

Some new results on the state liveness of open guidepath-based traffic systems

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

Tomáš Masopust
Palacky University Olomouc &
Institute of Mathematics – Czech Academy of Sciences
masopust@math.cas.cz

Abstract—Guidepath-based traffic systems is a pertinent abstraction that has been used extensively by the Discrete Event Systems (DES) community for the study of the traffic dynamics that take place in the automated unit-load material handling systems (MHS) encountered in various production and distribution facilities. A particular problem that has drawn extensive attention in the DES-based investigation of these systems, is the preservation of the system “liveness”; i.e., the preservation of the ability of the system agents to complete successfully their currently allocated tasks and engage repeatedly to similar tasks in the future. The first part of this paper establishes that for a large subclass of the considered traffic systems, the preservation of their traffic liveness in a maximally permissive manner reduces to the observation of a particular property that must be possessed by the admitted traffic states. The second part of the paper provides some complexity analysis for assessing the aforementioned property on any given traffic state, under some further assumptions regarding the operation of the considered traffic systems and the structure of the traffic states under consideration.¹

Keywords: Guidepath-based traffic systems; traffic liveness enforcement; deadlock avoidance; discrete event systems

I. INTRODUCTION

Guidepath-based traffic systems is a pertinent abstraction that has been used extensively by the Discrete Event Systems (DES) community for the study of the traffic dynamics that take place in the automated unit-load material handling systems (MHS) encountered in various production and distribution facilities [1], [2], [3], [4], [5]. Perhaps the most familiar instantiations of these MHS are the automated guided vehicle (AGV) systems [6], [7], where a set of mobile robots are transferring parts among a set of well-defined locations in the underlying facility while moving on a set of lanes that isolate the traffic of these robots from the activity that takes place in the surrounding environment. Another popular realization of such automated MHS is the overhead monorail systems that are used in modern semiconductor manufacturing facilities (also known as “fabs”); in this second case, the guidepath network is physically defined by the monorails that are used to support the motion of the system vehicles [8]. Finally, a third major instantiation of the aforementioned unit-load MHS are the complex gantry crane systems that have been deployed at various logistical hubs, in particular, some major ports and railway yards [7].

In order to ensure collision-free traffic in all of the aforementioned environments, the various links of the underlying guidepath network are segmented into “zones” and

¹The first author was partially supported by NSF grant ECCS-1707695 and the second author by the GARC project GC19-06175J.




Fig. 1: An AGV deadlock.

it is stipulated that each zone can be occupied by at most one vehicle at a time. Occupation of a free zone by a traveling vehicle must be negotiated with a central traffic controller. In this way, the entire vehicle trip between an origin and a destination location becomes a “resource (i.e., zone) allocation process”, where the traveling vehicle must request and secure the different zones that it needs for the execution of its trip in a sequential manner.

From a control-theoretic standpoint, an important requirement for the aforementioned zone-allocation process, is *the preservation of the “traffic liveness”, i.e., the ability of all the system vehicles to complete their current assignments and engage successfully to similar assignments in the future operation of the system.* This ability can be compromised by the formation of deadlocks and livelocks among the traveling vehicles that might result from: (i) an arbitrary topology of the underlying guidepath network; (ii) an arbitrary specification of the vehicle paths in this network; (iii) a presumed irreversibility of the vehicle motion in their currently occupied zones; and (iv) the further inability of two vehicles that are located on neighboring zones to “swap” their zones.² A stylized deadlock formation taking place in the context of an AGV system is depicted in Fig. 1.

For a systematic investigation of the notion of the “traffic liveness” that was defined in the previous paragraph, and the corresponding control problem of liveness-enforcing su-

²All these assumptions about the structure of the considered MHS and their operational capabilities are standard assumptions in the corresponding DES literature; their justification can be traced in past publications in this area, including the various survey papers that are published periodically on the corresponding results, e.g. [9], [10].

pervision (LES)”³ for the considered traffic systems, it is pertinent to classify these systems into “open” and “closed”. The defining characteristic of the “open” subclass of such systems is the presence of a location in the underlying guidepath network where all idle vehicles can retire, possibly recharging their batteries and receiving other types of service. In the sequel, we shall refer to this location as the “home” location of the considered traffic system, and we shall model it as a zone in the underlying guidepath network with infinite accommodating capacity.

Furthermore, in the next section we shall show that, under certain natural assumptions regarding the connectivity of the underlying guidepath network, the requirement for liveness in open, zone-controlled, guidepath-based traffic systems of the type that are considered in this work, reduces to the requirement for an ability to collect all traveling vehicles to the “home” zone. A traffic state where such a routing scheme is possible, will be characterized as a “live” traffic state in the sequel. The notion of the “live traffic state” subsequently defines the structure of the maximally permissive LES for the considered traffic systems: this LES must admit or reject a zone-allocation request on the basis of whether the resulting traffic state is live or not. Hence, the ability to resolve effectively and efficiently the liveness of any given traffic state is critical for the effective deployment of the maximally permissive LES in the considered traffic systems.

