2005 American Control Conference
June 8-10, 2005. Portland, OR, USA

On the Time Complexity of Conflict-Free Vehicle Routing

Michael Savchenko and Emilio Frazzoli

***Abstract*— In this paper, we study the following problem: given n vehicles and origin-destination pairs in the plane, what is the minimum time needed to transfer each vehicle from its origin to its destination, avoiding conflicts with other vehicles? The environment is free of obstacles, and a conflict occurs when the distance between any two vehicles is smaller than a velocity-dependent safety distance. We derive lower and upper bounds on the time needed to complete the transfer, in the case in which the origin and destination points can be chosen arbitrarily, proving that the transfer takes $\Theta(\sqrt{n}\bar{L})$ time to complete, where $\bar{L}$ is the average distance between origins and destinations. We also analyze the case in which origin and destination points are generated randomly according to a uniform distribution, and present an algorithm providing a constructive upper bound on the time needed for complete the transfer, proving that in the random case the transfer requires $O(\sqrt{n \log n})$ time.**

I. INTRODUCTION

Problems involving the coordinated motion of several mobile agents in a shared environment are ubiquitous, appearing in many safety-critical application domains, such as automated highways, air traffic control, and automated factories. As a consequence, the design of algorithms for the resolution of such problems has attracted a great interest in the recent years, especially in view of the increasing role of autonomous decision making in the development of complex and information-rich man-made systems [1].

Many multiple-vehicle coordination algorithms have been proposed by researchers from robotics, computer science, systems and control, and optimization. While a thorough review of the literature is out of the scope of this paper, we mention a few recently-proposed approaches in the literature. Centralized approaches in robotics solve a motion planning problem in an extended configuration space obtained as the Cartesian product of the individual robots' configuration spaces [2]. Prioritization-based schemes design a feasible trajectory for a robot considering higher-priority robots as time-varying obstacles [3]. Coordination solutions based on Nash equilibria were proposed in [4]. In order to reduce the complexity of the problem, the motion of the robots is often restricted on a graph (roadmap), as in [5]; Pareto-optimal solution for coordination of robots moving on a roadmap have been characterized in [6], [7]. A market-based approach for conflict negotiation has been proposed in [8]. Semi-definite relaxations of the coordination problems are considered in [9]. In [10], decentralized optimization schemes are used to ensure safe coordination of aircraft.

In spite of the richness of the literature on the subject, it still remains unclear what the fundamental limits in terms of achievable performance are for multiple-robot systems; in particular, it remains unknown how the efficiency of the coordinated motion scales with the number of robots. As a consequence, it is difficult to objectively characterize the effectiveness of the algorithms available in the literature in solving coordination problems, especially when the number of agents is very large and an optimal solution cannot be practically computed. As a matter of fact, the importance of the development of a model of computation for distributed robotic systems has been emphasized in a number of recent papers. For example, [11] introduces a notion of communication complexity, in the context of the development of a control and communication language for robotic networks; in [12], time and communication complexity of a number of coordination tasks are investigated.

The objective of this paper is to make a further contribution in this direction. We consider a particular class of multiple-vehicle coordination problem, which we call Conflict-Free Vehicle Routing. Informally, we wish to investigate the minimum time needed to safely transfer n mobile agents between pre-assigned origin and destination points, avoiding conflicts. A conflict is generated when two agents get closer than a velocity-dependent safety distance. As will be discussed, the problem is particularly interesting when agents are allowed to get arbitrarily close, provided that their speed is low enough. The main results of the paper include: (i) a characterization of the minimum time needed to solve certain instances of the problem, and (ii) algorithms for conflict-free motion coordination, with an asymptotic analysis of their performance as a function of the number of vehicles. In this paper, we use the expression *time complexity* to indicate the time needed to complete the physical transfer of all the n robots from their origin to their destination points. This is not necessarily related to the *algorithmic complexity* of the problem, i.e., the number of operations or memory space needed to compute a set of conflict-free trajectories between origin and destination points.

This paper builds on recent results obtained by Gupta and Kumar [13] on the capacity of wireless networks. Perhaps surprisingly, the insight gained through the analysis of wireless networks can be used to yield novel results

This material is based upon work partially supported by NSF under grants CCR-0133869 and CCR-0325716. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organization.

