction, and, in general, they can hold no more than one agent at any time.

However, in many of these transport systems, the edge set E might contain a "self-loop" edge, denoted by h, and modeling a *"home"* location that can hold an arbitrary number of agents; in particular, edge h acts as a "parking" location for agents that are not on an active "mission" trip. Following rather standard terminology borrowed from the AGV literature, we shall refer to those guidepath-based transport systems that possess the aforementioned "home" edge h as *"open"*, and to the remaining ones as *"closed'*.

A "mission" trip for some agent $a \in \mathcal{A}$ is defined by a sequence of edges $\Sigma_a = \langle e \in E \rangle$ that must be visited by agent a in the specified order.[4] Furthermore, for open transport systems, the last edge in sequence Σ_a is the "home" edge h, in line with the aforementioned role of this edge.

While traversing an edge $e \in E$ with $e = \{v_i, v_j\}$, an agent a will have a certain direction of motion that will be indicated by the corresponding ordered pair (v_i, v_j) or (v_j, v_i). This representation of the agent motion also enables the following definition of the system *state*, that will be useful in this work:

Definition 2.1: At any point in time, the *state* s of the considered traffic system is defined by: (i) the edges $e_a = \{v_{i_a}, v_{j_a}\}$ that define the currently occupied zone by each agent $a \in \mathcal{A}$, and the agent's direction of motion in that zone; and (ii) the remaining visitation requirements, $\hat{\Sigma}_a$, for each agent.

State s evolves by having (at least) one of the system agents moving to a neighboring edge that is compatible with its direction of motion on its current edge, and by further updating the visitation sequences every time that a new target edge is reached by such a transition. Moreover, in some of the considered transport systems, an agent can reverse its direction of motion on its current edge, in which case this is an alternative mechanism for evolving the running traffic state s. In the following, we shall refer to transport systems that support reversibility of the agent motion on any given edge as *"reversible"*, and to the remaining ones as *"irreversible"*.

Since edges $e \in E$ model the specified zones of the underlying guidepath network G, they are assumed to be of

[4]E.g.,in the MHS operational setting, the edges $e \in \Sigma_a$ typically model pairs of pick-up and deposition locations that must be visited by the material handling agent during its trip.

equal length. This assumption, together with the presumed uniformity of the traveling agents $a \in \mathcal{A}$, allow us to further assume that the corresponding edge-traversal times are deterministic and uniform across all edges. This last duration defines a natural *"time unit"* for the considered models, and enables the discretization of the traffic dynamics of the underlying transport system.[5]

In the resulting discretized dynamics of the considered traffic, it is further stipulated that an agent a cannot move in an edge e at time t from a neighboring edge e', unless e was empty at time $t-1$. This assumption prevents potential agent cohabitation on a given edge during the transitional phases that lead from (discrete) epoch $t-1$ to epoch t, and it also implies that two agents cannot "swap" the occupation of two neighboring edges.[6] Most importantly for the subsequent developments of this work, these traffic restrictions, when combined with the arbitrary topology of the graph G and the bidirectional traversal of its edges by the traveling agents, can potentially give rise to deadlocking situations that will permanently stall the further advancement of the agents involved, and they necessitate the proactive management of the underlying traffic with the additional objective of preserving its liveness [16].

The considered traffic scheduling problem and the simplifying MPC framework. Having established the operational dynamics of the considered transport systems in the previous paragraphs, next we proceed with the definition of the basic traffic scheduling problem that motivates the main developments presented in this work. For this, we need the following definitions:

Definition 2.2: A *route* $\mathcal{R}_a$ for any given agent $a \in \mathcal{A}$ over a (discrete) time span T is a sequence $\langle (v_i, v_j)_a^t : t = 0, \ldots, T \rangle$ defining, for each epoch $t \in \{0, \ldots, T\}$, the edge of the guidepath network G that is occupied by agent a during that epoch, and the agent orientation (or direction of its motion) in this edge.

Definition 2.3: A set of routes $\{\mathcal{R}_a : a \in \mathcal{A}\}$ that satisfy the posed visitation requirements Σ_a for each traveling agent, is a *routing schedule* $\mathcal{S}$ (w.r.t. these visitation requirements). Schedule $\mathcal{S}$ is *feasible*, if the corresponding routes $\mathcal{R}_a$ are compatible with the topology of the underlying guidepath network G, and they abide to the aforestated operational assumptions and stipulations that characterize the agent maneuverability and the employed zone allocation protocol.

Furthermore, for any given routing schedule $\mathcal{S}$, the maximal traveling time required to meet all the posed visitation requirements across all the corresponding routes $\mathcal{R}_a,\ a \in \mathcal{A}$, will be characterized as the schedule *makespan*, and it will be denoted by $T^{\mathcal{S}}$.

