2017 IEEE International Conference on Service Operations and Logistics, and Informatics
September 18-20, 2017
Bari, Italy

Optimized Multi-Agent Routing in Shared Guidepath Networks

Greyson Daugherty and Spyros Reveliotis
School of Industrial & Systems Engineering
Georgia Institute of Technology
Atlanta, GA, USA

Greg Mohler
Georgia Tech Research Institute
Georgia Institute of Technology
Atlanta, GA, USA

***Abstract*— This paper overviews a research program of ours that concerns the routing of a number of traveling agents over a transportation medium that is abstracted as a "zone"-controlled guidepath network. From an application standpoint, this class of problems arise in the real-time management of the traffic that takes place in many contemporary material handling systems. But quite interestingly, the same problems also arise in the routing of the ionized atoms that are the elementary information carriers in the context of quantum computing. The paper provides a conceptual description and some motivational applications for the aforementioned problems, places these problems in the context of the existing literature, and outlines the methodological base and the key results that have been developed in the presented research; detailed expositions of these technical developments can be found in the provided references.**

Keywords: Guidepath-based Traffic Management, Combinatorial Scheduling Theory, Lagrangian Duality

I. Introduction: Problem Description and Motivation

The research program presented in this document addresses the problem of routing a number of traveling agents along the edges of a connected graph, which is referred to as a "guidepath network." Every agent has a set of designated destinations to which it is being routed. The objective of this research is to find ways to minimize the time required to route all agents in the system to their respective destinations, simultaneously. Since the considered guidepath network is shared by the entire set of agents, the routes followed by these agents will be subject to certain congestion – or "coupling" – constraints that prevent any agent's route from causing conflict with the routes of the remaining agents. Thus, this research aims at providing efficient, conflict-free, multi-agent routes between arbitrary origins and destinations within a highly restricted transportation environment.

A particular "real-world" application that motivates this research is that of flexibly automated, unit-load, zone-controlled material handling (MH) systems implemented in various production and distribution settings [1]. In this case, the vehicle traffic is restricted to a well-defined network (i.e., a "guidepath network"). The structure of this guidepath network may be defined naturally by the physical structure of the MH system, as in the case of (i) the crane or gantry systems used at certain major ports, and some heavy-industrial manufacturing plants, and (ii) the monorail systems used in modern semiconductor fabs. Alternatively, the guidepath network may also be an externally imposed structure, as in the case of automated guided vehicles (AGVs); in this last case, the traveling vehicles are restricted to certain corridors of a production or distribution layout in order to avoid collisions with workers and other shop-floor equipment.

The guidepath network traversed by a set of AGVs (or other MH systems) is frequently divided into "zones" designed to prevent collisions among the traveling vehicles, by allowing only one vehicle to reside in a zone at any given time. Allocation of a zone to a traveling vehicle must be coordinated through a control mechanism. Furthermore, the employed zoning scheme induces a natural discretization of the vehicle trips into "legs," wherein a route for any given vehicle can be described by a sequence of adjacent zones that define the agent's discrete location at each time step.

Another problem instance that motivates this work (and maps neatly into the discretization process just described) is that of coordinating the movement of quantum bits (known as "qubits") inside the central processing unit of a quantum computer [2]. Much like AGVs, qubits must travel between storage locations and "gate" locations, where they can interact with one another or be subjected to phase-change operations. And much like AGVs, qubits move along a specific guidepath network. The exact physical implementation may vary (e.g., from "ion-trap" to "quantum-dot" configurations [2]), but the underlying routing problem is essentially independent from the type of quantum computer.

The above two application contexts share certain operational characteristics, such as (i) the restrictions against agent co-residency at points on the guidepath network that do not serve either as source or destination locations, along with (ii) a "no-swap" rule that prevents agents from switching locations between two consecutive time steps. This last restriction ensures, among other things, that the model will not expect two AGVs traveling in opposite directions down a narrow aisle to switch places instantaneously, an operation that is not physically possible.