In view of the remarks of the previous paragraphs, a more complete description of the contribution of this work to the problem of characterizing and assessing the notion of “(state) liveness” in open, zone-controlled guidepath-based traffic systems is as follows:

- 1) In the first part of the work, we provide a systematic characterization of the considered traffic systems, the notion of “traffic liveness” that must be observed by them, and the further notion of “state liveness” that is induced in this context.
- 2) The second part of the work focuses on the complexity of the decision problem of assessing state liveness. The main results of the paper along these lines differentiate further the considered class of traffic systems on the basis of whether the routes followed by the system vehicles during their various assignments are predetermined or not; the corresponding cases are respectively characterized as “static” and “dynamic” routing. In the context of this further classification, our main results can be described as follows:
 - a) For open, zone-controlled guidepath-based traffic systems with *static* vehicle routing, we provide a proof that the decision problem of (assessing) “state-liveness” is NP-complete. A similar result already exists in the literature [5]. But, in this paper, the aforementioned result and its proof are based on a slightly different set of assumptions from the assumptions that underlie the corresponding result of [5]. Furthermore, the new developments that are presented in this paper

³In the following, the “LES” acronym will mean either “liveness-enforcing supervision” or “liveness-enforcing supervisor”, according to the context.

reveal more clearly the strong affinity between the LES problem considered in this work and the broader problem of liveness-enforcing supervision for complex resource allocation that has been studied extensively in the DES literature [11].

- b) For open, zone-controlled guidepath-based traffic systems with *dynamic* vehicle routing, we consider the particular case of traffic states where the guidepath network is fully congested. For this case, we show that state liveness can be resolved with worst-case computational complexity that is polynomially related to the size of the underlying guidepath network.

On the other hand, currently, it is unclear whether it is possible to resolve polynomially the liveness of any *arbitrary* state coming from an open, zone-controlled guidepath-based traffic system with dynamic vehicle routing. The systematic investigation of this question is part of our ongoing research on the considered traffic systems. But at the same time, an additional value of this work is that it functions as an “exponent” of this particular research problem to the corresponding research community.

The rest of the paper is organized as follows: Section II provides a formal characterization of the considered traffic systems, and the various notions of “liveness” that are employed in this work. Subsequently, Section III provides our results on the complexity analysis of “state liveness” that were outlined in the earlier parts of this section. Finally, Section IV concludes the paper, and highlights some directions in our ongoing research on the class of problems that are considered in this work.

II. THE CLASS OF GUIDEPATH-BASED TRAFFIC SYSTEMS CONSIDERED IN THIS WORK AND THE CORRESPONDING NOTION OF LIVENESS

A formal modeling of the considered traffic systems:

An instance of the particular sub-class of the guidepath-based traffic systems considered in this work can be formally defined by a pair (\mathcal{A}, G) , where the elements of this pair denote, respectively, (a) the set of the system vehicles (or “agents”) circulating in it, and (b) the guidepath graph $G = (V, E \cup \{h\})$ that is traversed by these agents.

Graph G is assumed to be undirected, connected, and with a minimum vertex degree of 2.⁴ The edges $e \in E$ of G model the “zones” of the underlying guidepath network. These edges can be traversed by a traveling agent $a \in \mathcal{A}$ in either direction, and they can hold no more than one agent at a time. On the other hand, edge h models the “home” zone of the guidepath network. This edge is connected to the rest of the guidepath network through a single vertex (i.e., edge h is a self-loop of G), and it can hold an arbitrary number of agents that either have not initiated or have completed their assigned missions. Furthermore, in the following, we shall denote by v_h the vertex of graph G that is the single terminal vertex for the self-loop edge h . Finally, in the considered application context, it is also natural to assume that two vertices v_1, v_2 of graph G may be connected by more than

⁴As it will be revealed in the following, the imposed requirement of a minimal vertex degree of 2 is necessitated by the presumed irreversibility of the agent motion.

one zones, and therefore, in stricter terms, graph G is actually a multi-graph; but this feature does not impact substantially our subsequent developments, and we shall keep referring to G as a graph in the sequel.

A “mission” trip for an agent $a \in \mathcal{A}$ is defined by a sequence of edges $\Sigma_a = \langle e_i \in E \setminus \{h\} \rangle$ that must be visited by agent a in the specified order.⁵ Furthermore, edge h can be perceived as an implicit last edge in sequence Σ_a , a fact that signifies the requirement of retiring those agents a that have completed their mission trips to the “home” location.

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) . Furthermore, we stipulate that agents cannot switch the direction of their motion in the edges that are currently allocated to them; hence, an agent a entering edge $e = \{v_i, v_j\}$ from vertex v_i must leave this edge through vertex v_j , and vice versa.

