

Scheduling Automated Traffic on a Network of Roads

Arvind Giridhar and P.R. Kumar

Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory

University of Illinois at Urbana-Champaign

1308 West Main Street

Urbana, IL 61801

E-mail: giridhar@uiuc.edu, prkumar@uiuc.edu

Abstract— We consider the problem of scheduling automated traffic in a city. Such automatic scheduling can improve efficiency of the system by decreasing delays, increasing capacity and easing congestion. The algorithms described in this paper can be thought of as belonging to a layer in the corresponding control-communication-computational system which is responsible for generating timed trajectories for individual vehicles [1].

Each vehicle has a specified route from its origin to its destination, and the task of the scheduler is to provide timed trajectories for all vehicles, which follow the respective vehicles' routes and further ensure that no collisions or deadlock will result. Our approach reduces the problem to a discrete-time graph scheduling problem, by defining an appropriate graph to model the road network. Our main result is a sufficient condition on the graph of the road network and on the initial distribution of vehicles, under which there exists a scheduling algorithm that is guaranteed to clear the system in finite time. The nature of this result allows the design of provably correct scheduling algorithms that require only a small portion of future routes of all vehicles to be known, and are consequently able to work in real time.

We also address the optimization of performance with respect to delay, by focusing on the “one-step move” problem. Finding an optimal solution for the one-step problem would provide a greedy solution of the original network-wide scheduling problem, but is itself NP-hard. We present a polynomial time heuristic algorithm and evaluate its performance through simulations.

Index Terms: Traffic scheduling, smart vehicles, deadlocks, traffic networks, graph scheduling .

I. INTRODUCTION

Future traffic systems could include vehicles which are automatically driven along smart highways and/or city roads [2], [3]. Such automation of traffic movement, along with a control infrastructure for the overall coordination and scheduling of traffic, could improve efficiency, reduce travel time, and prevent traffic deadlocks.

A miniature and highly simplified version of such a futuristic scenario is the vehicular testbed recently constructed in the

Information Technology Convergence Lab at the University of Illinois [4]. This testbed consists of remote controlled vehicles controlled individually by dedicated laptop computers, cameras providing vision feedback, and additional computers which process this feedback and communicate instructions to the individual laptops. The laptops and the servers are connected together to form a wireless network. One of the main tasks required of this distributed system is to automatically schedule movements of the vehicles in such a way that no collision occurs, and no vehicle is ever required to deviate away from its pre-specified route.

Such an automated system features control at several layers [5,6]. At the lowest layer, closed loop control based on the sensor inputs ensures that each vehicle follows a timed trajectory given by a higher layer for a short duration of time. At a higher layer, this timed trajectory is itself generated, based on the layout of the roads and the intersections, the desired routes of all the vehicles, and the positions of all the vehicles. The main component of this layer is a scheduling algorithm, which must be provably correct in that it must ensure that at no time in the future will the system ever reach a deadlock, and must also yield low delays for the vehicles in reaching their destinations. It is this scheduling problem which is the focus of this paper.

The contribution of this paper can be summarized as follows. We construct a simple discrete-space and discrete-time model for the movements of vehicles, which ensures a minimum safe distance between adjacent vehicles. We then derive a class of scheduling algorithms that are provably deadlock free given a certain sufficient condition. Our main result is a sufficient condition on the configuration of vehicles on the road network and the next “step” in the route of each vehicle, that guarantees that deadlock can be avoided and the network cleared, irrespective of future (beyond the one-step) route information. This allows the design of “online” algorithms which only require current and immediate future information to produce schedules which are guaranteed to be loop free. We further address the important issue of optimization with respect to delay, and describe a delay efficient deadlock free algorithm that has been implemented in the testbed described above.

This material is based upon work partially supported by AFOSR under Contract No. F49620-02-1-0217, DARPA under Contract Nos. F33615-01-C-1905 and N00014-01-1-0576, USARO under Contract Nos. DAAD19-00-1-0466 and DAAD19-01010-465, DARPA/AFOSR under Contract No. F49620-02-1-0325, and NSF under Contract Nos. NSF ANI 02-21357 and NSF CCR-0325716.

Corresponding author: Arvind Giridhar, giridhar@uiuc.edu.

Traffic flow modeling has been extensively studied in the transportation literature. There are macroscopic approaches (e.g. [7, 8]) which model traffic flow on the basis of “macro” parameters such as throughput, mean velocity, global density, etc. In contrast, microscopic modeling approaches, such as vehicle following models [9], cell transmission models [10] and cellular automaton models [11–13], propose simple rules for individual vehicle behavior, and proceed to analyze the dynamic behavior of aggregates of such vehicles. Our modeling approach is microscopic in nature, and can be regarded as a simplification of the cellular automaton model. However, our model does not seek to capture the macroscopic behavior of current human driven traffic, but rather to provide an abstraction for the control algorithm.