Michael Savchenko is with the Department of Aerospace Engineering at the University of Illinois at Urbana-Champaign, Urbana, IL 61801, savchenk@uiuc.edu.

Emilio Frazzoli is with the Department of Mechanical and Aerospace Engineering at the University of California at Los Angeles, Los Angeles, CA 90095, frazzoli@ucla.edu.

0-7803-9098-9/05/$25.00 ©2005 AACC

in terms of the time complexity of a class of motion coordination problem. In fact, the capacity-limiting factor in wireless networks is interference, occurring when nearby network nodes attempt to broadcast information at the same time; its effects can be reduced by appropriately tuning the transmission schedule and the broadcast power. Similarly, in our model, congestion is caused by vehicles moving too close to one another, and generating a conflict; collisions can be avoided by careful maneuvering, which in our model requires a speed reduction. The main, substantial difference with the wireless network case can be found in the effects of the agent mobility on the topology of the vehicle network.

The paper is organized as follows. In Section II we introduce some notation, formulate and motivate the problem we wish to address. In Section III we provide lower and upper bounds on the time complexity of the coordination problem for the case in which the origin-destination pairs can be chosen arbitrarily, thus showing that the time complexity of the problem is $\Theta(\sqrt{n}\bar{L})$. In Section IV, we provide an upper bound on the time complexity of the motion coordination problem in the case in which origin-destination pairs are chosen randomly from a uniform distribution, thus proving that the problem has time complexity $O(\sqrt{n \log n})$. Finally, in Section V we draw some conclusions and discuss directions of future research.

II. Problem formulation and motivation

Let the environment $\mathcal{Q}$ be a convex, compact region in the plane. Without loss of generality, we shall assume that $\mathcal{Q}$ is a square of unit area. Consider n pairs of points in $\mathcal{Q}$, to which we will refer to as origin-destination (O, D) pairs, with $(O, D)_i \in \mathcal{Q} \times \mathcal{Q}$, $i \in \{1, \ldots, n\}$. The i-th (O, D) pair is assigned to a mobile agent $\mathcal{A}_i$, $i \in \{1, \ldots, n\}$. We will refer to an (O, D) pair and the agent assigned to it interchangeably.

Initially, each agent is *inactive*, i.e., is not considered to be in the environment, and cannot be involved in a conflict (defined in the following). Let $t_{0,i} \geq 0$ be the time at which the i-th agent is activated and enters the environment at location O_i; the agent reverts to the inactive state upon arrival at its destination D_i, at time $t_{0,i} + T_i$. While active, the i-th agent moves within $\mathcal{Q}$ along a continuous, time-parameterized path $\gamma_i : [0, T_i] \to \mathcal{Q}$. The position of an agent as a function of time is given by the function $x_i : t \mapsto x_i(t) = \gamma_i(t - t_{0,i})$; by convention, we will set $x_i(t) = \gamma_i(0) = O_i$ for $t < t_{0,i}$, and $x_i(t) = \gamma_i(T_i) = D_i$ for $t > t_{0,i} + T_i$. Finally, let $v_i(t)$ be the velocity of agent i at time t, and assume its magnitude is bounded by $v_{\max} > 0$. The active/inactive status of agents can be thought of as representing the fact that origin and destination locations are "safe havens" in which vehicles are removed from the environment shared with traffic, and safety is guaranteed; for example, these might represent parking spots for automobiles, and airports for aircraft.

For each active agent, we define an *exclusion region* C, modeled as a disk centered at the agent's position, and with radius depending on the agent's velocity, i.e.,

$$C_i(t) = \{z \in \mathbb{R}^2 : |z - x_i(t)| < r_0 + k|v_i(t)|\}, \quad (1)$$

for given constants $r_0, k \geq 0$, $r_0 k > 0$; $|\cdot|$ represents the Euclidean norm. We say that a *conflict* occurs between agents $\mathcal{A}_i$ and $\mathcal{A}_j$ if there exists a time t_c such that:

- Both $\mathcal{A}_i$ and $\mathcal{A}_j$ are active at time t_c, and
- $C_i(t_c) \cap C_j(t_c) \neq \emptyset$.

