

Safe and Stabilizing Distributed Multi-Path Cellular Flows

Taylor T. Johnson*, Sayan Mitra

Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Abstract

We study the problem of distributed traffic control in the partitioned plane, where the movement of all entities (robots, vehicles, etc.) within each partition (cell) is coupled. Establishing liveness in such systems is challenging, but such analysis will be necessary to apply such distributed traffic control algorithms in applications like coordinating robot swarms and the intelligent highway system. We present a formal model of a distributed traffic control protocol that guarantees minimum separation between entities, even as some cells fail. Once new failures cease occurring, in the case of a single target, the protocol is guaranteed to self-stabilize and the entities with feasible paths to the target cell make progress towards it. For multiple targets, failures may cause deadlocks in the system, so we identify a class of non-deadlocking failures where all entities are able to make progress to their respective targets. The algorithm relies on two general principles: temporary blocking for maintenance of safety and local geographical routing for guaranteeing progress. Our assertional proofs may serve as a template for the analysis of other distributed traffic control protocols. We present simulation results that provide estimates of throughput as a function of entity velocity, safety separation, single-target path complexity, failure-recovery rates, and multi-target path complexity.

Keywords: distributed systems, swarm robotics, formal methods, traffic control

1. Introduction

Highway and air traffic flows are nonlinear switched dynamical systems that give rise to complex phenomena such as abrupt phase transitions from fast to sluggish flow [1, 2, 3]. Our ability to monitor, predict, and avoid such phenomena can have a significant impact on the reliability and capacity of physical traffic networks. Traditional traffic protocols, such as those implemented for air

*Corresponding author
Email addresses: `taylor.johnson@acm.org` (Taylor T. Johnson), `mitras@illinois.edu` (Sayan Mitra)

traffic control are *centralized* [4]—a *coordinator* periodically collects information from the vehicles, decides and disseminates waypoints, and subsequently the vehicles try to blindly follow a path to the waypoint. Wireless vehicular networks [5, 6, 7, 8] and autonomous vehicles [9, 10] present new opportunities for *distributed* traffic monitoring [11, 12, 13] and control [14, 15, 16, 17, 18, 19]. While these protocols may still rely on some centralized coordination, they should scale and be less vulnerable to failures compared to their centralized counterparts. In this paper, we propose a fault-tolerant distributed traffic control protocol, formally model it, and formally prove its correctness.

A *traffic control protocol* is a set of rules that determines the routing and movement of certain physical *entities*, such as vehicles, robots, or packages, over an underlying *graph*, such as a road network, air-traffic network, or warehouse conveyor system. Any traffic control protocol should guarantee: (a) (*safety*) that the entities always maintain some minimum physical separation, and (b) (*progress*) that the entities eventually arrive at a given a destination (or target) vertex. In a distributed traffic control protocol, each entity determines its own next-waypoint, or each vertex in the underlying graph determines the next-waypoints for the entities in an appropriately defined neighborhood.

In this paper, we study the problem of distributed traffic control in a partitioned plane where the motions of entities within a partition are *coupled*. The problem can be described as follows (refer to Figures 1 and 2). The *environment*—the geographical space of interest—is partitioned into regions or *cells*. Each entity is assigned a certain type or *color*. For each color, there is one *source cell* and one *target cell* of the same color. The source cells produce entities of some color, and the target cells only consume entities of a particular color, so the goal is to move entities of color c to the target of color c . The motion of all entities within a cell are coupled, in the sense that they all either move identically, or they all remain stationary (we discuss the motivation for this below). If some entities within some cell i touch the boundary of a neighboring cell j , those entities are transferred to j . Thus, the role of the distributed traffic control protocol is to control the motion of the cells so that the entities (a) always have the required safe separation, and (b) reach their respective targets, when feasible.

The coupling mentioned above that requires entities within a cell to move identically may appear strong at first sight. After all, under low traffic conditions, individual drivers control the movement of their cars within a particular region of the highway, somewhat independently of the other drivers in that region. However, on highways under high-traffic, high-velocity conditions, it is known that coupling may emerge spontaneously, causing the vehicles to form a fixed lattice structure and move with near-zero relative speed [1, 20]. In other scenarios, coupling arises because passive entities are moved around by active cells. For example, this occurs with packages being routed on a grid of multi-directional conveyors [21, 22], and molecules moving on a medium according to some controlled chemical gradient. Finally, even where the entities are active and cells are not, the entities can cooperate to emulate a virtual active cell expressly for the purposes of distributed coordination. This idea has been ex-

plored for mobile robot coordination in [23] using a cooperation strategy called *virtual stationary automata* [24, 25].

In this paper, we present a distributed traffic control protocol that guarantees *safety at all times*, even when some cells fail permanently by crashing. The protocol also guarantees *eventual progress* of entities toward their targets, provided (a) that there exists a path through non-faulty cells to the entities' respective targets, and (b) failures have not introduced unrecoverable deadlocks. Specifically, the protocol is *self-stabilizing* [26, 27], in that once new failures stop occurring, the composed system automatically returns to a state from which progress can be made. The algorithm relies on the following four mechanisms.

- (a) There is a *routing* rule to maintain local routing tables to each target at each non-faulty cell. This routing protocol is self-stabilizing and allows our protocol to tolerate crash failures of cells.
- (b) There is a *mutual exclusion* and scheduling mechanism to ensure moving entities over distinctly colored overlapping paths do not introduce deadlocks. The locking and scheduling mechanism ensures one-way traffic can make progress over shared routes (traffic intersections).
- (c) There is a *signaling* rule between neighbors that guarantees safety while preventing deadlocks. Roughly speaking, the signaling mechanism at some cell fairly chooses among its neighboring cells that contain entities, indicating if it is safe for one of these cells to apply a movement in the direction of the cell doing the signaling. This permission-to-move policy turns out to be necessary, because movement of neighboring cells may otherwise result in a violation of safety in the signaling cell, if entity transfers occur.
- (d) The *movement* policy causes all entities on a cell to either move with the same constant velocity in the direction of their destination, or remain stationary to ensure safety. This policy abstracts more complex motion modeling.

We establish these safety and progress properties through systematic assertional reasoning. For safety properties, we establish inductive invariants and for stabilization we use global ranking functions. To show that all entities reach their destinations (when feasible), we use a combination of ranking functions and fairness-based reasoning on infinite executions. These proof techniques may serve as a template for the analysis of other distributed traffic control protocols. Our analysis is generally independent of the size of the environment, number of cells, and number of entities. Additionally, only neighboring cells communicate with one another and the communication topology is fixed (aside from failures). For these reasons, this problem can serve as a case study in automatic parameterized verification of distributed cyber-physical systems [28, 29, 30, 31].

We present simulation results that illustrate the influence (or the lack thereof) of several factors on throughput. (a) Throughput decreases exponentially with

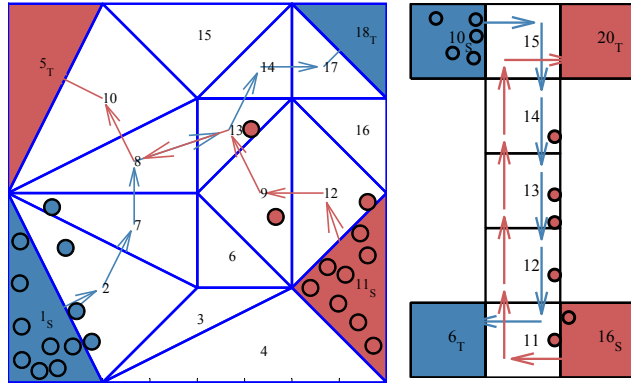


Figure 1: Source cells (1 and 11) produce entities that flow toward the target cell (18 and 5) of the appropriate color. Source-to-target paths overlap at cells 8 and 13. In this execution, the blue entity on cell 7 is waiting for the red entities to leave the overlapping cells.