There are also some differences between the AGV operational context and that of quantum-computing. For instance, it is often assumed that AGVs do not reverse direction within an aisle, but this behavior is acceptable for qubits, and it may, indeed, be helpful in order to solve the problems addressed in this work. Another consideration is that AGV systems are often assumed to have no destination locations that can also be used as transit locations by other vehicles. In the quantum-computing context, however, this overlap between transit and destination locations may be very common since this possibility offers a scheduler extra freedom when routing agents to their respective destinations. From the standpoint of

978-1-5090-5847-1/17/$31.00 ©2017 IEEE

the traffic management problems considered in this work, the aforementioned overlap implies a number of quite subtle and difficult constraints: given that some edges of the guidepath network that constitute destinations for certain agents are also used for transit by other agents, it is imperative that no set of agents residing in their destination locations form a cut of the guidepath network that separates an agent still in transit from its own destination.

The bottom line for all of the aforementioned applications is that agents, irrespective of whether they are qubits or AGVs, must travel efficiently between various origin and destination locations while avoiding certain detrimental or forbidden behaviors. In particular, multiple agents cannot occupy the same location at the same time during transit, and two agents cannot swap places from neighboring locations in the guidepath network. The system must also prevent the formation of "deadlocks" – i.e., states where multiple agents block each other's advancement within the guidepath network – that might arise due to the previous two restrictions. The model and the algorithms developed for this research are designed with sufficient generality to cover both types of problem instances and are even adaptable to other contexts (e.g., flexible job-shop scheduling) that have slightly different system behaviors or constraints than the ones discussed here.

II. Literature Review

A number of researchers have studied in depth the problem of eliminating deadlocks from the considered traffic systems, along with a similar problem known as "livelocks"; in both of these problems, a set of traveling agents reach a set of traffic states from which they are incapable of reaching their destinations. Hence, these two types of problems are united by one imperative: The system state (which can sufficiently be described by the location and the direction of all agents on the guidepath network at a given moment in time, and also their remaining visitation requirements) must always be "safe"; i.e., there must exist a deadlock- and livelock-free schedule that can return every agent to the system "depot" after they have successfully visited all their destinations. The system "depot" itself can be thought of as a single or even a broader set of locations in the guidepath network, with the capacity to retain all agents in the system and to dispatch them in arbitrary order.

Reveliotis [3] formalizes the problem of deadlock avoidance for the considered traffic systems by modeling it as a supervisory control problem of Discrete Event Systems (DES) theory [4], [5], and provides a computationally efficient method for enforcing liveness within AGV systems based on a non-trivial extension of Dijkstra's Banker's algorithm [6]. Reveliotis and Roszkowska [7] show that, for the considered traffic systems, maximally permissive deadlock avoidance (i.e., deadlock avoidance that guarantees that a system state will be classified as safe if and only if at least one multi-agent schedule exists to route all agents back to the "depot" location) is an NP-hard problem. Additional work in the DES-related literature that has coped with the problem of deadlock avoidance in the considered traffic systems can be found in [8], [9], [10], [11]; a significant part of these works adapts to the considered traffic systems more general results regarding the broader problem of deadlock avoidance for sequential resource allocation that is extensively treated in [12], [13].

While the aforementioned deadlock-avoidance algorithms can be used to guarantee that any prescribed traffic schedule cannot cause deadlocks, they do not address the problem of optimizing the performance of these schedules with respect to some time-based objective. In the considered operational context, such a frequently used objective is the time that will take all agents to complete their running trips and return to the system depot, a concept that is technically characterized as the schedule "makespan"; the research program outlined in this document uses the schedule makespan as the primary objective function.

One perspective that can be used to approach the resulting scheduling problem is that of "job-shop" scheduling [14]. When the considered multi-agent routing problem is mapped to the "job-shop" scheduling paradigm, "machines" are replaced by the zones of the guidepath network, and the executed "jobs" correspond to the agent trips between different locations. The literature on job-shop scheduling offers a variety of methods to approach this problem. For instance, these methods include certain adaptations of the more general "branch-and-bound" algorithm to this particular problem [14], [15]. The corresponding algorithms have the potential to find optimal solutions, but usually they will execute in prohibitively long times; this last effect is especially true for the real-time traffic-management context that is considered in this work. Job-shop scheduling algorithms also include disjunctive programming [14], [16] and various efficient heuristics such as the "shifting-bottleneck" heuristic [17], [18]; but these algorithms cannot accommodate easily the extraordinary flexibility that exists in many guidepath networks.