Then, the basic traffic scheduling problem that is considered in this work can be defined as follows:

Definition 2.4: The basic traffic scheduling problem for the considered traffic systems: Given a state s_0 of the considered guidepath-based transport system, determine a routing schedule $\mathcal{S}$ that will satisfy the corresponding visitation requirements that are defined by state s_0, with minimal makespan $T^{\mathcal{S}}$.

In the following, we shall denote an optimal schedule by $\mathcal{S}^*$ and the corresponding makespan by T^*. Furthermore, as remarked in the introductory section, an optimal schedule $\mathcal{S}^*$ can be obtained, in principle, by formulating and solving a mixed integer program. But this MIP formulation does not scale well for many real-world applications. To cope with these practical challenges, we have proposed to use an MPC scheme towards the solution of the aforementioned scheduling problem [21]. This MPC framework decomposes the overall scheduling problem to a sequence of subproblems that seek the effective and efficient routing of the traveling agents towards their *most immediate destinations* in the corresponding sequences $\Sigma_a,\ a \in \mathcal{A}$. The details of this decomposition and the formulation of the corresponding subproblems must also ensure that the traffic flow generated by this decomposition remains *"live"*, i.e., that it maintains the ability of all agents to reach their corresponding target nodes, and in the case of open transport systems, eventually retire to the corresponding "home" edge h. We address the satisfaction of this "liveness" requirement in the next section, while in Section IV we further discuss the implications of the results of Section III for the overall organization of the considered MPC framework.

III. Preserving traffic liveness in the considered guidepath-based transport systems

In this section we provide conditions that must be observed by the MPC framework that was outlined in the previous section, in order to ensure the liveness of the resulting traffic. Furthermore, we want these conditions to be efficiently implementable in the real-time operational context of the considered transport systems. The derived results are contingent upon (i) the presence of a "home" edge h in the underlying guidepath network G, and (ii) the "reversibility" of the guidepath-based transport system;[7] hence, we organize their presentation accordingly.

A. Preserving traffic liveness in open and reversible guidepath-based transport systems

For this class of guidepath-based transport systems, we have the following fundamental result:

[5]Even in the case that the zone-traversal times are non-uniform in terms of the zone set E and/or the agent set $\mathcal{A}$, a discrete-time model for the dynamics of the corresponding traffic can be obtained by utilizing the greatest common divisor of the zone-traversal times, and refining the edge definition for the guidepath network G accordingly.

[6]We also notice that in a considerable part of the relevant literature, especially some lines of work that come from the Robotics and the Artificial Intelligence communities (c.f. [8]), agents cannot swap their locations, but an agent a can move into a zone while another agent a' is moving out of the same zone during the same discretized time step. This possibility essentially implies an ability of locally coordinated motion on the part of the agents, and impacts in a substantial, qualitative manner the dynamics of the underlying traffic, as it enables behaviors that are not possible under the more restrictive zone allocation protocol that requires agents to advance only into empty zones. These remarks define the third classification dimension that was suggested in Footnote 3.

[7]We remind the reader that the "reversibility" of the guidepath-based transport systems considered in this work was defined in the previous section.

Proposition 3.1: Let s_1, s_2 denote two traffic states of an open, reversible guidepath-based transport system. Then, there is always a feasible routing schedule $\mathcal{S}$ that can lead from state s_1 to state s_2.

The result of Proposition 3.1 was first established in [4]. A constructive argument establishing the validity of this proposition can be briefly structured as follows: First, the assumed reversibility of the considered transport system implies the existence of a feasible routing schedule $\mathcal{S}_1$ that collects all agents in the "home" edge h; a simple such schedule will route agents to edge h one at a time, starting with the one(s) that have the closest distance from h. Once all agents have been collected in edge h, another routing schedule $\mathcal{S}_2$ will take them to their locations specified by state s_2. Again, schedule $\mathcal{S}_2$ will route the agents to their respective locations one at a time, starting with the agent(s) with the longest destination from the "home" edge h.

In the context of the MPC schemes that are considered in this work, the practical implication of Proposition 3.1 is that the generated traffic will remain live, irrespective of the intermediate target states that define the decomposing logic of the applied MPC scheme; some reasonable realizations of this decomposing logic are discussed in Section IV.

B. Preserving traffic liveness in open and irreversible guidepath-based transport systems

We start our discussion for this class of guidepath-based transport systems by noticing that any meaningful realization of these systems must satisfy the following additional condition:

Condition 3.1: The employed guidepath network G has a minimal nodal degree of 2.