Figure 2: If cell 10 moved its blue entities onto the shared one-lane “bridge” (11, 12, 13, 14, 15), then all entities would be deadlocked.

path length until saturation, as which point it decreases roughly linearly with path length. (b) Throughput decreases roughly linearly with required safety separation and cell velocity. (c) Throughput decreases roughly exponentially until it saturates as a function of path complexity measured in number of turns along a path. (d) Throughput decreases roughly exponentially with failure rate, and increases linearly with recovery rates, under a model where crash failures are not permanent and cells may recover from crashing. (e) Throughput decreases roughly exponentially until it saturates as a function of the percentage of overlapping cells between different colored targets.

Contributions over Previous Work. In previous work [32], we analyzed a similar problem, but have significantly generalized our results in this paper.

- (A) We consider general tessellations (including triangulations) that define the partitioning, while we considered uniform square partitions in [32]. We also present results on partitioning schemes that cannot work for our formulation of the problem.
- (B) We allow entities of multiple colors, each flowing to a different target, while in [32], we only allowed entities of one color, all of which flowed to the nearest target. This generalization lets source-to-target paths of different colors overlap, creating intersections, and requires several changes to the algorithm, including adding a mutual exclusion and scheduling mechanisms used to control traffic intersections. This generalization is significant because it makes the problem applicable to a much wider class of systems.

- (C) We extended our simulation results to allow for these generalizations, and characterized the cost on throughput due to the extra coordination required to allow multiple colors.

Paper Organization. The rest of the paper is organized as follows. First, Section 2 introduces the model of the physical system. Next in Section 3, we present the distributed traffic control algorithm. Then in Section 4, we define and prove the safety and progress properties. Subsection 4.1 establishes safety. Subsequently, we establish a progress property that shows entities eventually reach their targets in spite of failures (when possible). First in Subsection 4.2, it is shown that the routing protocol to find any target from any cell with a physical path through non-faulty cells to that target is self-stabilizing. Then in Subsection 4.3, we show how overlapping paths to different targets (traffic intersections) can be scheduled. Finally, in Subsection 4.4, it is shown that entities on any cell with a feasible physical path to their target eventually reach their target. Simulation results and interpretation are presented in Section 5, followed by a brief discussion of related work and further extensions, and a conclusion in Sections 6, 7, and 8.

2. Physical System Model

We describe the physical system in this section. For a set K , we define $K_{\perp} \triangleq K \cup \{\perp\}$ and $K_{\infty} \triangleq K \cup \{\infty\}$. For $N \in \mathbb{N}$, let $[N] \triangleq \{1, \dots, N\}$. The $\|\cdot\|$ brackets are used for the Euclidean norm of a vector.

Partitioning. The system consists of N convex polygonal *cells* partitioning a polygonal environment. Let $ID \triangleq [N]$ be the set of unique identifiers for all cells in the system. The planar environment Env is some given simply connected polygon. A partition P of Env is a set of closed, convex polygonal cells $\{P_i\}_{i \in ID}$ such that:

- (a) the interiors of the cells are pairwise disjoint,
- (b) the union of the cells is the original polygonal environment, and
- (c) cells only touch one another at a point or along an entire side.

The first two conditions are the standard definition of a partition, while the third restricts any cell from being adjacent along one of its sides to more than one other cell. Thus, cell i occupies a convex polygon P_i in the Euclidean plane. The boundary of cell i is denoted by ∂P_i . We denote the vertices (extreme points) of P_i as V_i . We denote the number of sides of P_i as $ns(i)$. Let $Side(i, j) \triangleq \partial P_i \cap \partial P_j$ be the common side of adjacent cells i and j —we will refer to $Side(i, j)$ as both an index and a line segment (set of points).

Communications. Cell i is said to be a *neighbor* of cell j if the boundaries of the cells share a common side. The set of identifiers of all neighbors of cell i is denoted by $Nbrs_i$. This definition of neighbors can naturally be represented as a graph, so let Δ be the worst-case diameter of such a neighbor communication graph¹. For each cell $i \in ID$ and each neighboring cell $j \in Nbrs_i$, let the *side normal vector* from i to j , denoted $n(i, j)$, be the unit vector orthogonal to $Side(i, j)$ and pointing into cell j from the common side $Side(i, j)$.

Each cell is controlled by software that implements the distributed traffic control algorithm described in the next section. We consider synchronous protocols that operate in rounds. At each round, each cell exchanges messages bearing state information with its neighbors. Then, each cell updates its software state and decides the (possibly zero) velocity with which to move any entities on it. Until the beginning of the next round, the cells continue to operate according to this velocity, which may lead to entity transfers.

Entities. Each cell may contain a number of *entities*. Each entity occupies a circular area and represents a physical object (or overapproximation of) such as an aircraft, car, robot, or package. Every entity that may ever be in the system has a unique identifier drawn from an index set I . This assumption is for presentation only, and the algorithm does not rely on knowing entity identifiers. For an entity $p \in I$, we denote the coordinates of its center by $\bar{p} \triangleq (p_x, p_y) \in \mathbb{R}^2$.

The open circular area (disc) centered at \bar{p} of radius r representing entity p is denoted $B(p, r)$. The radius of an entity is l , and r_s is the minimum required inter-entity safety gap. We define the total safety spacing radius as $d \triangleq r_s + l$. For simplicity of presentation, we work with uniform entity radii l and safety gaps r_s . If they differ, we would take l and r_s to be the maximums over all entities. We instantiate $B(p, l)$, which represents the physical space occupied by entity p , and we also instantiate $B(p, d)$, which is entity p 's total safety area.

Entity Colors, Source Cells, and Target Cells. There are $|C|$ types (or *colors*) of entities, where C is some finite, ordered set. The color of some entity $p \in I$ is denoted as $color(p)$. For each $c \in C$, there is a *source cell* sid_c and a *target cell* tid_c . All other cells are *ordinary cells*. For simplicity of presentation, we assume there is a unique source and target, but the algorithms and the results generalize for when sid_c and tid_c are sets.

Entity p 's color $color(p)$ designates the target cell entity p should eventually reach. The source sid_c produces entities of color c and the target tid_c consumes entities of color c . The sets of target and source identifiers are denoted $ID_T \subseteq ID$ and $ID_S \subseteq ID$, respectively.

Entity Movement. All the entities within a cell move identically—either they remain stationary, or they move with some constant velocity $0 < v < l$ in the

¹The diameter of this graph is not static, it may change due to failures, but the worst case is always a path graph, so $\Delta = N$.

direction of one of the sides of the cell. Here v is the maximum cell velocity, or the greatest distance traveled by any entity over one synchronous round. We require $v > 0$ to ensure progress. We require that $v < l$ to ensure entities do not collide when transfers occur. The actual cell velocity may differ in each cell so long as each is upper bounded by v . This movement is determined by the algorithm controlling each cell. When a moving entity touches a side of a cell, it is instantaneously transferred to the neighboring cell beyond that side, so that the entity is entirely contained in the new cell.