The motivation for a velocity-dependent exclusion region can be found, for example, in the need to ensure safety in the presence of position and velocity uncertainties, and delays in sensing or communication of the state of other agents. We will henceforth use the word "agent" to refer interchangeably to the mobile vehicle and to its exclusion region.

A conflict-free routing policy is a map $\pi : (O, D) \mapsto (t_0, T, \gamma)$ that, given a set of (O, D) pairs, assigns to each agent an activation schedule, and a time-parameterized path. We say that the policy π is *safe* if it does not generate conflicts. Let us indicate with $T_\pi(O, D)$ the time at which the last agent is deactivated according to policy π; we will define the *time complexity* of the Conflict-Free Vehicle Routing Problem (CF-VRP) defined by the (O, D) pairs as the infimum of this time over all safe policies, i.e.,

$$T^*(O, D) = \inf_{\pi \text{ safe}} T_\pi(O, D).$$

In the remainder of the paper, we will aim at establishing asymptotic bounds on the time complexity of the problem, as $n \to \infty$ (for a review of asymptotic notation, as well as big-O notation, see a standard text such as [14]).

A. Some preliminary results

We have the following trivial bound:

Proposition 2.1: For any set of n origin-destination pairs, the time complexity of the conflict-free vehicle routing problem is $O(n)$.

Proof: If agents are activated sequentially, i.e., agent $\mathcal{A}_{i+1}$ is activated upon deactivation of agent $\mathcal{A}_i$, no conflicts can arise, and the time needed for the i-th agent to reach its destination can be bounded as $T_i \leq \operatorname{diam}\mathcal{Q}/v_{\max}$. Hence, the time at which the last agent arrives at its destination is

$$T_{\text{seq}}(O, D) = \sum_{i=1}^{n} T_i \leq n \frac{\operatorname{diam}\mathcal{Q}}{v_{\max}},$$

which proves the claim. ■

Let us consider now the case in which $r_0 > 0$ in the definition of the exclusion region; we have the following:

Theorem 2.2: Suppose that $r_0 > 0$ in (1). For any set of n origin-destination pairs, with $\min_i |O_i - D_i| \geq l > 0$, the time complexity of the conflict-free vehicle routing problem is $\Theta(n)$.

Proof: Let us restrict our attention to the case $r_0 \leq 1$. Since each active agent reclaims a region of area at least $\pi r_0^2/4$, at most $n_d = \lfloor 4/(\pi r_0^2) \rfloor$ agents can be active at the same time. At most n_d new agents can be activated no sooner than every $l/v_{\max}$ time units, i.e., the minimum time needed for at least one agent to reach its destination. Hence the time needed to activate all agents can be bounded as

$$T_{r_0>0} \geq n \frac{l}{\lfloor 4/(\pi r_0^2) \rfloor v_{\max}},$$

which, together with Proposition 2.1, proves the claim. ■

This result condemns routing problems with agents of non-vanishing size to linear time complexity, that is, not significantly better than what can be achieved via sequential agent activation, when n becomes large. In the remainder of the paper, we will study the case in which $r_0 = 0$ in the definition of the no-conflict constraints, i.e., the case in which the radius of the exclusion region is directly proportional to the agent's velocity[1].

While this is not—strictly speaking—a physically realistic modeling assumption, one must keep in mind that, in most problems of interest, conflicts are generated when vehicles get closer than some safety distance that is much bigger than the physical dimensions of the vehicle. For example, in air traffic control, a conflict is generated whenever two aircraft come to within 5 nautical miles from each other. A common rule of thumb for defensive driving in automotive traffic requires to maintain a two-second buffer from a leading car; at 50 km/h, this corresponds to roughly 30 meters, several times the length of a typical car.

In addition, setting $r_0 = 0$ lets us study the effect of velocity on traffic congestion. The intuition is that as agents move faster, they need a bigger buffer to avoid collisions with others, hence reclaiming a larger portion of a shared resource (the environment), and thereby imposing severe constraints on the motion of other agents.