There is also some prior research that addresses more explicitly routing problems similar to those considered in this work, coming from the communities of industrial engineering (IE) and operations research (OR). The most prominent of these works can be found in [19], [20], [21]. However, the perusal of this material reveals that, for the most part, these works essentially seek to adapt to the considered problem context, perspectives and methodology that were originally developed for the more traditional "vehicle routing" problem [22], that has been extensively addressed by those two communities. As a result, the proposed solutions (i) fail to provide a thorough treatment of the more behavioral / structural issue of deadlock avoidance, and (ii) they are also computationally intensive, to the point that they may not be practically applicable to the "large-scale" and "real-time" operational context that is defined by the underlying applications.

As a case in point of the limitations that are experienced when trying to apply the aforementioned lines of work to the guidepath-based traffic systems that are considered in the presented research program, next, we focus, in particular, to the approach that is presented in [21] (see also [23]). This approach routes agents iteratively, one at a time, while imposing a set of time-window constraints, in order to find solutions that are efficient and (often) deadlock-free. More

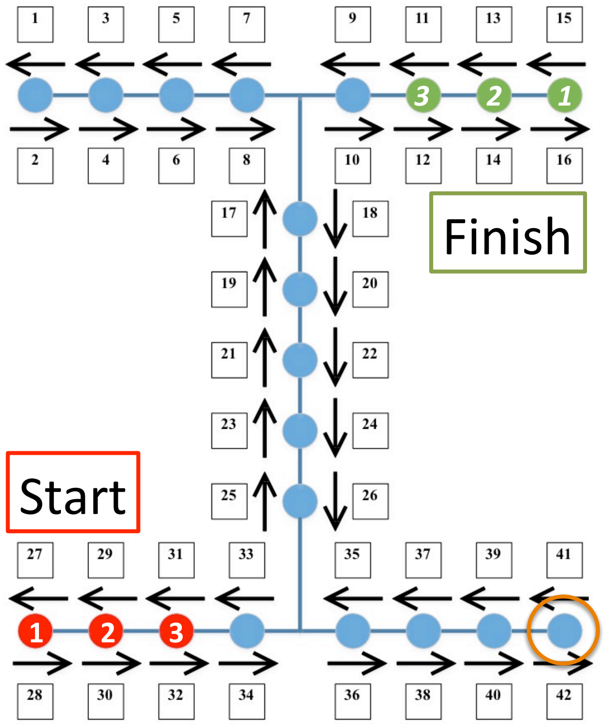

Fig. 1: This figure represents a problem where Agents 1, 2, and 3 must travel from the lower-left corner of the guidepath network to the upper-right, after reversing their order. A "depot" location is circled in the lower-right corner, where agents may be pooled and re-dispatched, but doing so would not be a strong solution to this problem. The arrows depicted in Figure 1 constitute a modeling abstraction that is useful especially for contexts where agents cannot necessarily reverse their direction of travel. The resulting directed graph, and the associated labeling and notation, exemplify the abstracting representation of the dynamics of the considered traffic systems and of their state, that is used by the considered research program.

specifically, the route of each agent is determined according to a shortest-path problem with time windows (SPPTW), where the time windows prevent conflicts with the agents routed earlier. The *ad hoc* decomposition of the overall routing problem that is adopted by this routing method, renders it computationally more efficient than some of the other methods that were cited in the previous paragraph, and possibly tractable in the real-time context of the routing problems that are considered in this work. Yet, there are problem instances where this method will fail to provide any solution.

For a more concrete example of the aforementioned limitations, consider Figure 1. The routing problem depicted in this figure requires three agents to start in the lower-left corner and finish in the upper-right, *after reversing their order.* If one were to attempt to route these three agents using the aforementioned SPPTW approach for each agent, then one would first find that Agent 3 *must* be the first to move. After that, *no conflict-free path will exist for either Agent 2 or Agent 3 to reach its destination.* Hence, approaching this problem instance as a series of SPPTWs would return no solution.

The ability of a small problem like that discussed in the previous paragraph to defeat some of the most promising existing algorithms hints at (i) the very high difficulty of the considered routing problems, and (ii) the need to pay close attention to all the operational constraints that are defined by the spatio-temporal dynamics of the underlying traffic. In the considered example, it is important that no subset of the agents that have come to rest at their respective destinations forms a cut blocking any agents still in transit from their own destinations. One of the key contributions of the presented research is its ability to recognize and address systematically all these complications and nuances.

III. Goals and Outline of the Presented Research Program

As already stated in the previous section, the presented research program addresses the problem of finding routing schedules for the considered multi-agent traffic systems that are conflict-, deadlock- and livelock-free, and also have minimized makespans.