Safety and Transfer Regions. The safety region on side s of a cell is the area within the cell where (the centers of) new entities entering the cell from side s can be placed. For a side s of some cell i , the *safety region on side s* $SR_i(s)$ is the area on P_i at most $3d$ distance measured orthogonally from side s . Analogously, the transfer region on side s of a cell is the area within a cell where (the centers of) entities reside when those entities will be transferred to the neighboring cell on that side. The *transfer region on side s* $TR_i(s)$ is the region in the partition P_i at most l distance measured orthogonally from side s . For a cell i , the *transfer region* TR_i and *safety region* SR_i are respectively the unions of $TR_i(s)$ and $SR_i(s)$ for each side s of P_i . We refer to the *inner side(s)* of TR_i , $TR_i(s)$, SR_i , or $SR_i(s)$, as the side(s) touching the inside of the annulus, and denote them by ITR_i , $ITR_i(s)$, etc.

For example, in Figure 3, the transfer region for the square cell 3 is the square annulus between the smaller cyan square and the larger blue square (the boundary ∂P_3 of cell 3). Similarly, for the triangular cell 1 in Figure 3, the transfer region is the triangular annulus between the smaller cyan triangle and the larger blue triangle. Thus, the distance measured orthogonally between the sides of the cyan polygons representing the boundary of the transfer region, and the sides of the blue polygons is always l . In Figure 3, for the square cell 3, the safety region is the square annulus between the smaller red square and the larger blue square.

Geometric Assumptions. We assume that the polygonal environment Env and its partition P have shapes and sizes such that each cell in the partition is large enough for an entity to lie completely on it. Particularly, we require for each cell $i \in ID$ that the transfer region TR_i is nonempty. We also assume the following assumptions to ensure transferring entities between cells is well-defined.

Assumption 1. (*Projection Property*): For each $i \in ID$, for each side s of P_i , there exists a constant vector field over P_i that drives every point in P_i to some point on side s without exiting P_i .

By definition, the cells form a partition. However, partially because there is “empty space” between the transfer regions of the cells, the transfer regions do not form a partition. Even if we remove this empty space by translating the transfer regions so the sides of transfer regions of neighboring cells coincide, they still may not form a partition (see Figure 3 for an example where the transfer regions cannot form a partition). This is because, for the shared side s

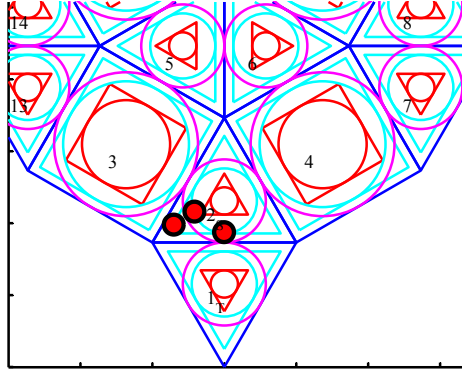


Figure 3: Safety regions (areas between red and blue) and transfer regions (areas between cyan and blue) for the squares and triangles composing the snub square tiling tessellation. The magenta circles are incircles and are used to compute the transfer and safety circles and regions.

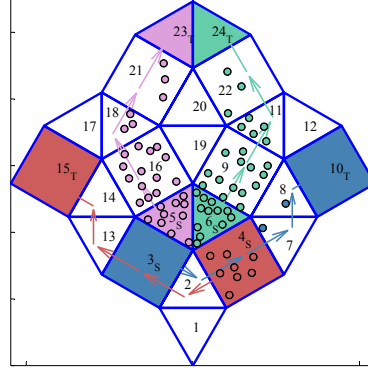


Figure 4: Blue and red paths overlap at cells 2, 3, and 4. Blue entities on cells 7 and 8 have traversed the intersection and then the red source (4) produces entities. Red and blue sources producing entities simultaneously would cause a deadlock.

of neighboring cells i and j , the inner sides of the transfer regions on P_i and P_j may have different lengths, even though the shared side s obviously had the same length for P_i and P_j .

Assumption 2. (*Transfer Feasibility*): For any $i \in ID$ and any $j \in Nbrs_i$, consider their common side $Side(i, j)$. The length of the inner side $ITR_i(Side(i, j))$ line segment equals the length of the inner side $ITR_j(Side(i, j))$ line segment.

3. Distributed Traffic Control Algorithm

Next, we describe the *discrete transition system* $Cell_i$ that specifies the software controlling an individual cell P_i of the partition P .

Preliminaries. A *variable* is a name with an associate type. For a variable x , its type is denoted by $type(x)$ and it is the set of values that x can take. A *valuation* (or state) for a set of variables X is denoted by \mathbf{x} , and is a function that maps each $x \in X$ to a point in $type(x)$. Given a valuation \mathbf{x} for X , the valuation for a particular variable $v \in X$, denoted by $\mathbf{x}.v$, is the restriction of \mathbf{x} to $\{v\}$. The set of all possible valuations of X is denoted by $val(X)$. Many variables return cell identifiers that we use to access variables of other cells using subscripts, and if the valuation of these variables are restricted to the same state, we will drop the particular state on the subscripted variables for more concise notation. For instance, suppose $\mathbf{x}.next_i \in ID$, then $\mathbf{x}.next_{\mathbf{x}.next_i}$ would be written $\mathbf{x}.next_{next_i}$.

A *discrete transition system* \mathcal{A} is a tuple $\langle X, Q_0, A, \rightarrow \rangle$, where:

- (i) X is a set of variables and $val(X)$ is called the set of *states*,

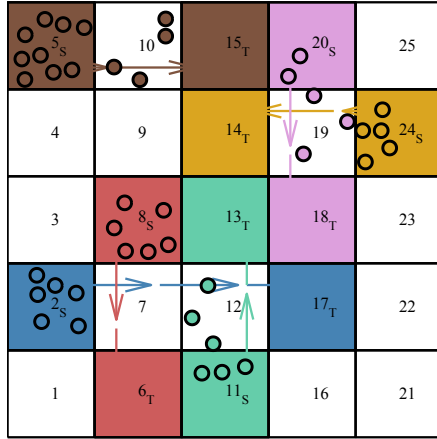


Figure 5: Example illustrating the computation of the color-shared cells and shared colors, stored in the $pint[c]$ and $lcs_i[c]$ variables, respectively. The color-shared cells are any cells on overlapping paths, and $lcs_i[c]$ corresponds to the colors of each disjoint set of color-shared cells.

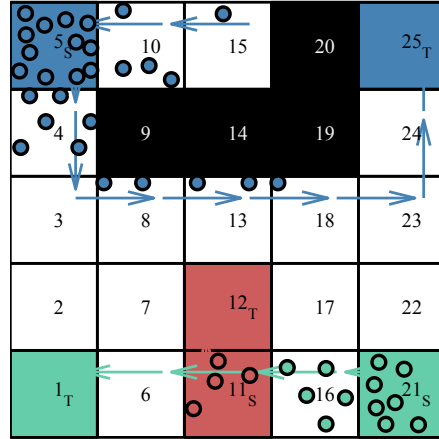


Figure 6: Example illustrating the two fairness requirements (Assumption 4) for proving liveness. Cells 9, 14, 19, and 20 failed, causing the original source-target path for blue to change from cells 5, 10, 15, 20, 25. If source cell 5 does not place new entities fairly, then entities on cells 10 and 15 may never reach the target. A similar situation occurs with paths of multiple colors in the lower part of the image.

- (ii) $Q_0 \subseteq val(X)$ is the set of *start states*,
- (iii) A is a set of *transition names*, and
- (iv) $\rightarrow \subseteq val(X) \times A \times val(X)$ is a set of *discrete transitions*. For $(\mathbf{x}_k, a, \mathbf{x}_{k+1}) \in \rightarrow$, we also use the notation $\mathbf{x}_k \xrightarrow{a} \mathbf{x}_{k+1}$.