III. ARBITRARY (O,D) PAIRS

In this section we will compute upper and lower bounds on time complexity for a routing problem in which the (O, D) pairs can be chosen arbitrarily. Note that "arbitrarily" here must be understood as "in such a way that the time complexity is minimized." In other words, this section provides *best-case* bounds for the coordination problem. We also note that with arbitrary (O, D) pairs, the average distance between origins and destinations is an arbitrary variable directly affecting the time complexity. The relevant measure for the arbitrary case is therefore not completion time, but time per unit distance traveled by each agent. This is equivalent to describing time complexity in terms of both n and the average distance between origins and destinations.

A. A lower bound on the time complexity

In the arbitrary case, we have the following lower bound.

Lemma 3.1: For any set of n (O, D) pairs, the time complexity of the conflict-free motion coordination problem is $\Omega(\sqrt{n}\bar{L})$, where $\bar{L}$ is the average distance between origin and destination points.

Proof: Let us assume that the motion of all the agents can be represented as a set of straight-line motions, over a common, synchronized time schedule of length h. For simplicity, let us assume that each time interval has the common duration τ. By definition, all agents reach their destination within time $T^* = h\tau$. Let us indicate with r_i^j the length of the straight-line segment along which the i-th agent moves during the j-th time interval. Obviously, we have $\sum_{j=1}^{h} r_i^j \geq L_i$, where $L_i = |O_i - D_i|$, and

$$\sum_{i=1}^{n} \sum_{j=1}^{h} r_i^j \geq n\bar{L}. \tag{2}$$

Defining $\delta_i^j = kv_i^j = kr_i^j/\tau$, the area of the set spanned by the exclusion region of the i-th agent during the j-th time interval can be computed as

$$A_i^j = \pi(\delta_i^j)^2 + 2\delta_i^j r_i^j = \left(\frac{r_i^j}{\tau}\right)^2 k(\pi k + 2\tau).$$

Since at each time instant, at least one fourth of each exclusion region is within $\mathcal{Q}$ (chosen to be a unit square), the sum of the areas of the regions claimed by all agents at each time interval is bounded as

$$\sum_{i=1}^{n} A_i^j = \frac{k(\pi k + 2\tau)}{\tau^2} \sum_{i=1}^{n} \left(r_i^j\right)^2 \leq 4.$$

Summing over all intervals in the time schedule, and rearranging, we get

$$\sum_{i=1}^{n} \sum_{j=1}^{h} \left(r_i^j\right)^2 \leq \frac{4h\tau^2}{k(\pi k + 2\tau)}. \tag{3}$$

Consider a convex function $f : \mathbb{R} \to \mathbb{R}$; Jensen's inequality states that

$$f\left(\frac{1}{P}\sum_{p=1}^{P} x_p\right) \leq \frac{1}{P}\sum_{p=1}^{P} f(x_p).$$

Since the function $x \mapsto x^2$ is convex, we can apply Jensen's inequality to (3) to obtain

$$\left(\sum_{i=1}^{n} \sum_{j=1}^{h} r_i^j\right)^2 \leq hn \sum_{i=1}^{n} \sum_{j=1}^{h} \left(r_i^j\right)^2 \leq \frac{4h^2\tau^2 n}{k(\pi k + 2\tau)},$$

that is,

$$\sum_{i=1}^{n} \sum_{j=1}^{h} r_i^j \leq 2\sqrt{\frac{(T^*)^2 n}{k(\pi k + 2\tau)}}.$$

Thus, from (2) we get $T^* \geq \frac{1}{2}\sqrt{k(\pi k + 2\tau)n}\bar{L}$. In the limit as $\tau \to 0$, i.e., for continuous schedules, we get $T^* \geq \frac{1}{2}\sqrt{\pi k^2 n}\bar{L}$, which proves the result. ■

B. A constructive upper bound

In this section, we present an algorithm for choosing (O, D) pairs in order to achieve the same time complexity that appears in the lower bound.

[1]Similar considerations hold when $r_0 = O(1/\sqrt{n})$, but we will not pursue this direction in this paper for the sake of conciseness.