Conceptually, each agent trip starts and finishes with the agent at the depot, and the trip itself is formally specified as a sequence of "target" locations that must be visited by the traveling agent, in the specified order. Such a formal specification of the "trip" concept is necessary for formalizing the underlying traffic dynamics and for stating and proving important properties for the agent motion, like its safety and its liveness. However, from a more practical operational standpoint, the synthesized solution will retain the possibility to redirect agents that have completed the key milestones of their trip to another mission, before actually returning to the depot. On the other hand, the problem of the optimized assignment of the various mission trips to the available agents is not part of the considered research program, but it constitutes a natural theme for follow-up work.

From a methodological standpoint, the considered research program must strike a pertinent balance between (i) the aforestated objective of minimizing the makespan of the adopted traffic schedules while ensuring the safety and the liveness of the resulting traffic, and (ii) the super-polynomial computational complexity of this task. We seek to establish this balance through a Model Predictive Control (MPC) framework [24], that decomposes the overall routing problem on the basis of the most immediate destinations of the traveling agents. Hence, this MPC framework seeks to develop an optimized route for each traveling agent, but only with respect to its next destination. At the same time, the resulting decomposition must ensure the safety and the liveness of the entire traffic that is generated by this approach; this last set of requirements can be perceived as a notion of "stability" customized to the considered operational context.

Within this MPC framework, the currently available results from the presented research program can be summarized as follows:

1) The program has formalized the aforementioned MPC

framework for the considered traffic scheduling problems, at a level of abstraction that applies to all the application contexts that were discussed in the previous sections of this document.
2) It has developed a complete Mixed Integer Programming (MIP) formulation for the "core" scheduling (sub-)problem of the considered MPC framework, that concerns the expedient routing of the traveling agents to their immediate destinations. This formulation renders the addressed problem amenable to the perspectives and the techniques of combinatorial scheduling theory [14].
3) It has developed a Lagrangian relaxation approach for the MIP formulation of item #2 above [25], [26], that can provide (a) strong lower bounds to this formulation, and (ii) potential "seeds" for the construction of near-optimal solutions to the original scheduling problem.
4) A particularly strong development in the context of item #3 above is the development of a set of customized approaches for the exact and expedient solution of the corresponding (Lagrangian) dual problem, a challenging task in itself [26]; these approaches take advantage of a distributed representation of the sub-differentials involved, which allow (i) the effective identification of ascending directions in the context of certain ascent-type methods, and even (ii) the reduction of the entire dual problem to a single linear program.
5) An additional set of results, complementary to the results of items #3 and #4 above, concerns the development of novel heuristic algorithms for the considered scheduling problems that take the form of "local search"-based methods [15]. We should also notice at this point that for the considered scheduling problems, it can be shown that even the construction of just a feasible solution is, in general, an NP-hard problem [7]. However, the considered research has identified a number of conditions that will render quite tractable the construction of a first initial solution. Subsequently, the employment of a pertinent specification and representation of the "neighborhood" structure to be used in the pursued local search, together with some novel dynamic programming (DP) formulations for effecting this search, enable the expedient identification of some very efficient schedules, even for some very hard problem instances.

Detailed, technical treatments of the aforementioned results can be found in the following more technical publications of ours: [27], [28] for items #1 – #4 in the above list, and [29], [30] for items #1 and #5. In the next section, we also present results from some numerical experimentation that demonstrates and assesses more empirically the efficacy of the aforementioned developments.

IV. Some Numerical Results

In a first set of experiments, we have sought the application of the methodology that was outlined in Section III to the particular problem instance that is depicted in Figure 1. The intricacies of this problem instance that were discussed during its introduction in Section II, have rendered it as some sort of "benchmark" in the corresponding community and literature.