An *execution fragment* of a discrete transition system \mathcal{A} is a (possibly infinite) sequence of states $\alpha = \mathbf{x}_0, \mathbf{x}_1, \dots$, such that for each index k appearing in α , $(\mathbf{x}_k, a, \mathbf{x}_{k+1}) \in \rightarrow$ for some $a \in A$. An *execution* is an execution fragment with $\mathbf{x}_0 \in Q_0$. A state \mathbf{x} is said to be *reachable* if there exists a finite execution that ends in \mathbf{x} . \mathcal{A} is said to be *safe* with respect to a set $S \subseteq val(X)$ if all reachable states are contained in S . A set S is said to be *stable* if, for each $(\mathbf{x}, a, \mathbf{x}') \in \rightarrow$, $\mathbf{x} \in S$ implies that $\mathbf{x}' \in S$. \mathcal{A} is said to *stabilize* to S if S is stable and every execution fragment eventually enters S .

Cells. We assume messages are delivered within bounded time and computations are instantaneous. Under these assumptions, the system can be modeled as a collection of discrete transition systems. The overall system is obtained by composing the transition systems of the individual cells. We first present the discrete transition system corresponding to each cell, and then describe the composition.

The variables associated with each $Cell_i$ are as follows, with initial values of the variables shown in Figure 7 using the $:=$ notation.

- (a) $Entities_i$ is the set of identifiers for entities located on cell i . Cell i is said to be nonempty if $Entities_i \neq \emptyset$.
- (b) $color_i$ designates the entity colors on the cell, or \perp if there are none².
- (c) $failed_i$ indicates whether or not i has failed.
- (d) $NEPrev_i$ are the nonempty neighbors attempting to move entities (of any color) toward cell i .
- (e) $token_i$ is a token used for fairness to indicate which neighbor may move toward i .
- (f) $signal_i$ is the identifier of a neighbor of $Cell_i$ that has permission to move toward $Cell_i$.

Additionally, the following variables are defined as arrays for each color $c \in C$. The notation $next_i[c]$ means the c^{th} entry of the $next$ variable of cell i , and so on for the other variables.

- (a) $next_i[c]$ is the neighbor towards which i attempts to move entities of color c .
- (b) $dist_i[c]$ is the estimated distance—the number of cells—to the nearest target cell consuming entities of color c .
- (c) $lock_i[c]$ is a boolean variable for a lock of color c that some cells require to be able to move entities.
- (d) $path_i[c]$ is the set of cell identifiers from any source of color c (and any nonempty cell with entities of color c) to the target of color c . This variable and the next two are local variables for cell i , but they are storing some global information.
- (e) $pint_i[c]$ is the set of cell identifiers in traffic intersections with cells of color c (where $path_i[c]$ and $path_i[d]$ have nonempty intersection for some $d \neq c$).
- (f) $lcs_i[c]$ is the set of colors that are involved in traffic intersections with the color c path.

When clear from context, the subscripts in the names of the variables are dropped. A state of $Cell_i$ refers to a valuation of all these variables, i.e., a function that maps each variable to a value of the corresponding type. The complete system is an automaton, called **System**, consisting of the composition of all the cells. A state of **System** is a valuation of all the variables for all the cells. We refer to states of **System** with bold letters \mathbf{x} , \mathbf{x}' , etc.

²It will be established that cells contain entities of only a single color, see Invariant 3.

<pre> 1 variables Entities : Set[P] := {} 3 NEPrev : Set[ID_⊥] := {} signal, token : ID_⊥ := ⊥ 5 color : C_⊥ := ⊥ failed : B := false 7 next : [C → ID_⊥], init ∀c ∈ C, next[c] := ⊥ dist : [C → N_∞], init ∀c ∈ C, dist[c] := ∞ 9 path : [C → Set[ID_⊥]], init ∀c ∈ C, path[c] := {} pint : [C → Set[ID_⊥]], init ∀c ∈ C, pint[c] := {} 11 nlock : [C → B], init ∀c ∈ C, nlock := true lock : [C → B], init ∀c ∈ C, lock := false 13 lcs : [C → Set[C]], init ∀c ∈ C, lcs[c] := {} </pre>	<pre> transitions 14 fail(i) eff failed := true 16 for each c ∈ C dist[c] := ∞; next[c] := ⊥ 18 update 20 eff Route; Lock; Signal; Move </pre>
<p>Figure 7: Specification of Cell_i listing its variables, initial conditions, and transitions. Subscripts are dropped for readability.</p>	

Variables $token_i$, $failed_i$, $lock_i$, and $NEPrev_i$ are private to Cell_i, while $Entities_i$, $dist_i$, $next_i$, $path_i$, $color_i$, and $signal_i$ can be read by neighboring cells of Cell_i. This has the following interpretation for an actual message-passing implementation. At the beginning of each round, Cell_i broadcasts messages containing the values of these variables and receives similar values from its neighbors. Then, the computation of this round updates the local variables for each cell based on the values collected from its neighbors.

Variable $Entities_i$ is a special variable because it can also be written to by the neighbors of i . This is how we model transfer of entities between cells. For a state \mathbf{x} , for some $a \in A$ such that $\mathbf{x} \xrightarrow{a} \mathbf{x}'$, for some $i \in ID$, for some $j \in Nbrs_i$, for some entity $p \in \mathbf{x}.Entities_i$, then entity p transfers from cell i to j when $p \in \mathbf{x}'.Entities_j$. We use the notation p' to denote the state of entity p at \mathbf{x}' where $\mathbf{x} \xrightarrow{a} \mathbf{x}'$ for some $a \in A$.

Actions for the Composed System. System is a discrete transition system modeling the composition of all the cells, and has two types actions: fails and updates. A $fail(i)$ transition models the crash failure of the i^{th} cell and sets $failed_i$ to *true*, $dist_i[c]$ to ∞ for each $c \in C$, and $next_i[c]$ to \perp for each $c \in C$. Cell i is called *faulty* if $failed_i$ is *true*, otherwise it is called *non-faulty*. The set of identifiers of all faulty and non-faulty cells at a state \mathbf{x} is denoted by $F(\mathbf{x})$ and $NF(\mathbf{x})$, respectively. A faulty cell does nothing—it never moves and it never communicates³.

An update transition models the evolution of all non-faulty cells over one synchronous round. For readability, we describe the state-change caused by an update transition as a sequence of four functions (subroutines), where for each non-faulty i ,

- (a) *Route* computes the variables $dist_i$ and $next_i$,

³ $dist_i = \infty$ can be interpreted as i 's neighbors not receiving a timely response from i .

```

1 if  $\neg failed_i$  then
    $color_i := \{c \in C : \exists p \in Entities_i \wedge color(p) = c\}$ 
3 if  $i \notin ID_T$  then
   for each  $c \in C$ 
5      $dist_i[c] := \left( \min_{j \in Nbrs_i} dist_j[c] \right) + 1$ 
   if  $dist_i[c] = \infty$  then  $next_i[c] := \perp$ 
7   else  $next_i[c] := \operatorname{argmin}_{j \in Nbrs_i} \langle dist_j[c], j \rangle$ 

```

Figure 8: *Route* function for $Cell_i$. This function computes a minimum distance vector routing spanning tree rooted composed of non-faulty cells for each color, rooted at each target.