An additional stipulation for the dynamics of the underlying traffic is that an agent a will move from its current edge e to a neighboring edge $e' \neq h$ only after it has been granted permission by the traffic controller, and such a permission can be granted by this controller only if the requested edge e' is free of any other agents. Besides preventing agent cohabitation in the different zones of the guidepath network, this last stipulation further implies that a set of agents cannot simultaneously swap their current locations.

As remarked in the introductory section, the impossibility of edge-swapping among the traveling agents, when combined with the arbitrary topology of the underlying guidepath network, can be a source of *deadlock* and *livelock* in the considered class of traffic systems [1], [4]. Such formations will prevent, or, more generally, restrict the future motion of the agents involved, and will impair the ability of these agents to complete their “mission” trips. Hence, an important task of the traffic controller is the preservation of the traffic “liveness”; i.e., the traffic controller must proactively prevent the development of deadlock and livelock by further restricting the admissibility of the “zone” allocations that are requested by the traveling agents. Next, we formalize further this “liveness” requirement for the considered traffic systems, by abstracting the corresponding dynamics through the formal modeling and analysis tools that are offered by qualitative DES theory [12].

An automaton-based representation of the considered traffic and a formal definition of the notion of “traffic liveness”: The qualitative – or “untimed” – dynamics of the guidepath-based traffic systems that were defined in the previous subsection, can be formally represented by an automaton $\Phi = \langle S, Q, f, s_0, S_M \rangle$ [12]. The state s of automaton Φ is defined by the following elements: (i) the placement of the system agents $a \in \mathcal{A}$ on the edges of the guidepath network G , formally represented by the function $e(\cdot; s) : \mathcal{A} \rightarrow E$; (ii) for agents $a \in \mathcal{A}$ with $e(a; s) \neq h$,

⁵In order to obtain a more concrete feeling of these “mission” trips, the reader can think of an AGV that, setting out from the “home” edge h , must perform a sequence of transports, where each transport involves the pick up of some material from the zone that is represented by edge e_i in sequence Σ_a and the deposition of this material to the zone represented by edge e_{i+1} .

their direction of motion in their allocated edges; and (iii) the remaining visitation requirements for each agent $a \in \mathcal{A}$, that are communicated in the corresponding edge-sequence Σ_a .

The resulting state set, S , that collects all possible states s of automaton Φ , will not be finite, in general, since the (remaining) visitation requirements associated with any agent $a \in \mathcal{A}$ can be any string $\Sigma_a \in (E \setminus \{h\})^*$, where $(E \setminus \{h\})^*$ denotes the Kleene closure⁶ of the edge set $(E \setminus \{h\})$.

On the other hand, the event set Q that advances the state s of the considered automaton is finite, and it contains two different types of events: (i) The first event type consists of all those events q that advance a single agent $a \in \mathcal{A}$ from its current edge $e(a; s)$ to a free neighboring edge e' , under the further condition that this advancement is also compatible with the direction of motion of the corresponding agent a on its current edge $e(a; s)$. Furthermore, in the case that this advancement also satisfies the next visitation requirement for agent a , the corresponding event will update accordingly the sequence of the remaining visitation requirements Σ_a . (ii) The second type of events in $q \in Q$ will expand the list of visitation requirements, Σ_a , for some agent $a \in \mathcal{A}$, by appending a single new requirement at the end of the current list Σ_a . Furthermore, in the rest of this discussion, we consider the type I events of Q as *controllable* by the controller that manages the traffic in the considered traffic systems, while type II events are *uncontrollable*.

The state transition function $f : S \times Q \rightarrow S$ of the automaton Φ provides a formal representation of the transitional dynamics that are implied by the above definition of state s and the event set Q . Furthermore, following [12], we assume f to be a partial function that is defined only for those (s, q) pairs for which the corresponding state transition is feasible within the scope of the aforesaid operational assumptions. We also extend f in the set $S \times Q^*$ in the natural manner, and we use the notation $R(s)$ to denote the states s' of Φ that are reachable from a given state s , through the dynamics that are defined by the extended function f ; i.e., $\forall s' \in S, s' \in R(s) \iff \exists \sigma \in Q^* : s' = f(s, \sigma)$.

Finally, we also notice, for completeness, that the elements s_0 and S_M in the tuple that defines the considered automaton Φ denote, respectively, the initial state of this automaton and the set of its marked states. These elements will be further specified in the later parts of this discussion.