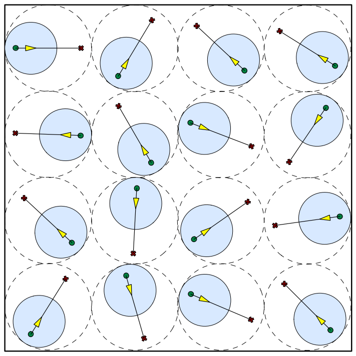

Fig. 1. Example of a selection of origin-destination pairs achieving the upper bound in Lemma 3.2. Circles and crosses represent origins and destinations, respectively, and the shaded disks are the exclusion regions.

Lemma 3.2: For any integer m there exists a set of $n = m^2$ (O, D) pairs, with average distance between origin and destination points equal to $\bar{L}$, such that the time complexity of the conflict-free motion coordination problem is $O(\sqrt{n}\bar{L})$.

Proof: Consider a disk $\mathcal{R}$ of radius $r < kv_{\max}$, and place a single pair of origin and destination points at the two points on a diameter at a distance ηr from the center ($\eta < 1$). As the agent moves along the diameter, its maximum velocity is constrained in such a way that its exclusion region is entirely contained within $\mathcal{R}$, that is,

$$\frac{d\xi(t)}{dt} \leq \frac{r}{k}\left(1 - \frac{|\xi(t)|}{r}\right),$$

where ξ is the distance of the agent from the center of $\mathcal{R}$ as a function of time. Given this constraint, the time needed to travel from the origin to the destination can be computed as

$$T = -2k\log(1-\eta). \quad (4)$$

Pack $n = m^2$ disjoint disks of radius r within $\mathcal{Q}$. This is possible with $r = 1/2\sqrt{n}$. Place one (O, D) pair on the diameter of each disk, in such a way that each point is at a distance ηr from the corresponding center. Move each agent inside its disk according to the same schedule considered above for the single-agent case. (See Figure 1 for a sketch.)

The time at which all agents will reach their destinations is then the same, given by (4). Noting that $\bar{L} = 2\eta r = \eta/\sqrt{n}$, and multiplying and dividing (4) by $\bar{L}$, we get our result:

$$T_{\text{arb}} = \frac{-2k\log(1-\eta)}{\eta}\sqrt{n}\bar{L}. \qquad \blacksquare$$

Let $\mathcal{OD}(n, \bar{L})$ be the set of all sets of n (O, D) pairs with average distance $\bar{L}$; the combination of Lemmas 3.1 and 3.2 proves the following:

Theorem 3.3: The time complexity of the CF-VRP satisfies

$$\inf_{(O,D)\in\mathcal{OD}(n,\bar{L})} T^*(O, D) = \Theta(\sqrt{n}).$$

We remark again that this is a result on the *best-case* time complexity: for all choices of (O, D) pairs in a given class, the time complexity is $\Omega(\sqrt{n})$, and there exists at least one choice of (O, D) pairs in the same class with time complexity $O(\sqrt{n})$.

IV. Random (O,D) pairs

In this section we derive an upper bound on the time complexity for the CF-VRP, for the case of randomly distributed origin and destination points. Such problems can arise in large-scale decentralized applications, in which many independent agents travel across a region with a large number of candidate origin or destination points. Specifically, suppose that the (O, D) points are chosen from a uniform distribution in $\mathcal{Q}$, identically and independently (*i.i.d.*). We will demonstrate the following bound:

Theorem 4.1: For any set of n (O, D) pairs, randomly chosen from a uniform distribution in $\mathcal{Q}$, the time complexity of the CF-VRP is $O(\sqrt{n \log n})$ with high probability.

We will use "with high probability" (*whp*) to mean with probability approaching 1 as $n \to \infty$. In order to prove this bound, we present an algorithm which achieves it *whp*. The scheme we present is based on a partitioning $\mathcal{P}_n$ of $\mathcal{Q}$ with a grid of identical square cells of area $4r^2(n)$, where $r(n)$ is chosen appropriately. We assume a time partition $\mathcal{T}$ into intervals of length τ. During each time interval, each agent will either remain in its current cell or transfer to a neighboring cell along its path. Thus the motion coordination problem is decoupled into *intra*cell and *inter*cell coordination, which we consider in order.