- (b) *Lock* computes the variables $path_i$, $pint_i$, lcs_i , and $lock_i$,
- (c) *Signal* computes (primarily) the variable $signal_i$, and
- (d) *Move* computes the new positions of entities.

We note that in the single-color case considered in [32], the *Lock* subroutine is unnecessary.

The entire update transition is atomic, so there is no possibility to interleave fail transitions between the subroutines of update. Thus, the state of System at (the beginning of) round $k + 1$ is obtained by applying these four functions to the state at round k . Now we proceed to describe the distributed traffic control algorithm that is implemented through these functions.

Route. For each cell and each color, the *Route* function (Figure 8) constructs a distance-based routing table to the target cell of that color. This relies only on neighbors' estimates of distance to the target. Recall that failed cells have $dist[c]$ set to ∞ for every color $c \in C$. From a state \mathbf{x} , for each $i \in NF(\mathbf{x})$, the variable $dist_i[c]$ is updated as 1 plus the minimum value of $dist_j[c]$ for each neighbor j of i . If this results in $dist_i[c]$ being infinity, then $next_i[c]$ is set to \perp , but otherwise it is set to be the identifier with the minimum $dist[c]$ where ties are broken with neighbor identifiers.

Next, we introduce some definitions used to relate the system state to the variables used in the algorithm. For a state \mathbf{x} , we inductively define the *color c target distance* ρ_c of a cell $i \in ID$ as the smallest number of non-faulty cells between i and tid_c :

$$\rho_c(\mathbf{x}, i) \triangleq \begin{cases} \infty & \text{if } \mathbf{x}.failed_i, \\ 0 & \text{if } i = tid_c \wedge \neg \mathbf{x}.failed_i, \\ 1 + \min_{j \in \mathbf{x}.Nbrs_i} \rho_c(\mathbf{x}, j) & \text{otherwise.} \end{cases}$$

A cell is said to be *target-connected* to color c if ρ_c is finite. We define

$$TC(\mathbf{x}, c) \triangleq \{i \in NF(\mathbf{x}) \mid \rho_c(\mathbf{x}, i) < \infty\}$$

as the set of cells that are target-connected to tid_c .

For a state \mathbf{x} and a color $c \in C$, we define the *routing graph* as $G_R(\mathbf{x}, c) = (V_R(\mathbf{x}, c), E_R(\mathbf{x}, c))$, where the vertices and directed edges are, respectively,

$$V_R(\mathbf{x}, c) \triangleq NF(\mathbf{x}) \text{ and}$$

$$E_R(\mathbf{x}, c) \triangleq \{(i, j) \in V_R(\mathbf{x}, c) : \rho_c(\mathbf{x}, j) = \rho_c(\mathbf{x}, i) + 1\}.$$

Under this definition, $G_R(\mathbf{x}, c)$ is a spanning tree rooted at tid_c . We will show that the graph induced by the $next_i[c]$ variables stabilizes to the routing graph $G_R(\mathbf{x}, c)$ at some state \mathbf{x} (). We previously introduced Δ as the worst-case diameter of the communication graph, and will refer to $\Delta(\mathbf{x})$ as the exact diameter at some state \mathbf{x} .

Lock. The *Lock* function (Figure 9) executes after *Route*, and schedules traffic over intersections (the cells where source-to-target paths of different colors overlap). To avoid deadlock scenarios, *Lock* maintains an invariant that entities of at most one color are on these intersections.

Moving entities over intersections requires some global coordination as illustrated by the following analogy. Consider the policy used to coordinate entities going in opposite directions over a one-lane bridge (see Figure 2), where there is a traffic signal on each side of the bridge. The algorithm chooses one traffic light, allowing some entities to safely travel over the bridge in one direction. After some time, the algorithm switches the lights (first turning green to red, and after the road is clear, turning red to green) allowing traffic to flow in the opposite direction. Then this process repeats.

Two parts of the previous example require global coordination and are included in the *Lock* function. The first is how to choose the direction in which entities are allowed to travel—this is accomplished through the use of a mutual exclusion algorithm. The second is when to allow entities to travel the other direction—this is accomplished by determining when the intersection is empty. We now describe this global coordination more formally.

For defining the locking algorithm, we first define intersections. For this we introduce the notion of an *entity graph*. Cell i is said to be in the *entity graph* of some color c at state \mathbf{x} if one of the following conditions hold: (a) i is sid_c , (b) in state \mathbf{x} , i has entities of color c , or (c) in state \mathbf{x} , i is the neighbor closest to tid_c of a cell already in the entity graph. Formally, we define the *color c entity graph* at state \mathbf{x} as $G_E(\mathbf{x}, c) = (V_E(\mathbf{x}, c), E_E(\mathbf{x}, c))$, which is the following subgraph of the color c routing graph $G_R(\mathbf{x}, c)$. The vertices of $G_E(\mathbf{x}, c)$ are inductively defined as

$$V_E(\mathbf{x}, c) \triangleq \{i \in NF(\mathbf{x}) : i = sid_c \vee \mathbf{x}.color_i = c \vee (\exists j \in V_E(\mathbf{x}, c). (i, j) \in E_R(\mathbf{x}, c))\}.$$

The edges of $G_E(\mathbf{x}, c)$ are $E_E(\mathbf{x}, c) \triangleq \{(i, j) \in V_E(\mathbf{x}, c) \times V_E(\mathbf{x}, c) : (i, j) \in E_R(\mathbf{x}, c)\}$. For example, if all cells are empty, then $V_E(\mathbf{x}, c)$ is the sequence of cell identifiers defined by following the minimum distance (as defined by ρ_c)

```

1 if  $\neg$ failedi
  for each  $c \in C$ 
3   if  $i = sid_c \vee color_i = c \vee i \in path_i[c]$  then
      $path_i[c] := path_i[c] \cup \{i\} \cup \{next_i[c]\}$ 
5   // gossip the entity graph
     for each  $j \in Nbrs_i, path_i[c]$   $:= path_i[c] \cup path_j[c]$ 
7   // compute the set of color-shared cells
      $pint_i[c] := \{j \in path_i[c] \cap path_i[d] : \exists c \neq d \in C\}$ 
9   if  $pint_i[c] \neq \emptyset$ 
      $lcs_i[c] :=$ 
11     $\{d \in C : c \neq d \wedge path_i[c] \cap path_i[d] \neq \emptyset\}$ 
     // graphs stabilized and i needs a lock for color c
13   if  $round > 2\Delta \wedge i \in pint_i[c] \wedge \neg lock_i[c]$ 
     Initiate mutual exclusion algorithm between all
15     color-shared cells in  $pint_i[c]$  using  $lcs_i$  as input
     Eventually, a color  $d$  is returned.
17     On return, if  $d = c$  then  $lock_i[c] := true$ 
     // detect if color-shared cells are empty
19   if  $round > 2\Delta \wedge i \in pint_i[c] \wedge lock_i[c]$ 
     Initiate distributed snapshot algorithm to decide
21     if all color-shared cells are empty after previously
     being nonempty with entities of color  $c$ .
23   On return, if all cells are empty then
      $lock_i[c] := false$ 

```

Figure 9: *Lock* function for $Cell_i$. This function computes the color-shared cells—the cells in intersections—for each color, and then ensures liveness by giving a lock to only one color on each intersection.

from the source to the target of color c . That is, each $G_E(\mathbf{x}, c)$ is a simple path graph from source to target⁴.