Indeed, in the class of guidepath-based transport systems considered in this subsection, if an agent a reaches a node v with degree one, will not be able to reverse its direction upon reaching that node and it will remain there forever. On the other hand, the following lemma establishes that Condition 3.1 guarantees the existence of a feasible route taking an agent $a \in \mathcal{A}$ on a round-trip from the "home" edge h to any node $v \in V$ and back to h.

Proposition 3.2: Consider an open, irreversible guidepath-based transport system satisfying Condition 3.1, and let s_0 denote the state where every agent $a \in \mathcal{A}$ is located at the "home" edge h. Then, for any given agent a and a node $v \in V$, there exists a feasible route $\mathcal{R}_a$ that takes agent a from edge h to node v and back edge h, while the remaining agents $a' \in \mathcal{A} \setminus \{a\}$ remain at edge h.

Proof: Since graph G is connected, there is a path p leading from edge h to vertex v. Let e' denote the last edge on path p. Since the minimal degree of G is 2, there is an edge $e_1 \equiv \{v, v_1\} \neq e'$. If v_1 belongs in path p, then, edge e_1 together with the segment of p leading from v_1 to h defines a path p' that can be followed by agent a for its return trip from node v to the "home" edge h. Otherwise, define $p_1 \equiv pe_1$, and notice that there exists an edge $e_2 \equiv \{v_1, v_2\} \neq e_1$. Repeating the previous argument with respect to node v_2 and path p_1, we can either identify a return path for agent a, or extend path p_1 to path $p_2 \equiv p_1 e_2$. Continuing in the same manner, and recognizing the finiteness of the nodal set V, eventually we shall reach to a node v_k that will belong in the constructed path p_{k-1} and will define the required return path for agent a. □

The next result is an immediate implication of Proposition 3.2.

Corollary 3.1: Consider an open, irreversible guidepath-based transport network satisfying Condition 3.1 at the state s_0 where every agent $a \in \mathcal{A}$ is located at the "home" edge h. Then, there is a feasible routing schedule for every set of visitation requirements $\{\Sigma_a : a \in \mathcal{A}\}$.

Proof: A simple (although not necessarily efficient) schedule will satisfy each visitation requirement for an agent $a \in \mathcal{A}$ and a node $v \in \Sigma_a$ by a distinct "round trip" of agent a from the "home" edge h to node v and back to h. The feasibility of this schedule is guaranteed by Proposition 3.2. □

Corollary 3.1 and its proof essentially establish that Condition 3.1 is necessary and sufficient for the satisfaction of any set of visitation requirements $\{\Sigma_a : a \in \mathcal{A}\}$ when the considered guidepath-based transport system is started from state s_0 where all agents are located in the "home" edge h. On the other hand, the computational complexity of assessing the feasibility of a set of visitation requirements $\{\Sigma_a : a \in \mathcal{A}\}$ when the system is started from some arbitrary state s, is currently an open problem.[8] Hence, we propose to restrict the operation of this class of guidepath-based transport systems in a subset of their state space S for which reachability of the state s_0, where every agent $a \in \mathcal{A}$ is located at the "home" edge h, can be established with polynomial complexity. Such a subset of states is provided by the concept of the *"h-ordered"* state.

Definition 3.1: A state s of an open, irreversible guidepath-based transport network satisfying Condition 3.1 is *"h-ordered"* if there exists an ordering $[\cdot] : \{1, \ldots, |\mathcal{A}|\} \to \mathcal{A}$, of the agent set $\mathcal{A}$, such that each agent $a_{[i]}$, $i = 1, \ldots, |\mathcal{A}|$, can advance to the "home" edge h from its current edge while agents $a_{[j]}$, $j = i+1, \ldots, |\mathcal{A}|$, maintain their original positions in state s.

The notion of h-ordered state adapts to the considered problem context the notion of "ordered" traffic-states for AGV systems that was introduced in [16]. In the following, we shall denote the set of h-ordered states by S_{ho}. Assessing whether any given state $s \in S$ is h-ordered reduces to the construction of a corresponding ordering. The search of such an ordering for the agent set $\mathcal{A}$ can be performed in a "greedy" manner (i.e., without the need for any backtracking) since the placement of any agent $a \in \mathcal{A}$ in the "home" edge h

[8] More specifically, the work of [20] provides a proof of the NP-completeness of the problem of assessing the feasibility of a set of visitation requirements $\{\Sigma_a : a \in \mathcal{A}\}$ in open irreversible guidepath-based transport systems, when the system is started from an arbitrary state s, and the agent routes are predetermined. On the other hand, under free agent routing, we can pursue a routing scheme similar to the routing scheme that was discussed in the proof of Proposition 3.1. But the complexity of determining the existence of a feasible routing plan that will collect all traveling agents from their current locations to the "home" location h, remains an open issue.

increases the set of free edges that can be used by the remaining agents for reaching the "home" edge h. The realization of this possibility goes back to the seminal work of Dijkstra and his development of *Banker's algorithm* for more general resource allocation systems [22]. However, in the considered operational setting of the open, irreversible guidepath-based transport systems, things are further complicated by the extensive levels of routing flexibility that is available to any given agent a. But this additional complication can be handled with polynomial complexity w.r.t. the size of the underlying guidepath network and the number of traveling agents, through the adaptation of some relevant algorithms that are presented in [16].