A. Intracell coordination scheme

Consider the following "nested orbiting" coordination scheme within a square cell of area $4r^2(n)$. Each agent in a cell is assigned to an orbit (i.e., an annular region); the first (outermost) orbit is contained between circles of radius $r_0 = r$ and $r_1 = \eta r$, for a given $\eta \in (0, 1)$. At most two agents are assigned to a single orbit. The j-th orbit is contained between circles of radius $r_{j-1} = \eta^{j-1} r$ and $r_j = \eta^j r = \eta r_{j-1}$ (i.e., the orbits are logarithmically spaced). An agent assigned to the j-th orbit can move along its central circle at speed up to $v_j = \omega r_j$, with

$$\omega = \frac{1-\eta}{2\eta k},$$

without interfering with agents in nearby orbits; in particular, agents can reposition themselves to anywhere in their assigned orbit in no more than $\tau_{\text{a}} := \pi/\omega$ time. An arbitrary number of agents in neighboring orbits can move radially at speed proportional to their instantaneous distance from the center of the cell without conflicts; all agents will reach the center of the next innermost (or outermost) orbit in at most $\tau_{\text{r}} := -\ln(\eta)/\omega$ time.

Hence, the choice of logarithmically-spaced orbits has the following important consequences:

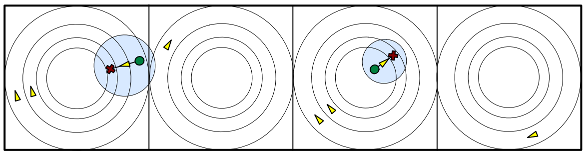

Fig. 2. Sketch of the initialization algorithm, on a row of $\mathcal{P}_n$: note that non-neighboring cells are initialized concurrently, which prevents conflicts between moving agents.

1) All agents can move along the central circle of their orbit at the same angular speed, without interfering with agents on different orbits or agents in different cells. Angular reconfiguration will not take more than a constant time τ_{a}.
2) Agents can move to a higher or lower orbit within constant time τ_{c}.
3) The operations of adding a new agent (pushing all other agents one orbit down), or removing an agent ("popping" the outermost agent, and moving the remaining agents one orbit up) can be performed in constant time τ_{c}.
4) An arbitrary number of agents can be "stored" in a single cell.
5) At least two agents can share a single orbit.

In other words, the nested orbiting scheme allows us to treat each cell as a buffer with infinite capacity, and with constant-cost insertion and extraction operations.

The agents are reconfigured from the origin points to a nested configuration (with any arbitrary assignment of agents to orbits) during an *Initialization* phase. The time necessary to initialize a single cell can be bounded as follows.

Lemma 4.2: Let N be the number of agents with origin point within a square cell. For any assignment of agents to orbits, there exists an algorithm to safely activate and reconfigure the agents in a nested orbit configuration in time $O(N^2)$.

Proof: In order to prove the statement, we present an algorithm that achieves the claimed performance; see Figure 2 for a sketch. At each iteration, the algorithm activates the agent assigned to the outermost empty orbit, and moves it to a position within that orbit, without generating conflicts with agents in the same cell. Consider the j-th iteration, $1 \leq j \leq N$, and let O_j be the origin of the agent assigned to the j-th orbit. All agents active at the beginning of this iteration are in orbits lie outside of the j-th orbit. Before activation of the new agent, all the already active agents are moved to a point in their orbit opposite to O_j with respect to the center of the cell: this takes constant time. Moving on the radius through O_j, the new agent will be able to reach the j-th orbit in time at most $k \log(2\eta^{1-j}/(\eta+1)+1) = O(N)$, without generating conflicts with active agents in the same cell. Since there are at most N agents in a cell, it will take $O(N^2)$ time to initialize nested orbits in it. ■

The initialization algorithm cannot be executed concurrently at all cells, since conflicts may be generated between agents in neighboring cells. A straightforward coloring argument shows that a schedule for concurrent initialization of non-neighboring cells adds only a constant factor to the time needed to initialize all cells, i.e.,

Proposition 4.3: Let N be the maximum number of agents with origin point within a single cell in the partition $\mathcal{P}_n$. The initialization of all cells in $\mathcal{P}_n$ takes $O(N^2)$ time.