In the context of the traffic dynamics that are described by the above automaton Φ , a first intuitive characterization of the notion of “traffic liveness” is as follows: We shall say that the traffic generated by any instance of the considered traffic systems is “live” when initialized at some traffic state s_0 , if and only if (*iff*) for every reachable state $s \in R(s_0)$, there exists an event sequence σ , consisting of type I events only, that will satisfy all the agent visitation requirements at state s . On the other hand, the presumed uncontrollability of the type II events of the automaton Φ also implies that, at the state $s' = f(s, \sigma)$, the visitation requirements for each agent $a \in \mathcal{A}$ might be updated arbitrarily to some new sequences Σ'_a , and since the state s'' that will result from these updates

⁶We remind the reader that the Kleene closure, X^* , of a finite set X consists of all the finite strings with elements from the set X , including the empty string ϵ .

is a reachable state of Φ , it must also satisfy the “liveness” condition that was stipulated for state s . The insights that are provided by these remarks can be formalized by defining the notion of “liveness” for the considered traffic systems as follows:

Definition 2.1: Consider the automaton Φ abstracting the dynamics of a guidepath-based traffic system considered in this work, and let $s_0 \in S$ be an arbitrary initial state for this automaton. Then, the traffic that is represented by automaton Φ is *live*, *iff*

$$\forall s \in R(s_0), \forall a \in \mathcal{A}, \forall e \in E \cup \{h\}, \exists s' \in R(s) : e(a; s') = e$$

□

Definition 2.1 has two important properties:

- 1) It does not consider explicitly the visitation requirements Σ_a , $a \in \mathcal{A}$.
- 2) The assessment of the “reachability” condition that eventually defines the notion of liveness in this definition, can be resolved by analyzing the projected dynamics of the automaton Φ on the event set $Q^I \subset Q$, that contains only the type I events.

The realization of these two facts is important since they further imply that the liveness of the considered traffic systems can be studied by means of a simplified automaton Φ' that is obtained from the original automaton Φ by:

- eliminating the visitation requirements Σ_a , $a \in \mathcal{A}$, as a component of state s ,
- substituting the original event set Q by Q^I , and
- restricting accordingly the state transition function f .

The above redefinition of state s is especially important for the subsequent developments, since the new state set S' , for the new automaton Φ' , is finite; i.e., Φ' is a finite state automaton (FSA) [12].

Furthermore, in an effort to prevent a potential profusion of the employed notation, in the following, we shall denote this new FSA by Φ , as well, and we shall carry over to this new automaton all the notation that was originally introduced in the context of the original automaton Φ .

State liveness: In [4] it is further shown that the liveness condition of Definition 2.1 is equivalent to the following condition that must be satisfied by every state $s \in R(s_0)$.

Proposition 2.1: Consider the automaton Φ abstracting the dynamics of a guidepath-based traffic system of the type that is considered in this work, and let $s_0 \in S$ be an arbitrary initial state for this automaton. Then, the resulting traffic that is represented by automaton Φ is *live*, *iff* for every state $s \in R(s_0)$, the corresponding subspace $R(s)$ contains a strongly connected component $\Psi(s)$ that satisfies the following condition:

$$\forall (a, e) \in \mathcal{A} \times E \cup \{h\}, \exists s' \in \Psi(s) : e(a; s') = e$$

□

The liveness condition of Proposition 2.1 is more amenable to the control objective of liveness-enforcing supervision for the considered traffic systems, since it implies that such a traffic system has the potential to exhibit live behavior as long as it is kept in those traffic states s that satisfy this condition. More specifically, the *maximally permissive liveness-enforcing supervisor (LES)* for these traffic systems

will allow transition to any state $s \in R(s_0)$ *iff* state s satisfies the condition of Proposition 2.1. Hence, in the following, we shall characterize a state s that satisfies the liveness condition of Proposition 2.1, as “*live*”, and we shall denote the entire subset of live states by S_l .

However, assessing state liveness for any given traffic state s through the characterization of Proposition 2.1 requires a global view of the corresponding subspace $R(s)$, and therefore, such a test will not be easily tractable for most practical instantiations of the considered traffic systems. Fortunately, in the rest of this section we establish that, for the considered FSA Φ , state liveness can be succinctly characterized by some alternative (co-)reachability condition that must be satisfied by any state s of practical interest in the operation of the considered traffic systems.

A central position in the following developments of this section is held by the particular state of FSA Φ where all agents $a \in \mathcal{A}$ are collected in the “home” edge h . We shall characterize this state as the “home” state of Φ , and we shall denote it by s_h . Obviously, state s_h defines a natural initial state for the considered traffic systems. It also defines a natural “target” state for every state $s \in R(s_h)$, since edge h constitutes the final destination of every executed “mission” trip; in the more formal semantics of FSA-based modeling, the role of state s_h as the “target” state for any state $s \in R(s_h)$ is expressed by setting $S_M = \{s_h\}$ for FSA Φ . Next, we show that, for the case of open guidepath-based traffic systems, a necessary and sufficient condition for the liveness of any state $s \in R(s_h)$ is, indeed, its co-reachability to the “home” state s_h .