Furthermore, in [23] it is shown how the set of (h-)ordered states that is admitted by any efficient realization of Banker's algorithm can be effectively expanded through *controlled* partial search. Finally, it is easy to see that state s_0, as well as the states s that are involved in the proofs of Proposition 3.2 and Corollary 3.1, are h-ordered. Hence, starting from the natural initial state s_0, and using the algorithmic tools provided in [16] and [23], we can confine the operation of the underlying guidepath-based transport system in a subset S' of its state space S that is efficiently recognizable and ensures live operation for the generated traffic. In Section IV we also discuss how this capability can be utilized in the MPC framework that is proposed for the considered set of transport systems.

C. Preserving traffic liveness in closed and reversible guidepath-based transport systems

We start our discussion for this class of guidepath-based transport systems by noticing that for any meaningful realization of these systems we must have $|\mathcal{A}| < |E|$ (since, otherwise, no agent motion is possible). This inequality implies that there is always a free edge in the guidepath network; to facilitate the subsequent discussion, we shall refer to such a free edge as a *"hole"*. Then, we have the following lemma:

Lemma 3.1: For any given traffic state s of a closed, reversible guidepath-based transport system with $|\mathcal{A}| < |E|$, and an edge $e \in E$, there is a state s' reachable from s in which edge e is a "hole".

Proof: If edge e is a "hole" in state s, then we simply set $s' \equiv s$. Otherwise, let e' denote an edge containing a "hole" in state s that has the shortest possible distance from edge e, and let $p = \langle e \equiv e_0, e_1, \dots, e_l \equiv e' \rangle$ be a corresponding shortest path connecting e and e'. Then, according to the working assumptions, each edge $e_i,\ i = 0, 1, \dots, l-1$, is occupied by an agent a_i. Consider the state s' that is obtained from state s by advancing each agent $a_i,\ i = 0, \dots, l-1$, from its current edge e_i to edge e_{i+1}, starting with agent a_{l-1} and working in decreasing order of index i. Then, it is not hard to see that the hole at edge e' in state s has moved to edge e in state s'. □

However, while the condition $|\mathcal{A}| < |E|$ guarantees the effective move of a hole to any edge of network G, it is not sufficient to ensure that any agent $a \in \mathcal{A}$ can move from its current location to a target destination. A counter-example establishing the validity of this statement is presented in Figure 1. In the depicted situation, agent a_1 wants to move to edge e_2, and it also holds that $|\mathcal{A}| = 4 < |E| = 5$. But it is easy to check that the required transfer of agent a_1 is not feasible.

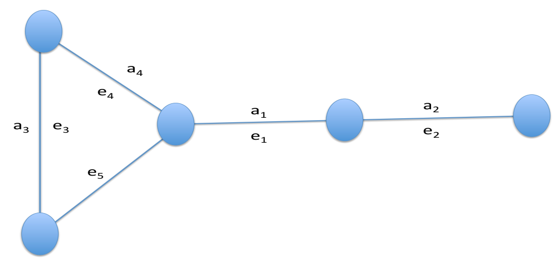

Fig. 1: A counter-example establishing that, for closed and reversible guidepath-based transport systems, the condition $|\mathcal{A}| < |E|$ is not adequate for ensuring the ability of an agent $a \in \mathcal{A}$ to advance from its current edge e to a target edge e'.

The problem in the example of Figure 1 arises from the presence of the path $e_1 e_2$. This path is characterized by the fact that all of its edges do not belong on any cycle,[9] and in the following discussion, we shall characterize these paths as *"singular"*. Also, we shall denote the set of singular paths in graph G by $\mathcal{P}_S$, and for any path $p \in \mathcal{P}_S$, $|p|$ will denote the "length" of p, i.e., the number of its edges.

Then, our main result for this class of guidepath-based transport systems can be stated as follows.

Proposition 3.3: In the class of closed and reversible guide-path-based transport systems, a sufficient condition guaranteeing that any agent $a \in \mathcal{A}$ can move from its current edge e to any other edge $e' \in E$ of the guidepath network G is that $|\mathcal{A}| \leq |E| - 1 - \max_{p \in \mathcal{P}_S}\{|p|\}$.