B. Intercell routing

At the end of the initialization phase all agents are arranged in an array of nested orbits, with each agent in the cell containing its origin point. A two-phase intercell routing algorithm is proposed in the following: in the first phase agents are moved to the column of $\mathcal{P}_n$ containing their destination; in the second phase they are delivered to the destination cell. In other words, we will consider two consecutive routing problems on a linear array of cells with infinite-capacity buffers.

1) Routing in a linear array: We have the following result, which is a simple extension of a widely-known fact in parallel computing, i.e. that the farthest-first routing strategy is time-optimal in linear arrays [15]. In the farthest-first strategy, if one or more packets need to leave a node, conflict is resolved giving precedence to packets that have to travel the farthest to reach their destination.

Theorem 4.4 (Adapted from [15]): Consider a linear array of l nodes with infinite buffers. Each node is able to transfer one packet to each neighbor in a single time step. If each node is the source and destination of at most N packets, all packets can be delivered using the farthest-first strategy within $\max\{l, Nl/4\} + O(\sqrt{Nl \log l})$ time steps, with high probability.

2) Intercell coordination scheme: Without loss of generality, let us assume that in the first phase only horizontal cell transfers are allowed. Choose the following criteria for assigning agents to orbits in the initialization phase: (i) Assign agents traveling "right" to even-numbered orbits, and agents traveling "left" to odd-numbered orbits; (ii) sort agents in order of distance to their destination column, placing on the outermost orbit the agent that has to go the farthest distance. From such a configuration, the agents can be reconfigured in constant time to a configuration in which all agents moving right are directly above the center of the cell, and all agents moving left are below the center of the cell. From such a configuration, at each time step all agents in the outermost orbit can move to next cell, with no conflicts between agents moving in opposite directions (see Figure 3). It can be shown that an agent will never stop moving until it reaches its destination cell; upon reaching the destination cell, the agent moves directly to the lowest empty orbit. The length of the time step is a constant, i.e., $\tau = 3r_1/v_1 = 4/(1-\eta)$. Theorem 4.4, and the fact that the

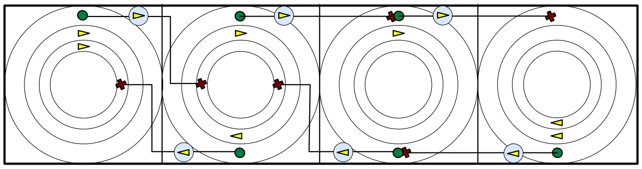

Fig. 3. Sketch of the routing algorithm, on a row of $\mathcal{P}_n$: note that once an agent starts moving between cells, it does not stop until it reaches its destination cell, at which point it moves to the lowest empty orbit.

number of cells in a row is $l = 1/(2r)$, show that the column of the destination cell is reached within $O(N/r)$ time. In the second phase, an analogous procedure is followed, allowing motions in the vertical direction.

C. Termination phase

At the end of the two intercell routing phases, all agents are in the cell containing their destination point. All agents are moved within their cell to their destination using an algorithm analogous to the initialization phase; the termination phase requires the same time as its initialization counterpart, i.e., $O(N^2)$ time.

D. Time complexity bound

Choose the dimension of the cells as follows:

$$r(n) = \sqrt{\frac{3}{4}\frac{\log n}{n}}.$$

In other words, the area of each cell is chosen as $A(n) = 3\log n/n$; we ignore edge effects (i.e. a nonintegral number of cells) since their significance vanishes as $n \to \infty$. We use the following result adapted from [16];

Proposition 4.5: Any cell in $\mathcal{P}_n$ contains no more than $3e\log n$ origin/destination points, almost surely.

In other words, $N \leq 3e\log n$ almost surely. We can finally prove the main result of this section.