Research into traffic control has included study of route guidance strategies as well as traffic signal control strategies. Route guidance has been extensively studied under the context of Dynamic Traffic Assignment, which is a component of Advanced Traveler Information Systems and Advanced Traffic Management Systems (ATIS/ATMS). Methodologies used include simulation based approaches [14, 15], optimal control based approaches [16], and optimization based approaches [17]. Traffic signal control strategies would fall in the class of what is referred to here as scheduling policies or algorithms. There has been research into traffic responsive signal control (see [18] for an overview), involving design of strategies to optimize the various signal control parameters such as cycle time, phase split and offset in order to maintain a high degree of saturation in the road network.

Scheduling problems similar to those in this paper have been considered in the context of flexible manufacturing systems (FMS). In such systems, processes concurrently competing for exclusive use of resources, combined with a circular chain of processes in which each process is waiting for the next process in the chain, can lead to deadlock. In this paper, we refer to such an occurrence as an occupied cycle. The problem of *deadlock avoidance*, or scheduling to avoid deadlocks (as opposed to other approaches which seek to prevent deadlocks by restraining request structure of resources, or approaches that seek to detect and recover from deadlock) has been studied in the FMS context [19–22]. A traffic network can be regarded as an instance of a single-unit resource allocation system (SU-RAS) [22], with vehicles being the jobs and discrete sections of streets being the resources. There are however some important distinctions between our model and the FMS model. First, in our traffic model, simultaneous transitions of jobs (vehicles) from one resource to another available resource may be forbidden due to “edge conflicts” in the corresponding graph. This feature does not exist in FMSs. Second, there is a difference in the notion of a deadlock in an FMS and that used here. In a traffic network with automated vehicles and centralized control, we contend that an occupied cycle does *not* imply deadlock; if every vehicle in the circuit moves in synchronized fashion, no collisions will result and each vehicle can move forward. Our notion of deadlock requires an even stronger condition involving conflicting edges. Interestingly, our main result, independently derived here, turns out to be an extension of a similar result derived for SU-RASs [21, 22].

The remainder of the paper is organized as follows. Section II describes the model of the system and elaborates on the problem statement. Section III provides conditions for the existence of a feasible scheduling algorithm. Section IV presents a greedy algorithm for one step in the above scheme and assesses its performance through simulations. Section V addresses the applicability of this approach to actual traffic networks, followed by some possible directions for future work.

II. SYSTEM MODEL AND PROBLEM STATEMENT

We construct a discrete-time discrete-space model of the system. Let each directed lane be divided into *cells*, and discretize time into slots so that at the beginning of each slot, at most one vehicle occupies a single cell. This essentially imposes a minimum safe distance between vehicles in the same lane (assuming each vehicle is situated in the middle of its cell at the start of every time slot), equal to the length of the cell less the length of the vehicle, assuming that the lengths of each cell and vehicle are fixed. Assume further that all vehicles move at a constant speed, or are stationary, and the length of a time slot is chosen such that a vehicle can only traverse from one cell to an adjacent one in a single time slot. For instance, if the vehicle length is c , minimum safe distance d and the common speed v , the length of a cell is $d + c$ and the duration of a single slot is $\frac{d+c}{v}$.

Additionally, the intersection of two streets is not modeled as a cell, but rather consider the cells immediately preceding and succeeding the intersection to be adjacent.

Further, if there are adjacent lanes in the same direction, the lanes can be sectioned so that a lane change operation from a cell in one lane to the “next” cell in the adjacent lane (with next denoting the cell adjoining the upstream neighbor) is permitted. This defines adjacency between cells in parallel lanes going in the same direction. It is important to note that the scheduling requirement of single time-slot cell-to-cell shifting does not require all cells to have exactly the same dimension, or for all vehicles to have the same speed all the time, or indeed for speed to jump from 0 to v instantaneously. For instance, cells adjoining an intersection can be of shorter length, since they are not physically adjacent. Further, this physical non-adjacency means that a vehicle must travel at a higher speed to cross an intersection in a single time slot. The salient fact is that this class of schedules ensures that vehicles in the same lane will be separated by a minimum safe distance.