Proof (Sketch): Consider first the case where $\mathcal{P}_S = \emptyset$. Then, the condition of Proposition 3.3 becomes $|\mathcal{A}| \leq |E| - 1$. Also, consider a path $p = \langle e \equiv e_0, e_1, \dots, e_l \equiv e' \rangle$ from e to e'. Then, the working assumption $\mathcal{P}_S = \emptyset$ implies that there is a path from edge e_1 to a "hole" that does not include edge e_0, and working as in the proof of Lemma 3.1, we can move this "hole" to edge e_1. Hence, agent a can advance across path p by one edge, to edge e_1. Furthermore, iterative invocation of the above argument implies that there is a routing schedule that can take agent a all the way to edge e'.

When singular paths are present, the entire graph G can be uniquely decomposed to a "tree"-like structure $\mathcal{T}$, where the nodes of $\mathcal{T}$ are the maximal subgraphs $G_k, k = 1, \dots, K$, that are not containing any singular paths themselves, and the edges of $\mathcal{T}$ correspond to the singular paths that interconnect

[9]Following standard terminology of graph theory, in this work we define a cycle in an undirected graph as a simple path with coinciding starting and ending nodes. Furthermore, a path is simple if it does not revisit any of its vertices (except possibly the first and the last ones, in the case of a cycle).

pairs $\{G_i, G_j\}$ of the aforementioned subgraphs G_k. Then, transferring the considered agent a from its current edge e to edge e' will involve, in general, the traversal of a "path" Q in this "tree"-like structure consisting of some subgraphs G_k and the interconnecting singular paths. Let G_1 denote the maximal subgraph containing edge e. Lemma 3.1 guarantees that we can move a "hole" to subgraph G_1, and, subsequently, an argument similar to that provided in the first part of this proof further establishes that agent a can move between any pair of edges of the considered subgraph G_1. Also, let G_2 denote the maximal subgraph that is second in the aforementioned "path" Q. The condition $|\mathcal{A}| \leq |E| - 1 - \max_{p \in \mathcal{P}_S}\{|p|\}$ guarantees (i) that it is always possible to empty the singular path $p \equiv \{G_1, G_2\}$ required by the traveling agent a, while preserving the accessibility of agent a to this path, and (ii) the agent ability to enter the next required subgraph G_2: all that needs to be done is first to bring agent a to an edge e_1 of G_1 that is adjacent to the singular path p, and subsequently empty the path p of any other agents while ensuring the presence of a "hole" in subgraph G_2. Working in this way, agent a can advance through the entire "path" Q that connects edges e and e' in the aforementioned "tree"-like structure. □

In the case where $\mathcal{P}_S = \emptyset$, the resulting condition of Proposition 3.3 is also necessary for ensuring the ability of any agent $a \in \mathcal{A}$ to move from its current edge e to any other edge $e' \in E$ of the guidepath network G. On the other hand, when $\mathcal{P}_S \neq \emptyset$, the condition of Proposition 3.3 is only sufficient; characteristically, the reader can check that in the example of Figure 1, the circulating agents can reach any edge of the depicted guidepath graph as long as $|\mathcal{A}| \leq |E| - \max_{p \in \mathcal{P}_S}\{|p|\}$.[10]

The reader should notice that the condition of Proposition 3.3 is a structural condition for the underlying guidepath-based transport system that can be validated off-line. Furthermore, once this condition is established, the proof of Proposition 3.3 also provides an effective mechanism for developing a routing schedule for the transferring of agent a from edge e to edge e'. And since this mechanism involves only the identification of (shortest) paths for the necessary transfers of agent a and the facilitating "holes" during the various legs of the agent trip, this mechanism is also computationally efficient.

From a more conceptual standpoint, the result of Proposition 3.3 for the considered class of guidepath-based transport systems resembles the result of Proposition 3.1 for the open and reversible case, in that once a certain structural condition is established,[11] then, the liveness of the underlying traffic is automatically guaranteed. It is, thus, natural to attribute this property to the motion reversibility that is possessed by these transport systems, and as indicated in the closing part of Section III-A, the availability of this property simplifies substantially the implementation of the MPC schemes that are pursued in this work.

D. Preserving traffic liveness in closed and irreversible guidepath-based transport systems

As in the case of open and irreversible guidepath-based transport systems, closed and irreversible guidepath-based transport systems need more active real-time supervision for ensuring the liveness of the underlying traffic. A set of results for facilitating such efficient control logic has been developed in [19]. In this section we overview the main points of these past developments, in a way that they can be related to the MPC schemes that are the focus of this work.