TABLE I: An optimal set of agent routes computed for the problem instance of Figure 1 – adapted from [30].

t	a_1	a_2	a_3	t	a_1	a_2	a_3	t	a_1	a_2	a_3
0	28	30	32	6	25	35	37	12	12	17	21
1	28	30	34	7	23	33	37	13	14	10	19
2	28	32	25	8	21	25	35	14	16	12	17
3	30	34	36	9	19	23	33	15	16	14	10
4	32	25	38	10	17	21	25	16	16	14	12
5	34	36	38	11	10	19	23				

The application of the heuristic algorithms discussed in item #5 of the previous section to the aforementioned problem instance gave an optimal set of routes that is tabulated in Table I. In particular, Table I reveals that the optimal makespan for the considered problem is 16 periods, and the agent routes are represented by the edge sequences that are occupied by each agent $a_i,\ i = 1, 2, 3$, over the horizon $\{0, \ldots 16\}$. The detailed tracing of the reported routes manifests the extent of coordination and "intelligence" that must be exhibited by the considered set of agents.[1] On the other hand, our algorithms were able to compute this optimal schedule in 25 msecs while running on a simple MacBook Pro.

Another line of experimentation, originally reported in [30], has sought to assess more systematically the performance of the routing algorithms that have been developed in our work, and to identify some factors that might impact this performance. This experiment was executed by means of the guidepath network that is depicted in Figure 2. More specifically, Figure 2 depicts a "dual" version of the employed guidepath graph, where the various zones are represented by the nodes of the depicted graph, while the edges that connect these nodes define the neighboring structure of these zones. As it can be seen in Figure 2, the employed guidepath network consists of 133 zones, organized in a grid structure.

The red node at the center of the graph depicted in Figure 2 indicates the location of the "depot". More specifically, in our experiments, the depot was placed at two different locations: (i) the middle of the guidepath network, as indicated in Figure 2, and also (ii) one of the four corners of the depicted graph.[2] The results that are reported in the rest of this section reveal that the depot placement can have a significant impact on the performance of the presented algorithm.

The problem instances addressed in the considered experiment for each of the two depot placements involved a number of agents ranging from 3 to 45, with a step-increase of three agents. The starting and the destination locations for each agent were determined randomly. Furthermore, the construction of these problem instances was such that the

[1] To facilitate this route tracing, we must also notice that (i) the considered problem instance assumes that agents can reverse the direction of their motion within any given edge, and (ii) the computed schedules enforce the separation requirement stated in Section I that an agent cannot move into a zone that was occupied in the previous period. However, none of these two assumptions is critical for the effective implementation of the presented methodology.

[2] Due to the symmetries of the graph of Figure 2, all these corners are topologically equivalent.

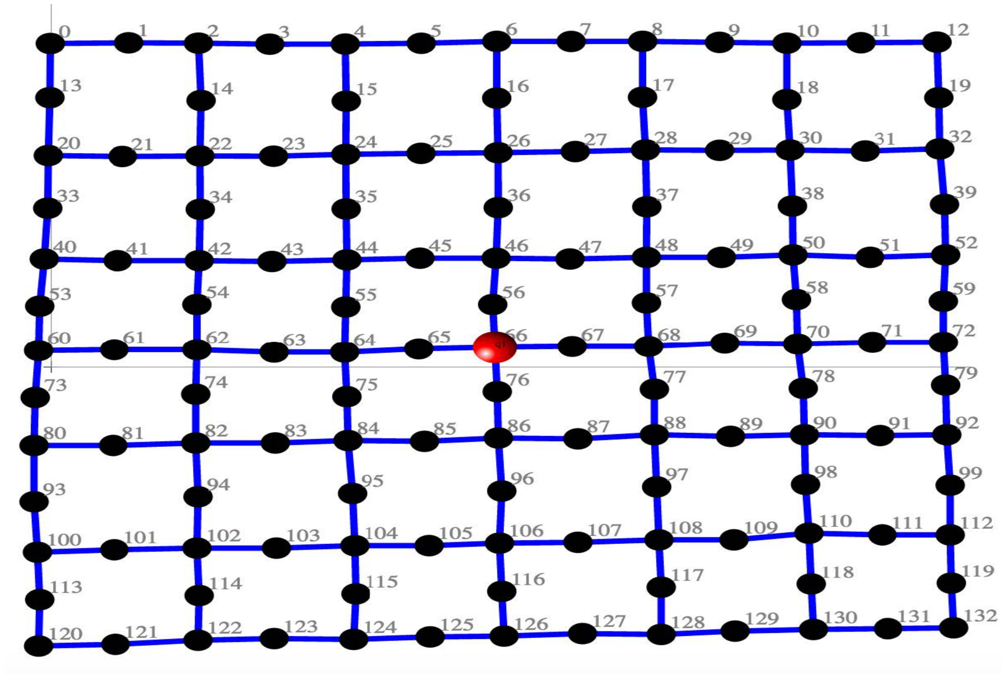

Fig. 2: The guidepath graph used in the numerical experiment discussed in Section IV – reproduced from [30].