The road network can be modeled as a directed graph \mathcal{G} by representing each cell by a vertex $v \in \mathcal{V}$, and by representing cell adjacencies by directed edges in a set \mathcal{E} . Edge uv oriented from vertex u to vertex v indicates the adjacency of the corresponding cells, and orientation of the lane joining them. The notation “ $u \rightarrow v$ ” will be used to mean that such an edge exists. Two edges are said to *conflict* if the paths between the corresponding pairs of cells intersect. The graph can thus be represented by a $|\mathcal{V}| \times |\mathcal{E}|$ incidence matrix:

$$I(i, j) = \begin{cases} 1, & \text{if } i \text{ is the tail vertex of edge } j; \\ -1, & \text{if } i \text{ is the head vertex of edge } j; \\ 0, & \text{otherwise,} \end{cases}$$

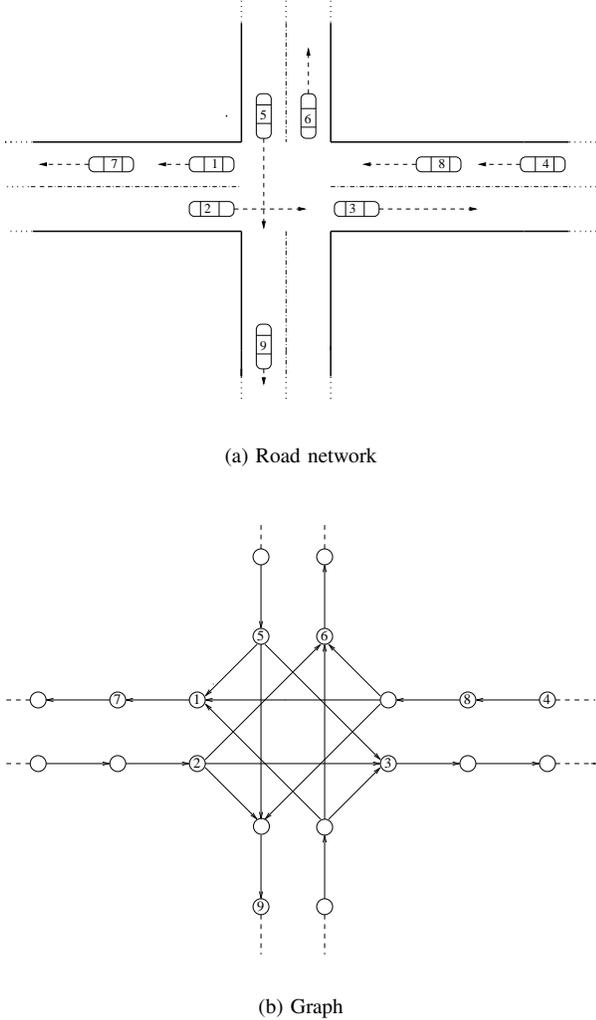


Fig. 1. Graph representation of a four-way intersection. Crossing edges represent edge conflicts

and a symmetric $|\mathcal{E}| \times |\mathcal{E}|$ conflict matrix:

$$C(i, j) = \begin{cases} 1, & \text{if edges } i \text{ and } j \text{ conflict;} \\ 0, & \text{otherwise.} \end{cases}$$

Next, let $[n] := \{1, 2, \dots, n\}$ denote the set of vehicles. The state of the system is described by the $X(t) = [x_1(t), x_2(t), \dots, x_n(t)]$, $t = 1, 2, 3, \dots$ where $x_i(t) \in \mathcal{V}$ is the vertex corresponding to the cell containing vehicle i at time t . Equivalently, vehicle i is said to *occupy* vertex $x_i(t)$ at time t . Figure 1 shows the graph obtained from a four-way intersection.

The evolution of the system over time is thus described by the sequence $\{X(t) : t = 1, 2, 3, \dots\}$. The constraints on the movements of vehicles are as follows:

- 1) Given any time t and vehicle i , either $x_i(t) = x_i(t+1)$, or, if not, $x_i(t) \rightarrow x_i(t+1)$. In other words, a vehicle can only *shift* to an adjacent vertex in accordance with the corresponding edge orientation.
- 2) If $i \neq j$, $x_i(t) \neq x_j(t)$. A vertex can be occupied by only one vehicle.

- 3) Vehicles cannot simultaneously shift across conflicting edges. That is, if $x_i(t) \neq x_i(t+1)$ and $x_j(t) \neq x_j(t+1)$, then edges $x_i(t)x_i(t+1)$ and $x_j(t)x_j(t+1)$ cannot conflict.

Each vehicle i has an associated route R_i , which is a directed walk on the graph \mathcal{G} , along which it is constrained to travel. Each vehicle begins at time 0 at its source and exits the system on reaching its destination. If vehicle i reaches its destination at time t' , $x_i(t) = \phi$ for all $t > t'$, where ϕ is a dummy vertex. Note that the above conditions describing state transitions are valid only for vehicles that have not exited the system.