Central in the developments of [19] is the *partially directed graph (PDG)* $G'(s)$ that is defined as follows:

Definition 3.2: Given a state s of the considered class of guidepath-based transport systems, the corresponding PDG $G'(s)$ induced from state s and the guidepath graph G, by substituting each edge e of G occupied by an agent $a \in \mathcal{A}$ in s with a directed edge that indicates the orientation / direction of motion of a on e.

A *(simple) path* p in $G'(s)$ is defined as any (simple) path p in the original graph G where, however, the directed edges introduced by the definition of $G'(s)$ have the same sense of direction. Furthermore, a *cycle* c in $G'(s)$ is a simple path with coinciding initial and terminal nodes. A *joint* between two cycles c and c' is a simple path that is a sub-path for both c and c'. On the other hand, a *pass* between two cycles c and c' is a singular path of the original guidepath graph G with its first node lying on c, its last node lying on c', and with all of its edges being undirected in the PDG $G'(s)$. Finally, the next set of concepts are at the core of the sought characterization of liveness for the considered class of guidepath-based transport systems:

Definition 3.3: A *chain* in PDG $G'(s)$ is the subgraph defined by a sequence $ch = \langle c_1, p_2, c_2, p_3, \ldots, p_n, c_n \rangle,\ n \geq 1$, such that (i) $c_i,\ i = 1, \ldots, n$, are cycles, (ii) $p_i,\ i = 2, \ldots, n$, are simple paths, and (iii) each path p_i is a joint or a pass between cycles c_{i-1} and c_i. Two edges $e, e' \in E$ are *chain-connected* – or, simply, *chained* – if there exists a chain that contains, both, e and e'. Furthermore, PDG $G'(s)$ is *chained* if every two edges $e, e' \in E$ are chained.

Chain connectivity defines an equivalence relationship on E, and the subgraphs of $G'(s)$ that are induced by the corresponding equivalence classes are the *chained components* of $G'(s)$. Also, the PDG $\mathcal{C}(s)$ that is obtained from $G'(s)$ by replacing each of its chained components by a simple node is called the *condensation* of $G'(s)$. Obviously, chained PDGs $G'(s)$ have condensations that correspond to a single node. An efficient algorithm for obtaining the condensation $\mathcal{C}(s)$ for any given PDG $G'(s)$ is provided in [19].

Finally, in order to state formally the main results of [19] that are of interest to this work, we also need to introduce the notion of a *live state* for the considered transport systems.

Definition 3.4: For the considered class of closed, irreversible guidepath-based transport systems, state s is *live*

[10]The perusal of the proof of Proposition 3.3 will also reveal that the condition $|\mathcal{A}| \leq |E| - 1 - \max_{p \in \mathcal{P}_S}\{|p|\}$ is also necessary as long as there exists a maximal singular path p that connects two nontrivial maximal subgraphs $G_i,\ G_j$ containing no singular paths.

[11]In the case of the open and reversible guidepath-based transport systems, this structural condition is the presence of the "home" edge h.

if and only if the state transition diagram $RG(s)$, that is induced by the reachable states from s, has a strongly connected component such that for each agent $a \in \mathcal{A}$ and each edge $e \in E$, there is a state s_a^e where agent a is located in edge e.

Then, utilizing also the notion of the singular paths $p \in \mathcal{P}_S$ of the guidepath network G that was introduced in Section III-C, the main result of [19] that is of interest to this work can be stated as follows:

Proposition 3.4: In a closed and irreversible guidepath-based transport system with $|\mathcal{A}| \leq |E| - \sum_{p \in \mathcal{P}_S} |p| - 2$, a given state s is live if and only if the set $R(s)$ of the states that are reachable from state s contains a chained state s'.

In [19] it is also argued that the condition $|\mathcal{A}| \leq |E| - \sum_{p \in \mathcal{P}_S} |p| - 1$ is necessary for being able to establish traffic liveness for closed and irreversible guidepath-based transport systems, and furthermore, the case of $|\mathcal{A}| = |E| - \sum_{p \in \mathcal{P}_S} |p| - 1$ can give rise to unavoidable livelocks. Hence, the corresponding condition of Proposition 3.4 can be perceived as practically necessary for being able to establish live traffic in the considered class of transport systems.

On the other hand, currently we do not avail of an efficient test to check the reachability of a chained state s' from any given state s. In view of this limitation, [19] proposes to confine the system operation in states that are either chained or *semi-chained*; the latter are obtained from chained states by transferring a single vehicle between two cycles c and c' over a pass p that connects these cycles. The developments of [19] guarantee that the aforementioned restriction will maintain the liveness of the underlying traffic. Of course, the resulting traffic coordinator is not maximally permissive anymore, but such a restriction is similar, in spirit, to the restriction that is imposed by the concept of the "h-ordered" state in the case of open and irreversible guidepath-based transport systems.