Now we describe how the entity graph of each color c is computed by each cell i as the $path_i[c]$ variable. If i is on the entity graph of color c , then we add i and i 's $next$ variable for color c to the entity graph (see Figure 9, lines 3 and 4). Once the $next_i[c]$ variables stabilize () and after an additional order of diameter rounds, the variable $path_i[c]$ contains all the entity graphs since we gossip these graphs (line 6). That is, the graph formed by the $path_i[c]$ variables stabilizes to equal $G_E(\mathbf{x}, c)$, and contains the sequence of identifiers from any source or nonempty cell of color c to the target of color c ().

Next, the variable $pint_i[c]$ is computed to be the set of cell identifiers on the color c entity graph that overlaps with any other colored entity graph (line 8). The cells involved in such non-empty intersections represent physical traffic intersections, and are called *color-shared cells*. These cells require coordinated locking for traffic flow to progress. Cell i is in $pint_i[c]$ if and only if it will need a lock for color c .

Formally, we define the c *color-shared cells*, for a state \mathbf{x} , for each $c \in C$, as

$$CSC(\mathbf{x}, c) = \{V_E(\mathbf{x}, c) : \exists d \in C. c \neq d \wedge V_E(\mathbf{x}, c) \cap V_E(\mathbf{x}, d) \neq \emptyset\}.$$

⁴Once cells have failed, this may stabilize to be a tree from any cell with entities of color c to the target of color c .

<pre> if $\neg failed_i \wedge round > 2\Delta$ then 2 $cn := \{d \in C : \exists j \in Nbrs_i \text{ s.t. } next_j[d] = i$ $\wedge color_j = d\}$ 4 if $color_i = \perp$ then $c := \text{choose from } cn$ 6 else $c := color_i$ 8 $NEPrev_i :=$ $\{j \in Nbrs_i : next_j[c] = i \wedge Entities_j \neq \emptyset\}$ </pre>	<pre> if $token_i = \perp$ then $token_i := \text{choose from } NEPrev_i$ 12 let $j = token_i$ 14 if $\forall p \in Entities_i : \bar{p} \notin SR(i, j)$ $\wedge (color_i \neq \perp \Rightarrow color_i = color_j)$ 16 $\wedge (j \in pint_i[c] \Rightarrow lock_j[c])$ then 18 $signal_i := j$ if $NEPrev_i > 1$ then 20 $token_i := \text{choose from } NEPrev_i \setminus \{j\}$ elseif $NEPrev_i = 1$ then 22 $token_i := \text{choose from } NEPrev_i$ else $token_i := \perp$ 24 else $signal_i := \perp; token_i := j$ </pre>
<p>Figure 10: <i>Signal</i> function for Cell_{<i>i</i>}. Cell <i>i</i> signals fairly to some neighbor <i>j</i> if it is safe for <i>j</i> to move its entities toward <i>i</i>.</p>	

In Figure 1, these are cells 8 and 13 for both colors. The $pint_i[c]$ variables stabilize to equal $CSC(\mathbf{x}, c)$, at some state \mathbf{x} , for any color c ().

Next, we need to determine the colors that will need to coordinate to schedule traffic through the color-shared cells. Then, a mutual exclusion algorithm is initiated between all cells for each disjoint set of cell colors in $pint_i[c]$. Formally, we define the *c shared colors*, for a state \mathbf{x} , for any $c \in C$, as

$$SC(\mathbf{x}, c) = \{d \in C : c = d \vee (c \neq d \wedge CSC(\mathbf{x}, d) = CSC(\mathbf{x}, c))\}.$$

The $lcs_i[c]$ variables stabilize at some state \mathbf{x} to equal $SC(\mathbf{x}, c)$, for any color c .

In general, up to $|C|$ colors could be involved in intersections, as well as all the smaller combinations. For instance, consider Figure 5 with 6 colors at some state \mathbf{x} . Here, the blue and red entity graphs overlap, green and blue entity graphs overlap, but red and green do not, and independently, the purple and yellow entity graphs overlap (that is, not with blue, red, nor green), but no colors overlap with brown. Then $SC(\mathbf{x}, c)$ is $\{blue, red, green\}$ for c equal to blue, red, or green, $SC(\mathbf{x}, c)$ is $\{yellow, purple\}$ for c equal to yellow or purple, and $SC(\mathbf{x}, c)$ is empty for c equal to brown. Two mutual exclusion algorithms would be initiated, one with blue, red, and green as the input set of values, and another with yellow and purple as the input set. Upon these two instances terminating, one element of the first set, say *green*, would be chosen and given a lock, and one element, say *yellow*, of the second set would also be given a lock. The entities of these colors progress over the color-shared cells toward their intended targets. Finally, once the color-shared cells are empty again, *green* and *yellow* would each be removed from the respective input sets for fairness, and another mutual exclusion algorithm is initiated.

Signal. The *Signal* function (Figure 10) executes after *Lock*. It is the key part of the protocol for maintaining safe entity separations, guaranteeing each cell has entities of only a single color, and ensuring progress of entities to the target. Roughly, each cell implements this through the following policies: (a) only

accept entities from a neighbor when it is safe to do so, (b) only accept entities with the same color as the entities currently on the cell (or an arbitrary color if the cell is empty), (c) if a lock is needed, then only let entities move if it is acquired, and (d) ensure fairness by providing opportunities infinitely often for each nonempty neighbor to make progress.

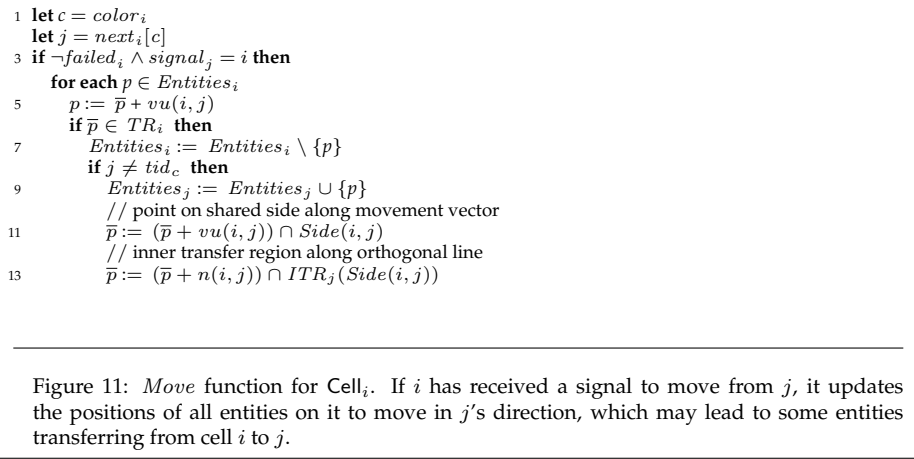
First i computes a temporary variable cn , which is the set of colors for any neighbor that has entities of some color, with the corresponding $next$ variable set to cell i . Next, cell i picks a color c from this set if it is empty, or the color of its own entities if it is nonempty, and will attempt to allow some cell with this chosen color to move toward itself. Then, cell i sets $NEPrev_i$ to be the subset of $Nbrs_i$ for which $next$ has been set to i and $Entities$ is nonempty. If $token_i$ is \perp , then it is set to some arbitrary value in $NEPrev_i$, but it continues to be \perp if $NEPrev_i$ is empty. Otherwise, $token_i = j$ for some neighbor j of i with nonempty $Entities_j$. This is accomplished through the conditional in line 6 as a step in guaranteeing fairness.