Thus, the next state $X(t+1)$ is fully specified by the current state $X(t)$, the remaining parts of the routes of all vehicles, and the set of vehicles $S(t) \subseteq [n]$ that shift during $[t, t+1]$. Such a subset is called *feasible* (with feasibility implicitly referencing a particular time instant) if states $X(t)$ and $X(t+1)$ satisfy the feasibility conditions described above, together with the route following constraint.

A *feasible system clearing schedule* S is a sequence of feasible subsets S_0, S_1, \dots, S_T of vehicles to be shifted at each time $0 \leq t \leq T$, such that for all $t' > T$ and each vehicle i , $x_i(t') = \phi$. Thus, a schedule is simply a sequence of sets of vehicles to be moved in each time slot, terminating when all vehicles reach their destination.

The objective of the scheduling problem can now be stated:

Given an initial state X_0 , to determine a feasible schedule S_0, S_1, \dots, S_T for some $T < \infty$ that clears the system, or to show that no such schedule exists.

It must be emphasized that this definition of a feasible schedule, and indeed the discrete model of the system, is a restriction of the most general class of feasible schedules, the latter consisting of any set of continuous time trajectories that, if followed by the corresponding vehicles, ensures that each vehicle will travel from the corresponding source to destination along the given route.

III. EXISTENCE OF FEASIBLE SCHEDULES

We now address the problem of finding conditions under which the system can be feasibly cleared. Naturally, we would like these conditions to be as weak as possible, with minimal assumptions on the underlying graph, as well as on the initial configuration of vehicles.

We define some terms for convenience.

A. Definitions

Given a state $X(t)$:

- An *occupied cycle* (defined implicitly at time t) is a directed cycle that satisfies the following conditions:
 - 1) Each vertex in the cycle is occupied by a vehicle, and
 - 2) If u, v are successive vertices in the cycle with $u \rightarrow v$, then v follows u in the remaining part of the route R_i of the vehicle i occupying u .
- An *occupied path* is a directed path that satisfies condition (2) above, in which the terminating vertex is either unoccupied or belongs to an occupied cycle, and in which

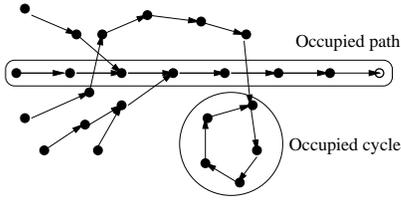


Fig. 2. Occupied path and cycle in a configuration.

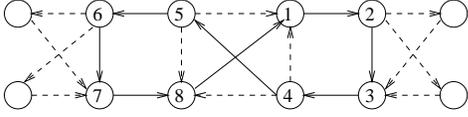


Fig. 3. Deadlock configuration.

every other vertex is occupied. Figure 2 shows instances of an occupied path and an occupied cycle (the edges shown are shift edges).

It is clear from the definitions that any occupied vertex must belong either to an occupied path or to an occupied cycle, and further that the occupied path beginning from that vertex, or the occupied cycle, is uniquely determined. This follows from the fact that the route of each vehicle uniquely specifies the next vertex to be shifted to.

B. Sufficient Conditions for Feasible Scheduling

The system is said to be in *deadlock* at time t if, given the state $X(t)$, there is a non-empty subset of vehicles which cannot be shifted by any future feasible schedule. Figure 3 gives an example of deadlock. The shift edges are solid, and the other edges are shown as dashed lines. The vehicles numbered 1 to 8 occupy a cycle. Clearly, if any one of them is to be shifted, every one must be shifted as well, which is, however, not feasible since two of the shift edges conflict.

It is thus clear that the presence of a cycle with mutually conflicting edges causes the system to be in deadlock. It may seem plausible that the absence of such a configuration at the initial time may be a sufficient condition for the existence of a feasible schedule. However, Figure 4 shows that that is not the case. Note that only the shift edges are shown. In Figure 4(a), the vertices occupied by vehicles 1, 2, 3, 4, 5, 6, 7, and 8 form an occupied cycle. Shifting any one of these vehicles forces every one of them to be shifted, resulting (due to the routes of each of these vehicles) in Figure 4(b). Here we have a self-conflicting occupied cycle consisting of the vertices occupied by 1, 2, 9, 10, 11, 12, 13, 6, 7, and 8. So, in fact, the configuration in Figure 4(a) is already in deadlock, since vehicles 9, 10, 11, 12, and 13 cannot be shifted.

However, a slightly stronger condition is sufficient to guarantee absence of deadlock, as well as existence of a feasible schedule. The following result is an extension of Theorem 2 in [21] (re-derived in [22]) to our model.