Theorem 2.1: In the class of open guidepath-based traffic systems that are considered in this work, a state $s \in R(s_h)$ is *live iff* it is co-reachable to the “home” state s_h .⁷

Proof: First we establish the necessity of the co-reachability condition of Theorem 2.1 for the liveness of the considered state s . Without loss of generality, suppose that $s \neq s_h$, and let a_1 denote an agent with $e(a_1; s) \neq h$. Then, according to the condition that defines state liveness in Proposition 2.1, there is a feasible event sequence σ that takes agent a_1 to the “home” edge h . Furthermore, the definition of edge h implies that we can obtain a subsequence σ' of σ that transfers agent a_1 to h without relocating the agents that are already on h in state s . Let the resulting state be denoted by s_1 . If $s_1 = s_h$, then, the co-reachability condition of Theorem 2.1 has been met. Otherwise, select an agent a_2 with $e(a_2; s) \neq h$, and repeat the above argument. Since every invocation of this argument increases the number of agents that are located on edge h by one, and the entire set of agents, \mathcal{A} , is finite, it follows that eventually we shall reach a state where all agents are located on edge h , i.e., state s_h .

For the sufficiency part of the proof, first we notice that, since the guidepath graph G of the considered traffic systems has a minimal vertex degree of 2, when the system is in state s_h , there is always a feasible event sequence σ that takes any given agent $a \in \mathcal{A}$ to any given edge $e \in E$ and

⁷We notice, for completeness, that the notion of “state liveness” of Theorem 2.1 is equivalent to the notion of “state safety” for open guidepath-based transport systems that has been used in some past works on liveness-enforcing supervision for these traffic systems (e.g., [1]).

brings a back to the “home” edge h .⁸ Hence, state s_h is part of a strongly connected component of $R(s_h)$ that satisfies the state-liveness condition of Proposition 2.1. But then, any state $s \in R(s_h)$ that is co-reachable to state s_h satisfies this condition as well. \square

In [14] it is shown that in the class of open, dynamically routed guidepath-based traffic systems where the traveling agents can reverse the direction of their motion in their allocated zones, it holds that $S_l = S$; i.e., the co-reachability condition of Theorem 2.1 will always be satisfiable for every state $s \in R(s_h)$. Indeed, if the traveling agents can reverse the direction of their motion in their current zones, it is always possible to send these agents back to the “home” edge h , starting with the agents that are located closer to this edge; then, the necessary edges for each agent route are guaranteed to be free.

But for irreversible guidepath-based traffic systems, such a routing scheme might not be possible, since the traveling agents that are closest to the “home” edge h , might have the “wrong” orientation in their current edges. In fact, in this new regime, the co-reachability of any state $s \in R(s_h)$ to state s_h is not guaranteed any more, due to potentially unavoidable formations of deadlocks and livelocks. Hence, for irreversible guidepath-based traffic systems, the set of live states, S_l , is usually a strict subset of the state set S , and as already remarked, the traffic controller must be able to resolve effectively and efficiently the condition $s \in S_l$, for any given state $s \in S$. In the rest of the paper we present some results regarding the computational complexity of this decision problem.

III. SOME COMPLEXITY RESULTS ON ASSESSING THE STATE LIVENESS IN THE CONSIDERED TRAFFIC SYSTEMS

A. An NP-completeness result

We start the developments of this section by showing that under a *static* routing of the system agents, the problem of assessing the state-liveness condition of Theorem 2.1 on any given state $s \in S$ is NP-complete [15]. As remarked in the introductory section, a similar result was first established in [5] as a corollary to some complexity results that concerned the assessment of state liveness for traffic systems where the traveling agents are free-ranging over a certain area. The current developments provide an alternative proof for this result that enhances its applicability, and also reveals more vividly the connection of the considered traffic management problems to the broader resource allocation problems that have been studied by certain groups of the DES community.

In more specific terms, the NP-completeness proof that we shall develop in the rest of this section, is based on a polynomial reduction to the considered (traffic) state-liveness problem of the problem of assessing the “state safety in linear single-unit resource allocation systems (L-SU-RAS)”;⁹ this last problem is formally defined as follows [11]:

Definition 3.1: L-SU-RAS state safety with unit resource capacities: Consider a set of m reusable resources $\mathcal{R} = \{R_1, \dots, R_m\}$ and another set of n process instances $\Pi = \{J_1, \dots, J_n\}$ that need to utilize these resources for their execution. More specifically, each process instance