It is then checked if there is any entity p with center \bar{p} in the safety region of $Cell_i$ on the side corresponding to $token_i$. If there is such an entity, then $signal_i$ is set to \perp , which blocks the neighboring cell with identifier $token_i$ from moving its entities in the direction of i , thus preventing entity transfers and ensuring safety. Otherwise, if there is no entity with center in the safety region on side $token_i$, then $signal_i$ is set to $token_i$ to allow $token_i$ to move its entities toward i . Subsequently, $token_i$ is updated to a value in $NEPrev_i$ that is different from its previous value, if that is possible according to the rules just described (lines 20–22).

Move. Finally, the *Move* function (Figure 11) models the physical movement of all the entities on cell i over a given round. For cell i , let j be $next_i[c]$, where c is $color_i$ (which may be \perp if cell i has no entities). Every entity in $Entities_i$ moves in the direction of j if and only if $signal_j$ is set to i . The direction followed from cell i to j is $u(i, j)$, which is any vector satisfying Assumption 1. For example, for a square (or rectangular) cell i , one choice for $u(i, j)$ is the unit vector orthogonal to $Side(i, j)$ and pointing into j . In the case of an equilateral triangular cell i , one choice for $u(i, j)$ is also any orthogonal vector pointing into j .

The movement toward cell j may lead to some entities crossing the boundary of $Cell_i$ into $Cell_j$, in which case, they are removed from $Entities_i$. If j is not the target matching the transferred entities' color, then the removed entities are added to $Entities_j$. In this case (line 9), any transferred entity p is placed so that $D_l(p)$ touches a single point of (is tangent to) $Side(i, j)$, the shared side of cells i and j , and lies on the inner side of the transfer region of cell j on side $Side(i, j)$. Resetting entity positions is a conservative approximation to the actual physical movement of entities. If j is the target matching the transferred entities' color, then the removed entities are not added to any cell and thus no longer exist in System.

The source cells $i \in ID_S$, in addition to the above, add a finite number of entities in each round to $Entities_i$, such that the addition of these entities does



not violate the minimum gap between entities at $Cell_i$. In the remainder of the paper, we will analyze System to show that in spite of failures, it maintains safety and liveness properties to be introduced in the next section.

4. Analysis of Distributed Traffic Control

In this section, we present an analysis of the safety and liveness properties of System. Roughly, the safety property requires that there is a minimum gap between entities on any cell, and the liveness property requires that all entities that reside on cells with feasible paths to the corresponding target eventually reach that target.

4.1. Safety and Collision Avoidance

A state is safe if, for every cell, the boundaries of all entities in the cell are separated by a distance of r_s . For any state \mathbf{x} of System, we define:

$$\begin{aligned}
Safe_i(\mathbf{x}) &\triangleq \forall p, q \in \mathbf{x}. Entities_i.p \neq q \Rightarrow \|\bar{p} - \bar{q}\| \geq 2l + r_s, \text{ and} \\
Safe(\mathbf{x}) &\triangleq \forall i \in ID, Safe_i(\mathbf{x}).
\end{aligned}$$

This definition allows entities in different cells to be closer than $2l + r_s$ apart, but their centers will be spaced by at least $2l$. We proceed by proving some preliminary properties of System that will be used for proving *Safe* is an invariant.

The first property asserts that entities' cannot come close enough to the sides of cells to reside on multiple cells. This is because any entity whose boundary touches the side of a cell is transferred to the neighboring cell on that side (if one exists), and then the entity's position is reset to be completely within the new cell. Assumption 2 restricts the allowed partitions to ensure entity transfers are well-defined. For instance, some of the cells in the snub square

tiling in Figure 3 do not satisfy Assumption 2. Consider an entity transfer from cell 3 to cell 5. There is no constant vector connecting the transfer regions of cell 3 to those of cell 5. This is because the side length of the transfer region of the triangular cell 5 is shorter than the side length of the transfer region of the square cell 3. However, in a transfer from cell 1 to cell 2 or vice-versa, the side lengths are the same. We also note that the assumption is only necessary for entity transfers from a cell with a longer transfer side length to a neighboring cell with smaller corresponding transfer side length. For example, a transfer from cell 5 to cell 3 is feasible.

Under Assumption 2, we have the following invariant, which states that the l -ball around each entity in a cell is completely contained within the cell.

Invariant 1. *In any reachable state \mathbf{x} , $\forall i \in ID, \forall p \in \mathbf{x}.Entities_i, D_l(p) \setminus P_i = \emptyset$.*

The next invariant states that cells' *Entities* sets are disjoint. This is immediate from the *Move* function since entities are only added to one cell's *Entities* upon being removed from a different cell's *Entities*.

Invariant 2. *In any reachable state \mathbf{x} , for any $i, j \in ID$, if $i \neq j$, then $\mathbf{x}.Entities_i \cap \mathbf{x}.Entities_j = \emptyset$.*

The following invariant states that cells contain entities of a single color in spite of failures. This follows from the *Signal* routine in Figure 10, where line 16 requires that if some neighbor j is attempting to move entities toward cell i , then the color of i is either \perp or equal to the color of j .

Invariant 3. *In any reachable state \mathbf{x} , for all $i \in ID$, for all $p, q \in \mathbf{x}.Entities_i$, $color(p) = color(q)$.*

Next, we define a predicate that states that if $signal_i$ is set to the identifier of some neighbor $j \in Nbrs_i$, then there is a large enough area from the common side between i and j where no entities reside in $Cell_i$. Recall that $Side(i, j)$ is the line segment shared between neighboring cells i and j . For a state \mathbf{x} , $H(\mathbf{x}) \triangleq \forall i \in ID, \forall j \in Nbrs_i, \text{if } \mathbf{x}.signal_i = j, \text{ then the following holds:}$

$$\forall p \in \mathbf{x}.Entities_i, \min_{x \in Side(i, j)} \|\bar{p} - x\| \geq 3d.$$

$H(\mathbf{x})$ is not an invariant property because once entities move the property may be violated. However, for proving safety, all that needs to be established is that at the *point of computation of the signal variable* this property holds. The next key lemma states this.

Lemma 4. *For all reachable states \mathbf{x} , $H(\mathbf{x}) \Rightarrow H(\mathbf{x}_S)$ where \mathbf{x}_S is the state obtained by applying the *Route*, *Lock*, and *Signal* functions to \mathbf{x} .*

Proof. Fix a reachable state \mathbf{x} , an $i \in ID$, and an $j \in Nbrs_i$ such that $\mathbf{x}.signal_i = j$. Let \mathbf{x}_R be the state obtained by applying the *Route* function to \mathbf{x} , \mathbf{x}_L be the state obtained by applying the *Lock* function to \mathbf{x}_R , and \mathbf{x}_S be the state obtained by applying the *Signal* function to \mathbf{x}_L .

First, observe that both $H(\mathbf{x}_R)$ and $H(\mathbf{x}_L)$ hold. This is because the *Route* and *Lock* functions do not change any of the variables involved in the definition of $H(\cdot)$. Next, we show that $H(\mathbf{x}_L)$ implies $H(\mathbf{x}_S)$. If $\mathbf{x}_S.\text{signal}_i \neq j$ then the statement holds vacuously. Otherwise, $\mathbf{x}_S.\text{signal}_i = j$, then since (a) $H(\mathbf{x}_L)$ holds, and (b) Figure 10, line 6 is satisfied, we have that $H(\mathbf{x}_S)$. \square

The following lemma asserts that if there is a cycle of length two formed by the *signal* variables—which could occur due to failures—then entity transfers cannot occur between the involved cells in that round.