Theorem 1: If the directed graph \mathcal{G} is such that each vertex has either in-degree 1 or out-degree 1, and the system has no occupied cycle at time 0, then there exists a schedule that clears the system in finite time.

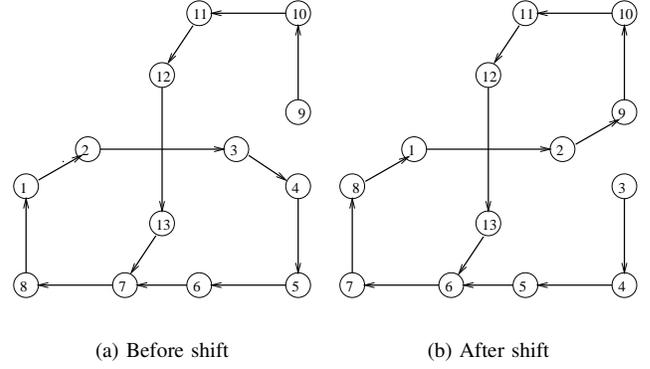


Fig. 4. Occupied cycle leading to deadlock.

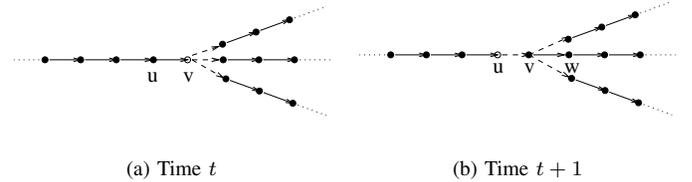


Fig. 5. Case 1.

Proof. It is enough to prove the following: If the system has no occupied cycle at time t , there is a feasible set S_t of size 1 which results in a state $X(t+1)$ having no occupied cycle. If this is proved, it follows that at each step it is possible to move at least one vehicle by one step along its route. Since the number of vehicles and the route lengths are finite, it follows that the system will eventually be cleared.

Consider any vertex, and the maximal occupied path beginning at that vertex. Such a path must end at a vertex u , and, since there is no occupied cycle, the vertex v following u in the route of the vehicle occupying u must be unoccupied.

The vertex v must either have in-degree or out-degree equal to 1, by assumption. Consider first the case in which the in-degree of v is 1 (Figure 5).

Let $S_t = \{i : x_i(t) = u\}$. We now claim that there is no occupied cycle at time $t+1$. Since there is no occupied cycle at time t , and the only vertex that goes from unoccupied to occupied at time $t+1$ is v , any occupied cycle must include v . Since the in-degree of v is one, the vertex preceding v , i.e., u , must also belong to this occupied cycle. However, u is unoccupied at time $t+1$, so this is not possible.

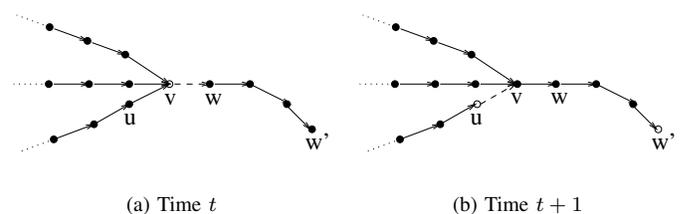


Fig. 6. Case 2.

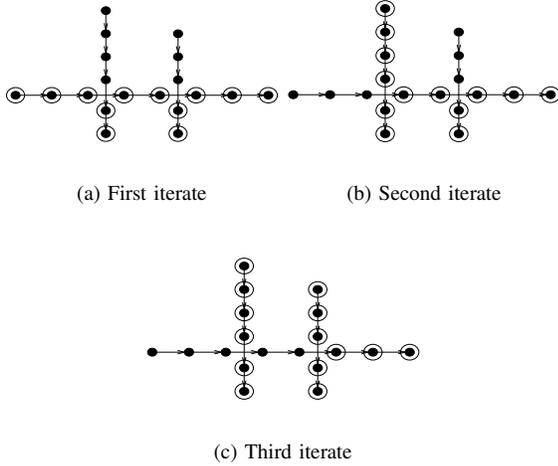


Fig. 8. Application of heuristic algorithm to example configuration.

shift edges that conflict with the added path, and deleting the occupied paths starting with these vertices.

S_1^m is said to be an m -neighbor of S_2 if it is a maximal feasible subset (i.e., it is not contained in any feasible subset), and it contains a neighbor of S_2 . Observe that the notions of neighbor and m -neighbor are asymmetric; if S_2 is a (m)-neighbor of S_1 , it need not be that S_1 is a (m)-neighbor of S_2 .