J_j , $j = 1, \dots, n$, is defined by a resource sequence $\mathcal{S}_j = \langle R[1; j], \dots, R[l_j; j] \rangle$; $R[k; j] \in \mathcal{R}$, $\forall k \in \{1, \dots, l_j\}$, that constitutes the corresponding “process plan” and must be interpreted according to the following semantics: Process instance J_j , $j = 1, \dots, n$, currently holds exclusively resource $R[1; j] \in \mathcal{S}_j$ and it further needs the sequential and exclusive allocation of the remaining resources in \mathcal{S}_j in order to advance to its completion. The allocation of the system resources to these process instances is coordinated by a central controller, and a requested resource allocation is feasible only if the considered resource is currently free. Furthermore, a process instance J_j will release its currently allocated resource, $R[k; j]$, only after it has been granted the next required resource, $R[k + 1; j]$, in the corresponding process plan \mathcal{S}_j . Finally, the system controller will grant any resource allocation requests that satisfy the aforesated conditions one at a time (and will recheck the feasibility of the remaining requests in the RAS state that will result from the execution of the selected allocation). We need to resolve whether there exists a resource allocation sequence for advancing process instances J_j , $j = 1, \dots, n$, through their various processing stages that are defined by the corresponding process plans \mathcal{S}_j ; more specifically, this resource allocation sequence must be feasible w.r.t. the aforesated resource allocation protocol, and it must allow each process instance J_j to complete successfully the corresponding process plan \mathcal{S}_j .

In [5] it is shown that the decision problem of Definition 3.1 is NP-complete in the strong sense. Next we use this result in order to establish the following result:

Theorem 3.1: The problem of assessing the liveness of any given traffic state $s \in S$ of an open, zone-controlled guidepath-based traffic system with pre-specified agent routes is NP-complete in the strong sense.

Proof: Consider a traffic state $s \in S$, and let $\sigma \in Q^*$ denote a feasible event sequence that leads from state s to the “home” state s_h . Under the working assumptions, the remaining route for each traveling agent is completely pre-specified. Hence, the length of sequence σ is exactly equal to the sum of the stages in all these remaining routes, and therefore, polynomially related to the problem data. Furthermore, the validity of sequence σ can be assessed through simulation, and this task is also of polynomial complexity w.r.t. the size of the underlying traffic system. Therefore, the considered problem is in NP.

In order to establish NP-completeness for this problem, we shall reduce to it the L-SU-RAS state safety problem that was introduced in Definition 3.1. So, consider an instance of this second problem, and let s denote the corresponding RAS state. The traffic state s' that will be constructed by the proposed reduction is depicted in Figure 2. The corresponding guidepath network possesses a central node N_c and $m + 1$ edges e_i , $i = 1, \dots, m + 1$, that are incident to this node in a “hub & spoke” sense. At the second node of each edge e_i there is a “self-loop” edge; for edges e_i , $i = 1, \dots, m$, the corresponding “self-loop” edge represents resource R_i , while the “self-loop” edge at the end of edge e_{m+1} is the “home” edge h . Each process instance J_j , $j = 1, \dots, n$, is represented in the constructed traffic state s' by an agent a_j located at the “self-loop” edge that corresponds to resource $R[1; j]$; this is

⁸A complete formal proof of this fact can be found in [13].




Fig. 2: The traffic state s' that is constructed in the reduction of the proof of Theorem 3.1.

indicated by representing this “self-loop” edge as a directed edge (the exact sense of direction is not important for this construction). Finally, the agent corresponding to process instance J_j must visit each of the “self-loop” edges that correspond to the resources $R[k; j]$, $k = 2, \dots, l_j$, according to the sequence that is specified by the corresponding process plan \mathcal{S}_j , and furthermore, it cannot visit any other edge that is not absolutely necessary for the realization of this process plan.

It is not difficult to see that, under the aforesaid specification of the route to be followed by each traveling agent, the construction of the previous paragraph essentially defines a bisimulation between the original RAS dynamics and the dynamics of the induced traffic system. Hence, the original RAS state s will be safe *iff* the induced state s' is live. Furthermore, it is clear that the size of the employed representation of the constructed state s' is related polynomially to the size of the employed representation for the RAS state s . Hence, the claim of Theorem 3.1 is true. \square

A careful study of the structure and the dynamics of the induced traffic system that is defined in Fig. 2, will also reveal that the above proof of Theorem 3.1 does not require the assumption of the irreversibility of the agent motion within their current zones. Hence, this assumption was not included in the statement of the result of this theorem. These remarks further imply that the high complexity of the “state liveness” problem that is addressed in this section essentially results from the complete pre-specification of the agent routes that is presumed by Theorem 3.1. On the other hand, the perusal of the proof of the corresponding complexity result that is provided in [5], will reveal that that proof relies substantially on the irreversibility of the agent motion within their allocated zones that is presumed by that result. This differentiation of the significance of the “motion irreversibility” assumption in each of these two results is due to the fact that, in the reduction of [5], the agent routes are partially restricted, but not completely pre-specified.