IV. DISCUSSION

The developments of Section III reveal that under motion reversibility for the guidepath-based transport systems considered in this work, the establishment of liveness for the underlying traffic requires only the satisfaction of a structural condition by the given system configuration. More specifically, in the case of open, reversible guidepath-based transport systems, the mere presence of the "home" edge h suffices for establishing the system liveness (c.f. Proposition 3.1). On the other hand, for closed, reversible guidepath-based transport systems, the corresponding condition is $|\mathcal{A}| \leq |E| - 1 - \max_{p \in \mathcal{P}_S}\{|p|\}$ (c.f. Proposition 3.3). When these conditions have been established for any given configuration from these two subclasses, accessing any reachable state in the resulting state space will not compromise the liveness of the underlying traffic.

The above remarks further imply that for the case of reversible guidepath-based transport systems, the intermediate target states to be employed by the considered MPC frameworks can be chosen arbitrarily. Furthermore, any other state that might be generated during the system advancement towards these intermediate target states, can function safely as an initial state in the case of a re-evaluation of the running routing schedule $\mathcal{S}$.

Realizing the extent of freedom that is described in the previous paragraph, [21] has pursued an MPC control scheme for open, reversible guidepath-based transport systems where the running target state is defined by the immediate targets of the traveling agents $a \in \mathcal{A}$ in their corresponding nodal sequences Σ_a. Furthermore, this target state is redefined every time that an agent reaches its current destination. At each intermediate phase of the resulting MPC scheme, the running schedule $\mathcal{S}$ is computed according to a local search scheme that first constructs an initial feasible schedule, $\mathcal{S}^0$, according to the constructive argument that was outlined in the discussion of Proposition 3.1, and subsequently improves iteratively the constructed schedule by seeking alternative shorter feasible routes for the agents that define the current makespan (i.e., the agents with the longest traveling times to their target destinations in the current schedule).

The extension of the MPC framework of [21] to the class of closed, reversible guidepath-based traffic systems, requires the provision of existence results and algorithmic procedures for the construction of the initial routing schedule in this new operational setting. In fact, the availability of such a routing schedule is contingent upon the selection of the next intermediate target state, $\tilde{s}(s)$, w.r.t. the current state s of the underlying traffic system. A set of seminal results for guiding this selection is provided by the following proposition, that initially appeared in [11]; the statement of this proposition is based on an alternative representation of the underlying guidepath structure and the supported traffic, where agents are located on the vertex set V' of the corresponding guidepath-graph G' and the edge set E' of G' defines the adjacency relationship among the available locations.

Proposition 4.1: Consider a closed, reversible guidepath-based transport system satisfying the condition of Proposition 3.3. Then, for any given pair (s, s') of traffic states, state s' is reachable from state s if any of the following conditions are satisfied:[12]

i. The corresponding guidepath graph G' is 1-connected.
ii. The corresponding guidepath graph G' is biconnected, not a polygon, and $|\mathcal{A}| < |V'| - 1$.
iii. The corresponding guidepath graph G' is biconnected, not a polygon or isomorphic to the graph T_0 depicted in Figure 2, $|\mathcal{A}| = |V'| - 1$, and the graph G' is not bipartite.
iv. The corresponding guidepath graph G' is biconnected, not a polygon or isomorphic to the graph T_0 depicted in Figure 2, $|\mathcal{A}| = |V'| - 1$, the graph G' is bipartite, and the permutation of the vertex set V' that is induced by the state pair (s, s') is even.

[12]We remind the reader that a connected graph G is *biconnected* iff the removal of a single vertex v from G retains the graph connectivity; otherwise, graph G is *1-connected*. Also, a permutation is *even* if it results from an even number of pairwise exchanges of the ordering of the underlying set.

Fig. 2: The special graphs involved in the statement of Proposition 4.1.

The proof provided in [11] for the above proposition also outlines an $O(|V'|^3)$ algorithm for constructing a feasible routing schedule from state s to state s'. Hence, this algorithm can provide the initial feasible schedule for extending the MPC framework of [21] to closed, reversible guidepath-based traffic systems. Furthermore, it is interesting to notice that the basic logic for the construction of these routing schedules is similar to the logic that provided the routing schedule in the proof of Proposition 3.1, but the role of the "home" edge h is now played by some other topological elements of the underlying guidepath graph G'.