In light of the above definitions, we slightly modify the above description of the algorithm by replacing feasible subsets by maximal feasible subsets, and neighbors by m -neighbors. In other words, we begin with a maximal feasible subset, and successively replace the current subset by a larger m -neighbor, stopping when we reach a local optimum.

The algorithm works as follows: Construct an initial set S_{init} by setting up a list of non-self-conflicting occupied paths in decreasing order of length, and sequentially picking out the path at the top of the list, adding it to S_{init} if the addition preserves feasibility, or else discarding it. The procedure ends when the list is empty. This gives a maximal feasible set S_{init} .

Given the current set S , pick any path in \bar{S} (the complement of S) and add it to S . Delete the elements of the original set that conflict with the chosen path. Using the list of paths in \bar{S} , repeat the procedure described in the previous paragraph to obtain an m -neighbor of S , say S' . If $|S'| > |S|$, then set $S = S'$ and begin a new iteration. If not, repeat the same operation with another path in S .

Each step in the algorithm results in a strictly larger feasible subset. The algorithm terminates when all the paths in S are checked without obtaining a larger m -neighbor subset S' .

We illustrate the working of the algorithm by applying it to an example configuration; see Figure 8. The vertices shown are occupied vertices, and the edges are the shift edges. Figure 8(a) shows the result of the initial greedy covering (the selected occupied vertices are circled). After two steps (Figure 8(b) and (c)), the algorithm achieves the local optimum, which in this case is also the global optimum.

It is not hard to obtain a crude polynomial bound (polynomial in the number of vehicles) on the running time of the

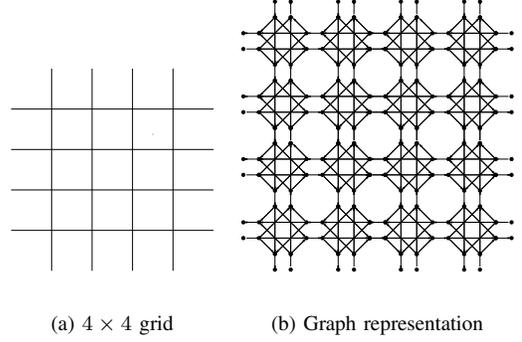


Fig. 9. Test graph for simulation.

algorithm. Let n be the number of vehicles. Checking whether a subset is feasible has complexity $O(n^2)$. The number of occupied paths is $O(n)$. To obtain the initial set S_{init} consequently takes $O(n^3)$ time. There are $O(n)$ neighbors of a feasible subset. Comparing with the set of neighbors takes $O(n^3)$ time, and this process is done $O(n)$ times (since the set S must grow strictly larger each time). The algorithm's running time is therefore bounded by $O(n^4)$. This is the time complexity per step. There are at most $O(R_{max}n)$ steps in all, where R_{max} is the length of the longest route.

V. SIMULATION RESULTS

Given an arbitrary configuration, in the worst case, only a single vehicle can be shifted in each move. Since the algorithm does not exercise any control on future configurations, apart from ensuring that they will contain no occupied cycle, we cannot provide a theoretical lower bound on the performance, except for the worst-case bound $O(R_{max}n)$ of moving one vehicle-section per time slot.

We have therefore tested through simulations the performance of the above heuristic algorithm and a greedy-optimal algorithm that chooses at each step the largest feasible subset. The greedy-optimal algorithm uses essentially an exponential-time brute force approach (with some slight refinements that reduce the exponent by a factor) at each time step, going through all possible subsets to find the largest feasible subset. This algorithm has exponential complexity.

In the example simulated, the graph is obtained from a road network grid (Figure 9). Each line in the grid is a two-way street. Random pairs of vertices are picked as sources and destinations, and the routes are randomly selected from among the minimum length paths between each source and destination pair. Figure 9 shows a 4×4 grid with two vertices per square edge.

The algorithm's performance is measured by the ratio

$$\text{delay_ratio}(n) = \frac{\text{schedule_sum}}{\text{route_sum}},$$

where schedule_sum is the total time for each vehicle to reach its destination (according to the schedule obtained by the algorithm) summed over all vehicles, route_sum is the sum over all vehicles of the route lengths, and n is the number of vehicles in the system. Clearly, $\frac{\text{route_sum}}{n}$ is a lower bound