The remarks in the previous paragraph reveal the sensitivity of the complexity results that are pursued in this paper and in [5] on the detailed operational assumptions for the underlying traffic system. With this realization in mind, in the next subsection, we shift attention to the complementary subclass of open, zone-controlled guidepath-based traffic

systems where the system agents are dynamically routed to their various destinations.

B. An instantiation of the considered “state-liveness” problem that admits polynomial solution

In this section, we consider the complexity of assessing state liveness in open, zone-controlled guidepath-based traffic systems where the system agents are dynamically routed to their various destinations. Furthermore, since according to the closing discussion of Section II, the combination of (i) an open structure for the guidepath network, (ii) reversibility of the agent motion in their allocated zones, and (iii) dynamic agent routing implies that $S_l = S$, we also assume that the direction of the agent motion in their currently allocated zones is irreversible. Finally, in order to formally state the main result of this section, we also need to introduce the following concept:

Definition 3.2: A state s of an open, zone-controlled guidepath based traffic system is *totally congested* *iff* the zone corresponding to every edge $e \in E$ of the guidepath network G is occupied by a traveling agent $a \in \mathcal{A}$.

Then, the main result of this subsection is stated as follows:

Theorem 3.2: Consider an open, zone-controlled guidepath based traffic system (\mathcal{A}, G) where the direction of the agent motion in their currently allocated zones is irreversible. Also, let $\hat{s} \in S$ denote a totally congested state of this traffic system, and consider the directed graph $\hat{G} = (V, E)$ that is induced from the original graph G and the considered state \hat{s} by assigning to each edge $e \in E$ the direction of motion on edge e of the agent $a \in \mathcal{A}$ that occupies the corresponding zone in state \hat{s} . Then, state \hat{s} is live *iff* every vertex $v \in V \setminus \{v_h\}$ of \hat{G} is co-reachable to the vertex v_h .

Proof: We remind the reader that, according to Theorem 2.1, in the considered class of guidepath systems, a state s is live *iff* it is co-reachable to the “home” state s_h . We shall utilize this characterization of the state liveness in order to prove the result of Theorem 3.2.

First, we prove the sufficiency of the condition that is stated in Theorem 3.2 for the liveness of the considered state \hat{s} . Since every vertex $v \in V$ is co-reachable to vertex v_h , for every vertex $v \in V$, there exists a simple directed path $\pi(v)$ that leads from v to v_h .⁹ Consider such a path $\pi(v)$ for some $v \in V$. Then, it is easy to see that all agents $a \in \mathcal{A}$ that are located on the edges of this path in state \hat{s} can be transferred to the “home” zone h , one at a time, starting with the agent a on the edge e of $\pi(v)$ that is incident to vertex v_h . Let \hat{s}' denote the state that results from the considered state \hat{s} through the clearance of the aforementioned paths $\pi(v)$, $v \in V$, from their occupying agents. Also, let $e = (v, v')$ denote an edge of the digraph \hat{G} that does not belong to any of the paths $\pi(v)$. From the above definition of state \hat{s}' , it should be clear that (i) edge (v, v') is occupied by an agent a in state \hat{s}' that is heading towards vertex v' , and (ii) agent a can be routed all the way to the “home” edge h through the freed path $\pi(v')$. Since edge (v, v') was chosen arbitrarily, it follows that all edges

⁹We remind the reader that in a directed graph G , a simple (directed) path π of length n is an edge sequence $(v_0, v_1), (v_1, v_2), \dots, (v_{n-1}, v_n)$ with $v_i \neq v_j$ for any pair (i, j) such that $i = 0, 1, \dots, n-1, j = i+1, \dots, n$.

$e \in E$ that are still occupied by agents in state \hat{s}' can have their agents routed to “home” edge h . Therefore, state s_h is reachable from state \hat{s}' , and consequently, from state \hat{s} .

Next, we prove the necessity of the co-reachability condition of Theorem 3.2 for the liveness of the considered state \hat{s} . For this, we prove the contrapositive result that if the co-reachability condition of Theorem 3.2 does not hold in state \hat{s} , then, state \hat{s} is not live. Since, by the working assumption, the co-reachability condition of Theorem 3.2 does not hold in state \hat{s} , the vertex set V of digraph \hat{G} can be partitioned to two nonempty subsets V_c and V_n , with subset V_c (resp., V_n) containing the vertices that are (resp., are not) co-reachable to vertex v_h . Furthermore, since the guidepath graph G is connected, the cut \mathcal{C} among the vertex sets V_c and V_n is non-empty. The definition of the partition $\{V_c, V_n\}$ also implies that every edge $e \in \mathcal{C}$ is occupied by an agent that moves towards the vertex v of e that belongs in V_n . Then, consider the subgraph of digraph \hat{G} that is induced by the edges $e \in \mathcal{C} \cup (E \cap V_n \times V_n)$. It is not hard to see that, in the considered state \hat{s} , each of these edges is occupied by an agent $a \in \mathcal{A}$, and all these agents are entangled in a deadlock. Hence, state \hat{s} is not live. \square

The next corollary is an immediate implication of Theorem 3.2.