Proof of Theorem 4.1: The result is proven by combining the algorithms for the distinct phases outlined above, along with the bound on N, the maximum number of origin/destination points in each cell. The initialization and termination phase will take $O((\log n)^2)$ each *whp*, where the horizontal and vertical intercell routing phases will take $O(\sqrt{n\log n})$ time *whp*. The combination of all these phases will require $O(\sqrt{n\log n})$ time with high probability. ■

V. Conclusions

In this paper, we have studied the time complexity of classes of conflict-free motion coordination problems, with conflicts defined by the intersection of velocity-dependent exclusion regions. We first showed that if the area of the exclusion region is bounded away from zero, the time complexity of the coordination problem is $\Theta(n)$, i.e., is no better than the trivial worst-case bound. We then focused on the case in which the exclusion region can be made arbitrarily small by reducing the agent's velocity, and showed that for the case in which origin and destination pairs can be chosen arbitrarily, the time complexity of the motion coordination problem is $\Theta(\sqrt{n})$. In the case of random origin-destination pairs, we showed that the time complexity is $O(\sqrt{n\log n})$. A tighter upper bound for the random case is now available, proving that the time complexity is $\Theta(\sqrt{n})$ [17]. Future directions include the convergence of our results with those available in the wireless communications community to characterize the capacity of mobile wireless networks under realistic constraints on the agents' motion.

Acknowledgments

The authors gratefully acknowledge the role of P. R. Kumar in inspiring this work, and thank him for insightful comments and discussions on the subject of this paper.

References

[1] R. M. Murray, Ed., *Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics, and Systems*. SIAM, 2003.

[2] J. Schwartz and M. Sharir, "On the piano movers' problem: III. coordinating the motion of several independent bodies." *International Journal of Robotics Research*, vol. 2, no. 3, pp. 97–140, 1983.

[3] M. Erdmann and T. Lozano-Pérez, "On multiple moving objects," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 1986, pp. 1419–1424.

[4] S. M. LaValle and S. A. Hutchinson, "Optimal motion planning for multiple robots having independent goals," *IEEE Trans. on Robotics and Automation*, vol. 14, no. 6, pp. 912–925, 1998.

[5] S. Akella and S. Hutchinson, "Coordinating the motions of multiple robots with specified trajectories," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2002, pp. 624–631.

[6] H. Chitsaz, J. M. O'Kane, and S. M. LaValle, "Exact pareto-optimal coordination of two translating polygonal robots on an acyclic roadmap," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2004.

[7] R. Ghrist, J. O'Kane, and S. M. LaValle, "Pareto-optimal coordination on roadmaps," in *Proc. Workshop on Algorithmic Foundation of Robotics*, 2004.

[8] B. P. Gerkey and M. J. Matarič, "Sold!: Auction methods for multi-robot coordination," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002.

[9] E. Frazzoli, Z. H. Mao, J. H. Oh, and E. Feron, "Aircraft conflict resolution via semi-definite programming," *AIAA J. of Guidance, Control, and Dynamics*, vol. 24, no. 1, pp. 79–86, 2001.

[10] G. Inalhan, D. M. Stipanovic, and C. J. Tomlin, "Decentralized optimization, with application to multiple aircraft coordination," in *cdc*, Las Vegas, NV, 2002.

[11] E. Klavins, "Communication complexity of multi-robot systems," in *Proc. Fifth International Workshop on the Algorithmic Foundations of Robotics*, Nice, France, 2002.

[12] S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli. (2005, jan) Synchronous robotic networks and complexity of control and communication laws. Short version submitted to the IEEE Conf. on Decision and Control. 2005. [Online]. Available: http://arxiv.org/pdf/math.OC/0501499

[13] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. on Information Theory*, vol. 46, no. 2, pp. 388–404, 2000.

[14] D. Papadimitriou, *Computational complexity*. Addison Wesley, 1993.

[15] M. Kaufmann, S. Rajasekaran, and J. Sibeyn, "Matching the bisection bound for routing and sorting on the mesh," in *Proc. ACM Symposium on Parallel Algorithms and Architectures*, San Diego, CA, 1992, pp. 31–40.

[16] S. R. Kulkarni and P. Vishwanath, "A deterministic approach to throughput scaling in wireless networks," *IEEE Trans. on Information Theory*, vol. 50, no. 6, pp. 1041–1049, June 2004.

[17] V. Sharma, E. Frazzoli, and P. Voulgaris, "On the time complexity of the multiple-vehicle coordination problem with random origin-destination pairs," University of Illinois at Urbana-Champaign, Technical Report AE 04-05 UILU ENG-04-0505, 2004.