In the case of irreversible guidepath-based transport systems, the intermediate target states $\tilde{s}(s)$ specified by the applied MPC scheme must be carefully selected in order to ensure the liveness of the generated traffic. Furthermore, these target states must be reachable by the running state s of the underlying traffic system. The particular sets of (i) h-ordered states and (ii) the chained and semi-chained states that were discussed in Section III, can provide an effective guideline for the generation of these target states for the respective cases of (a) open and (b) close irreversible guidepath-based transport systems in a way that ensures the traffic liveness. On the other hand, currently we lack any results similar to those of Propositions 3.1 and 4.1 that would resolve more general state-reachability requirements in the case of irreversible transport systems. Hence, the detailed implementation of a practical MPC-based traffic control framework for this class of guidepath-based transport systems remains quite open and it is part of our future investigations.

V. Conclusions

In this paper we have presented a set of results pertaining to the establishment of liveness for the traffic of guidepath-based transport systems that are controlled by an MPC framework. These results collect and extend a number of developments that have appeared in a sporadic / fragmented manner in the corresponding literature.

Our future work will seek to complete further the results presented in this paper, and to extend them to additional subclasses of guidepath-based transport systems. Such an interesting extension concerns the case defined by the more relaxed zone allocation protocol of Footnote 6. Furthermore, another case of considerable interest when moving in this direction, is that where the traveling agents must also observe "rendezvous" or other synchronization constraints.[13] Finally, a remaining important open issue concerns the exact characterization of the computational complexity of the decision problems of assessing state-liveness in the case of irreversible guidepath-based traffic systems.

References

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

[2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, UK: Cambridge University Press, 2010.

[3] G. Daugherty, S. Reveliotis, and G. Mohler, "Some novel traffic coordination problems and their analytical study based on Lagrangian Duality theory," in *Proceedings of the 55th IEEE Conf. on Decision and Control (CDC 2016)*. IEEE, 2016, pp. –.

[4] ——, "Optimized multi-agent routing in guidepath networks," in *Proceedings of the 2017 IFAC World Congress*. IFAC, 2017, pp. –.

[5] G. Daugherty, "Multi-agent routing in shared guidepath networks," Ph.D. dissertation, Georgia Tech, Atlanta, GA, 2017.

[6] T. Standley and R. Korf, "Complete algorithms for cooperative pathfinding problems," in *Proc. 22nd Intl. Joint Conf. Artif. Intell.*, 2011.

[7] Q. Sajid, R. Luna, and K. E. Bekris, "Multi-agent path finding with simultaneous execution of single-agent primitives," in *5th Symposium on Combinatorial Search*, 2012.

[8] 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.

[9] 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.

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

[11] 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.

[12] N. N. Krishnamurthy, R. Batta, and M. H. Karwan, "Developing conflict-free routes for automated guided vehicles," *Oper. Res.*, vol. 41, pp. 1077–1090, 1993.

[13] G. Desaulniers, A. Langevin, D. Riopel, and B. Villeneuve, "Dispatching and conflict-free routing of automated guided vehicles: An exact approach," *The Intl. Jrnl of Flexible Manufacturing Systems*, vol. 15, pp. 309–331, 2003.

[14] A. I. Corréa, A. Langevin, and L. M. Rousseau, "Scheduling and routing of automated guided vehicles: A hybrid approach," *Computers & Operations Research*, vol. 34, pp. 1688–1707, 2007.

[15] B. Kouvaritakis and M. Cannon, *Model Predictive Control: Classical, Robust and Stochastic*. London, UK: Springer, 2015.

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

[17] N. Wu and M. Zhou, "Resource-oriented Petri nets in deadlock avoidance of AGV systems," in *Proceedings of the ICRA'01*. IEEE, 2001, pp. 64–69.

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

[19] E. Roszkowska and S. Reveliotis, "On the liveness of guidepath-based, zoned-controlled, dynamically routed, closed traffic systems," *IEEE Trans. on Automatic Control*, vol. 53, pp. 1689–1695, 2008.

[20] S. Reveliotis and E. Roszkowska, "On the complexity of maximally permissive deadlock avoidance in multi-vehicle traffic systems," *IEEE Trans. on Automatic Control*, vol. 55, pp. 1646–1651, 2010.

[21] G. Daugherty, S. Reveliotis, and G. Mohler, "Optimized multi-agent routing in shared guidepath networks," in *Proceedings of SOLI 2017*. IEEE, 2017, pp. –.

[22] E. W. Dijkstra, "Cooperating sequential processes," Technological University, Eindhoven, Netherlands, Tech. Rep., 1965.

[23] M. Lawley, S. Reveliotis, and P. Ferreira, "The application and evaluation of Banker's algorithm for deadlock-free buffer space allocation in flexible manufacturing systems," *Intl. Jrnl. of Flexible Manufacturing Systems*, vol. 10, pp. 73–100, 1998.

[13] Such "rendezvous" requirements arise in the guidepath-based transport networks that model the CPU operations of quantum computers.