Corollary 3.1: Assessing the liveness of a totally congested state s of an open, zone-controlled guidepath based traffic system (\mathcal{A}, G) where the direction of the agent motion in their currently allocated zones is irreversible, is a task of linear complexity w.r.t. the size of the underlying guidepath network G .

Proof: Let \hat{G}^r denote the digraph that is obtained from the digraph \hat{G} of Theorem 3.2 by reversing the direction of its edges. Then, in digraph \hat{G}^r , the co-reachability condition of Theorem 3.2 reduces to the requirement for reachability of each vertex $v \in V \setminus \{v_h\}$ from vertex v_h through a simple directed path of \hat{G}^r . Clearly, this last condition can be tested in linear computational complexity w.r.t. the size of the digraph \hat{G}^r [16], and therefore, w.r.t. the size of the original graph G . \square

It is evident from the proof of Theorem 3.2 that this result, and therefore, the result of Corollary 3.1 as well, rely substantially on the assumption of dynamic routing of the system agents that is stated in that theorem. But equally substantial for the proof of Theorem 3.2 is the assumption that the considered state \hat{s} is fully congested. An important open problem is the computational complexity of assessing state liveness for open, zone-controlled guidepath based traffic systems (\mathcal{A}, G) where the direction of the agent motion in their currently allocated zones is irreversible but the considered state s is not totally congested. This issue is part of our ongoing investigations.

IV. CONCLUSIONS

The first part of the paper has defined a notion of “traffic liveness” for a class of guidepath traffic systems that abstracts the operation of the automated, unit-load MHS employed in modern production and distribution environments. This notion of “traffic liveness” has also been reduced to more localized conditions that must be satisfied by each traffic state that is admissible by the maximally permissive LES for

these systems. The second part of the paper has undertaken a systematic investigation of the computational complexity of assessing the state-based conditions that characterize traffic liveness, and it has demonstrated the strong dependence of this complexity on the particular assumptions that define the structure and the operation of the underlying system, and also the structure of the state under consideration.

A first part of our future work will seek to extend the aforementioned complexity study to states coming from open, zone-controlled, dynamically routed guidepath-based traffic systems and possessing more arbitrary structure than the structure that is presumed in the results of Section III-B. In fact, some additional results in this direction can be found in [17], while an efficient, but non-polynomial worst-case-complexity algorithm that can resolve the liveness of any given traffic state from the considered class of traffic systems, can be found in [18]. Furthermore, a longer-term objective of our research program is the embedding of all the obtained results into the Model Predictive Control (MPC) scheme for the real-time traffic management of the considered traffic systems that has been outlined in [13].

REFERENCES

- [1] S. A. Reveliotis, “Conflict resolution in AGV systems,” *IIE Trans.*, vol. 32(7), pp. 647–659, 2000.
- [2] 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.
- [3] M. P. Fantl, “Event-based controller to avoid deadlock and collisions in zone-controlled AGVS,” *Int. J. Prod. Res.*, vol. 40, pp. 1453–1478, 2002.
- [4] 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.
- [5] 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.
- [6] W. L. Maxwell and J. A. Muckstadt, “Design of automatic guided vehicle systems,” *IIE Trans.*, vol. 14, pp. 114–124, 1982.
- [7] S. S. Heragu, *Facilities Design (3rd ed.)*. CRC Press, 2008.
- [8] D. Pillai, “The future of semiconductor manufacturing: Factory integration breakthrough opportunities,” *IEEE Robotics & Automation Magazine*, vol. 13-4, pp. 16–24, 2006.
- [9] 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.
- [10] 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.
- [11] S. A. Reveliotis, *Real-time Management of Resource Allocation Systems: A Discrete Event Systems Approach*. NY, NY: Springer, 2005.
- [12] C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems (2nd ed.)*. NY, NY: Springer, 2008.
- [13] S. Reveliotis, “Preservation of traffic liveness in MPC schemes for guidepath-based transport systems,” in *IEEE CASE 2018*. IEEE, 2018.
- [14] 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.
- [15] 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.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, 2nd ed.* Boston, MA: MIT Press, 2001.
- [17] S. Reveliotis and T. Masopust, “Efficient liveness assessment for traffic states in open, irreversible, dynamically routed, zone-controlled guidepath-based transport systems,” *IEEE Trans. on Automatic Control*, vol. (cond. accepted).
- [18] —, “Efficient assessment of state liveness in open, irreversible, dynamically routed, zone-controlled guidepath-based transport systems: The general case,” School of Industrial & Systems Eng., Georgia Tech, Tech. Rep. (submitted for publication), 2019.