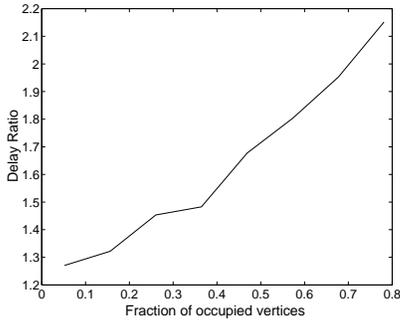
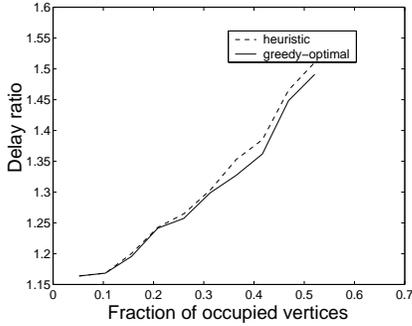
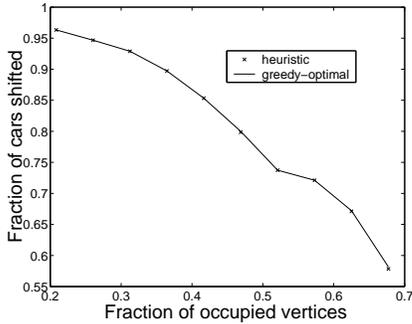


Fig. 10. Performance of the heuristic algorithm.



(a) Delay ratio comparison



(b) One-step comparison

Fig. 11. Comparisons of greedy-optimal and heuristic algorithms.

on the minimum average time per vehicle to clear the system. For example, a ratio of 2 implies that on average a vehicle is shifted in as many time slots as it is held stationary. These values are averaged over multiple random placements.

Figure 10 shows the plot of $delay_ratio(n)$ versus fraction of occupied vertices, for a 3×3 grid with 2 vertices per square edge (each way). The graph has 96 vertices. We are not able to show a comparison with the greedy optimal, since the exponential complexity of the latter makes it infeasible to calculate for a large number of vehicles. It shows that even when the density of vehicles per vertex approaches 1, $delay_ratio(n)$ does not increase significantly.

Figure 11(a) shows the comparison of the two algorithms for a smaller graph, a 2×2 grid with 4 vertices per square edge (a

96 vertex graph). Figure 11(b) shows the comparison in terms of number of vehicles moved in a single step. The results indicate that the heuristic algorithm closely, or exactly in the case of the one-step simulation, matches the performance of the greedy-optimal algorithm.

Note that these results do not tell the story for large dense configurations, in which the exponential complexity of the greedy-optimal algorithm makes it infeasible to run, and thus does not allow such comparisons. Also note that none of the plots reflect algorithm performance relative to the global optimum. We have no approximation or bound for the latter, and exhaustively searching over all possibilities is computationally infeasible even for small configurations.

VI. GENERALIZATIONS, REAL WORLD CONSIDERATIONS, AND FUTURE WORK

Some extensions and relaxations of assumptions in our model can be made without much difficulty. For example, the assumption that no new vehicles can enter the system can be relaxed to allow vehicles to enter the system at arbitrary times as long as the occupancy constraints are not violated, and no occupied cycles are introduced. The assumption that the routes are fixed and known beforehand can also be relaxed due to the online nature of Theorem 1: Routes can be changed at any time as long as occupied cycles are not introduced.

Some important elements, however, are absent in this model. These include multiple vehicle speeds, the cost attached to stopping and starting, and the possible requirement of different minimum distances depending on speed.

The class of allowable schedules can be expanded to include some of these elements. For instance, the requirement that the speeds of all vehicles are the same could be relaxed to allow speeds to be multiples of a minimum speed. The system could then be discretized according to this minimum speed, and multiple vertices would be allowed to be traversed in a single time slot. Also, a cost can be attached to stopping and starting, and the resulting optimization problem of picking which vehicles are allowed to shift would include the costs as weights. To guarantee different safe distances depending on speed, a rule enforcing multiple empty slots in between vehicles could be introduced. The mapping of the road into sections would then be according to the minimum safe distance at the minimum non-zero speed. Analyzing such a model could be a possible future direction to investigate.

With the likelihood of graphs of very large size, the computational complexity of any algorithm becomes crucial. A greedy approach in this class of algorithms results in an NP-complete problem. We present instead a polynomial time heuristic algorithm. The simulations show that in most of the randomly generated instances on the graph obtained from a grid road network, the heuristic performs almost as well as the greedy optimal solution. Finding lower order polynomial algorithms is definitely worth looking into. Another issue concerning scalability which arises is that providing scheduling instructions individually to all vehicles could be a hard problem from a communication point of view. It would be worth investigating whether such algorithms could be implemented in decentralized fashion.

On the theoretical side, given the assumption of arbitrary fixed routes, we have not been able to find useful lower bounds on the clearing time of the algorithm (i.e., an algorithm with a useful lower bound or approximation factor). It might be possible to obtain stronger results if another degree of freedom is added, i.e., if the routes are allowed to be variable.