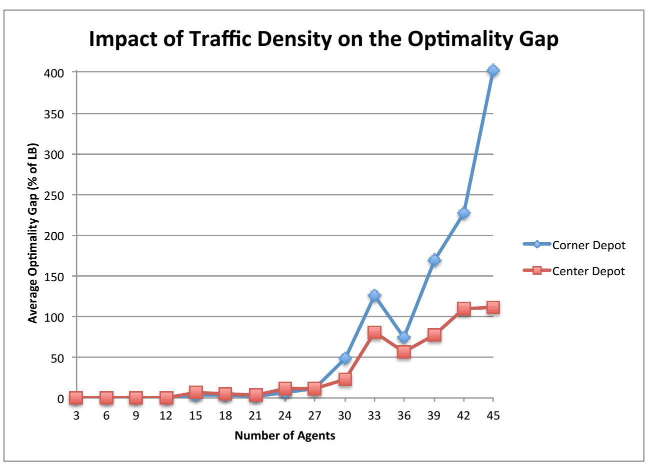

Fig. 3: Plotting the optimality gap for the numerical experiment of Section IV – reproduced from [30].

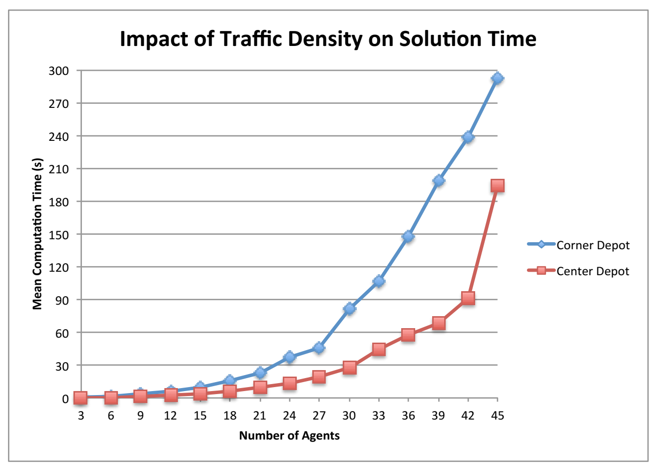

Fig. 4: Plotting the computational times for the numerical experiment of Section IV – reproduced from [30].

problem instance involving n agents subsumed the problem instance that was defined with $n-3$ agents, $n = 6, 9, \ldots, 45$. This structure of the experiment in terms of the number of the traveling agents and their routing specifications intended to assess the performance of the algorithm as the guidepath network became increasingly more congested; indeed, problem instances with 45 agents imply a pretty congested guidepath network, since, at each time period, the traveling agents occupy about 1/3 of the available zones.

We executed five replications of the aforementioned experiment on an HP Z230 workstation with an Intel core i7 processor and 8 GB RAM, running Fedora Linux. The obtained results are summarized in Figures 3 and 4.

Figure 3 reports the optimality gaps for the obtained schedules, averaged across the five replications for each considered problem instance. These optimality gaps were assessed through the solution of the Lagrangian dual problem of the considered problem instances that is studied in [28], [27]; in particular, for each replication, the reported value for the optimality gap was computed by the following formula:

$$\frac{\text{obt. sched. makespan} \; - \; \text{opt. value of Lagr. dual}}{\text{opt. value of Lagr. dual}} \times 100$$

The two plots presented in Figure 3 suggest that the performance of our algorithms is pretty close to the optimal in environments where the zone occupancy by the traveling agents is fairly low, but this performance degrades as the ratio of the number of the traveling agents to the number of zones of the guidepath network increases to some higher levels. Furthermore, the experienced degradation is higher in the case that the agent depot is located away from the "center" of the guidepath network.[3]

Figure 4 plots the algorithm computational times as a function of the traffic density. Again, the reported times are averaged across the five replications of each considered problem instance. As it can be seen in the provided plots, our algorithms execute very fast, even for problem instances that correspond to highly congested environments. Also, the plots of Figure 4 are consistent with the plots of Figure 3 in identifying (i) the traffic density, and (ii) the centrality of the depot location as important factors that determine the difficulty of the considered problem instance.