Indeed, there are a number of ways to generalize the problem definition, and some of them may be worth investigating for possible savings achievable through more flexible constraints.

REFERENCES

- [1] S. Graham, G. Baliga, and P. Kumar, "Issues in the convergence of control with communication and computing: Proliferation, architecture, design, services, and middleware," in *Proceedings of CDC '04*, Nassau, Bahamas, December 2004.
- [2] P. Varaiya, "Smart cars on smart roads: Problems of control," *IEEE Trans. Aut. Con.*, vol. 38, no. 2, pp. 195–207, February 1993.
- [3] S. Shladover, C. Desoer, J. Hedrick, M. Tomizuka, J. Walrand, W. Zhang, D. McMahon, H. Peng, S. Sheikholeslam, and N. McKeown, "Automated vehicle control developments in the PATH program," *IEEE Trans. Vehic. Tech.*, vol. 40, pp. 114–130, February 1991.
- [4] "The I.T. Convergence Lab, University of Illinois." [Online]. Available: <http://black1.csl.uiuc.edu/~prkumar>
- [5] S. Graham and P. Kumar, "The convergence of control, communication and computation," in *Proceedings of PWC 2003: Personal Wireless Communication, Venice, Italy*, 2003.
- [6] G. Baliga, S. Graham, L. Sha, and P. Kumar, "Etherware: Domainware for wireless control networks," in *Proceedings of the 7th IEEE Symposium on Object-oriented Real-time Distributed Computing*, Vienna, Austria, 2004.
- [7] G. Newell, "A simplified theory of kinematic waves in highway traffic, part 1: General theory," *Transpn. Res. B*, vol. 27, 1993.
- [8] B. Kerner and P. Kornhauser, "Structure and parameters of cluster in traffic flow," *Phys. Rev. E*, vol. 50, pp. 54–83, 1994.
- [9] G. Newell, "A simplified car-following theory: A lower order model," *Transpn. Res. B*, vol. 36, pp. 195–205, 2002.
- [10] C. Daganzo, "The cell transmission model, part 2: Network traffic," *Transpn. Res. B*, vol. 29B, no. 2, pp. 79–93, 1995.
- [11] K. Nagel and M. Schreckenberg, "A cellular automaton model for freeway traffic," *J. Phys. I*, vol. 2, pp. 2221–2229, 1992.
- [12] K. Nagel, "Particle hopping models and traffic flow theory," *Physical Review E*, vol. 3, pp. 4655–4672, 1996.
- [13] A. Schadschneider and M. Schreckenberg, "Cellular automaton models and traffic flow," *J. Phys. A*, vol. 26, pp. 679–683, 1993.
- [14] H. Mahamassani, T. Hu, S. Peeta, and A. Ziliaskopoulos, "Development and testing of dynamic traffic assignment and simulation procedures for ATIS/ATMS applications," Center for Transportation Research, University of Texas at Austin, Tech. Rep., 1994.
- [15] M. Ben-Akiva, M. Bierlaire, H. Koutsopoulos, and R. Mishalani, "DynaMIT: a simulation based system for traffic prediction," in *DACCORD Short Term Forecasting Workshop*, 1998.
- [16] T. Friesz, F. Luque, R. Tobin, and B. Wie, "Dynamic network traffic assignment considered as a continuous time optimal control problem," *Operations Research*, vol. 37, pp. 893–901, 1989.
- [17] D. Merchant and G. Nemhauser, "A model and an algorithm for the dynamic traffic assignment problem," *Transpn. Sc.*, vol. 12, pp. 183–199, 1978.
- [18] M. Bell, "Future directions in traffic signal control," *Transpn Res. A*, vol. 26, pp. 303–313, 1992.
- [19] Z. Banaszak and B. Krogh, "Deadlock avoidance in flexible manufacturing systems with concurrently competing process flows," *Ieee Trans. Rob. and Aut.*, vol. 6, no. 6, pp. 724–734, 1990.
- [20] R. Wysk, N. Yang, and S. Joshi, "Detection of deadlocks in flexible manufacturing cells," *IEEE Trans. Rob. and Aut.*, vol. 7, no. 6, pp. 853–859, 1995.
- [21] M. Fanti, B. Maione, and B. Turchiano, "Event-based feedback control for deadlock avoidance in flexible production systems," *IEEE Trans. Rob. and Aut.*, vol. 13, pp. 347–363, 1997.
- [22] M. Lawley and S. Reveliotis, "Deadlock avoidance for sequential resource allocation systems: Hard and easy cases," *In. J. Flex. Man. Sys.*, vol. 13, pp. 385–404, 2001.
- [23] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Murray Hill, NJ: W.H. Freeman and Co, 1979.