V. Conclusions

In this paper we have overviewed an ongoing research program concerning the real-time traffic management for a set of agents that travel over a zone-controlled guidepath network, in order to ensure the expediency of the agent trips, but also the safety and the liveness of the resulting motion. We have discussed a number of application contexts that motivate the considered problem, and we have positioned this problem and our technical developments in the relevant literature; along these lines, we have characterized the affinity of the considered problem to some previously studied scheduling problems, but also the new analytical and computational challenges that are defined by it. We have overviewed the pursued methodological framework, and the primary technical developments that have been attained in the context of this framework. In the last section, we also reported some numerical results that demonstrate the practical potency of the presented developments. Finally, we should add at this point that the derived results have been implemented in a high-fidelity simulator for quantum computing that is developed by a research team at the Georgia Tech Research Institute (GTRI).

Our future work will seek to strengthen further the reported developments by improving the currently attained trade-off between the computational efficiency of the developed methods and the efficacy of the derived solutions. It will also seek to embed more systematically the current developments in the MPC framework that was discussed in Section III, and to extend this framework in order to address more complicated traffic dynamics and higher-level decisions, like the agent matching with the emerging transport requests. Finally, we shall also seek the application of the derived results to some "real-life" instantiations of the material handling systems that were discussed in Section I.

Acknowledgement

The second author of this work acknowledges his support by the NSF grant ECCS-1405156.

[3]This effect is explainable by the detailed algorithmic logic that is employed for the construction of the corresponding schedules; the reader is referred to [30], [29] for the relevant details.

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

[4] P. J. G. Ramadge and W. M. Wonham, "The control of discrete event systems," *Proceedings of the IEEE*, vol. 77, pp. 81–98, 1989.

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

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

[7] 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.

[8] 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.

[9] 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.

[10] A. Giua, M. P. Fanti, and C. Seatzu, "Monitor design for colored Petri nets: an application to deadlock prevention in railway networks," *Control Engineering Practice*, vol. 10, pp. 1231–1247, 2006.

[11] 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.

[12] S. A. Reveliotis, "Coordinating Autonomy: Sequential Resource Allocation Systems for Automation," *IEEE Robotics & Automation Magazine*, vol. 22, pp. 77–94, 2015.

[13] S. Reveliotis, "Logical Control of Complex Resource Allocation Systems," *NOW Series on Foundations and Trends in Systems and Control*, vol. 4, pp. 1–223, 2017.

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

[15] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Mineola, NY: Dover, 1998.

[16] T. E. Morton and D. Pentico, *Heuristic Scheduling Systems*. NY, NY: John Wiley, 1993.

[17] J. Adams, E. Balas, and D. Zawack, "The Shifting Bottleneck Procedure for Job Shop Scheduling," *Management Science*, vol. 34, pp. 391–401, 1988.

[18] I. M. Ovacik and R. Uzsoy, *Decomposition Methods for Complex Factory Scheduling Problems*. Boston: Kluwer Academic Publ., 1997.

[19] 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.

[20] 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.

[21] J. Huang, U. S. Palekar, and S. G. Kapoor, "A labeling algorithm for the navigation of automated guided vehicles," *Journ. of Eng. for Industry*, vol. 115, pp. 315–321, 1993.

[22] G. Ghiani, G. Laporte, and R. Musmanno, *Introduction to Logistics Systems Planning and Control*. West Sussex, England: John Wiley & Sons, Ltd., 2004.

[23] E. Gawrilow, E. Köhler, R. H. Möhring, and B. Stenzel, "Dynamic routing of automated guided vehicles in real-time," in *Mathematics – Key Technology for the Future*. Springer, 2008, pp. 165–177.

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

[25] M. L. Fisher, "The Lagrangian relaxation method for solving integer programming problems," *Management Science*, vol. 27, pp. 1–18, 1981.

[26] D. P. Bertsekas, *Nonlinear Programming (2nd ed.)*. Belmont, MA: Athena Scientific, 1999.

[27] 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. –.

[28] ——, "A customized dual ascent algorithm for a class of traffic coordination problems," ISyE, Georgia Institute of Technology (submitted for publication), Tech. Rep., 2016.

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

[30] ——, "Optimized multi-agent routing in guidepath networks," ISyE, Georgia Institute of Technology (submitted for publication), Tech. Rep., 2017.