Lemma 5. *Let \mathbf{x} be any reachable state and \mathbf{x}' be a state that is reached from \mathbf{x} after a single update transition (round). If $\mathbf{x}.\text{signal}_i = j$ and $\mathbf{x}.\text{signal}_j = i$, then $\mathbf{x}.\text{Entities}_i = \mathbf{x}'.\text{Entities}_i$ and $\mathbf{x}.\text{Entities}_j = \mathbf{x}'.\text{Entities}_j$.*

Proof. No entities enter either $\mathbf{x}'.\text{Entities}_i$ or $\mathbf{x}'.\text{Entities}_j$ from any other $m \in \text{Nbrs}_i$ or $n \in \text{Nbrs}_j$ since $\mathbf{x}.\text{signal}_i = j$ and $\mathbf{x}.\text{signal}_j = i$. It remains to be established that $\nexists p \in \mathbf{x}.\text{Entities}_j$ such that $p' \in \mathbf{x}'.\text{Entities}_i$ where $p = p'$ or vice-versa. Suppose such a transfer occurs. For the transfer to have occurred, \bar{p} must be such that $\bar{p}' = (p_x, p_y) + vu(i, j)$ by Figure 11, line 5. But for $\mathbf{x}.\text{signal}_i = j$ to be satisfied, it must have been the case that $D_i(p) \cap P_i = \emptyset$ by Figure 10, line 6 and since $v < l$, a contradiction is reached. \square

Using the previous results, we now prove that System preserves safety even when some cells fail.

Theorem 1. *In any reachable state \mathbf{x} of System, $\text{Safe}(\mathbf{x})$.*

Proof. The proof is by standard induction over the length of any execution of System. The base case is satisfied by the assumption that initial states $\mathbf{x} \in Q_0$ satisfy $\text{Safe}(\mathbf{x})$. For the inductive step, consider any reachable states \mathbf{x} , \mathbf{x}' and an action $a \in A$ such that $\mathbf{x} \xrightarrow{a} \mathbf{x}'$. Fix $i \in ID$ and assuming $\text{Safe}_i(\mathbf{x})$, we show that $\text{Safe}_i(\mathbf{x}')$. If $a = \text{fail}_i$, then $\text{Safe}_i(\mathbf{x}')$ since no entities move.

For $a = \text{update}$, there are two cases to consider by Invariant 2. First, $\mathbf{x}'.\text{Entities}_i \subseteq \mathbf{x}.\text{Entities}_i$, that is, no new entities were added to i , but some may have transferred off i . There are two sub-cases: if $\mathbf{x}'.\text{Entities}_i = \mathbf{x}.\text{Entities}_i$, then all entities in $\mathbf{x}.\text{Entities}_i$ move identically and the spacing between two distinct entities $p, q \in \mathbf{x}'.\text{Entities}_i$ is unchanged. Let $j = \text{next}_i[c]$ where $c = \text{color}_i$ by Invariant 3. That is, $\forall p, q \in \mathbf{x}.\text{Entities}_i, \forall p', q' \in \mathbf{x}'.\text{Entities}_i$ such that $p' = p$ and $q' = q$ and where $p \neq q$, $\|(p'_x, p'_y) - (q'_x, q'_y)\| = \|(p_x, p_y) + vu(i, j), (q_x, q_y) + vu(i, j)\|$ (Figure 11, line 5). It follows by the inductive hypothesis that $\|(p'_x, p'_y) - (q'_x, q'_y)\| \geq d$. The second sub-case arises if $\mathbf{x}'.\text{Entities}_i \subsetneq \mathbf{x}.\text{Entities}_i$, then $\text{Safe}_i(\mathbf{x}')$ is either vacuously satisfied or it is satisfied by the same argument just stated.

The second case is when $\mathbf{x}'.\text{Entities}_i \not\subseteq \mathbf{x}.\text{Entities}_i$, that is, there was at least one entity transferred to i . Consider any such transferred entity $p' \in \mathbf{x}'.\text{Entities}_i$ where $p' \notin \mathbf{x}.\text{Entities}_i$. There are two sub-cases. The first sub-case is when p' was added to $\mathbf{x}'.\text{Entities}_i$ because i is a source, that is, $i \in ID_S$. In this case, the specification of the source cells states that the entity p' was added to $\mathbf{x}'.\text{Entities}_i$ without violating $\text{Safe}_i(\mathbf{x}')$, and the proof is complete. Otherwise,

p' was added to $\mathbf{x}'.Entities_i$ by some neighbor $j \in \mathbf{x}.Nbrs_i$, so $p' \in \mathbf{x}.Entities_j$ but $p' \notin \mathbf{x}.Entities_i$, and $p' \in \mathbf{x}'.Entities_i$ but $p' \notin \mathbf{x}'.Entities_j$. From line 9 of Figure 11, we have that that $(p'_x, p'_y) = ResetEntity(p, i, j)$. The fact that p' transferred from $Cell_j$ in \mathbf{x} to $Cell_i$ in \mathbf{x}' implies that $\mathbf{x}.next_j = i$ and $\mathbf{x}.signal_i = j$ —these are necessary conditions for the transfer by Figure 10, line 15. Thus, applying the predicate $H(\mathbf{x})$ at state \mathbf{x} and by Lemma 4, it follows that for every $q \in \mathbf{x}.Entities_i$, $(q_x, q_y) \notin FR(i, j)$. It must now be established that if p' is transferred to $\mathbf{x}'.Entities_i$, then every $q' \in \mathbf{x}'.Entities_i$, where $q' \neq p'$ satisfies $(q'_x, q'_y) \notin FR(i, j)$, which means that any entity q already on i did not move toward the transferred entity p that is now on i . This follows by application of Lemma 5, which states that if entities on adjacent cells move towards one another simultaneously, then a transfer of entities cannot occur. This implies that the discs of all entities q' in $\mathbf{x}'.Entities_i$ are farther than r_s of the borders of any transferred entity p' , implying $Safe_i(\mathbf{x}')$. Finally, since i was chosen arbitrarily, $Safe(\mathbf{x}')$. \square

Theorem 1 shows that System is safe in spite of failures.

4.2. Stabilization of Spanning Routing Trees

Next, we show under some additional assumptions, that once new failures cease to occur, System recovers to a state where each non-faulty cell with a feasible path to its target computes a route toward it. This route stabilization is then used in showing that any entity on a non-faulty cell with a feasible path to its target makes progress toward it. Our analysis relies on the following assumptions on cell failures and the placement of new entities on source cells. The first assumption states that no target cells fail, and is reasonable and necessary because if any target cell did fail, entities of that color obviously cannot make progress.

Assumption 3. *No target cells $t \in ID_T$ may fail.*

The next assumption ensures source cells place entities fairly so that they may not perpetually prevent any neighboring cell or any color-shared cell from making progress. The assumption is needed because it provides a specification of how the source cells behave, which has not been done so far. The assumption is reasonable because it essentially says that traffic is not produced perpetually without any break.

Assumption 4. *(Fairness): Source cells place new entities without perpetually blocking either (i) any of their nonempty non-faulty neighbors, or (ii) any cell $i \in CSC(\mathbf{x}, c)$, where c is the color of source s .*

Formally, the first fairness condition states, for any execution α of System, for any color $c \in C$, for any source cell sid_c , if there exists an $i \in Nbrs_s$, such that for every state \mathbf{x} in α after a certain round, $i \in \mathbf{x}.NEPrev_s$, then eventually $signal_s$ becomes equal to i in some round of α . The second fairness condition states, for any execution α of System, for any state $\mathbf{x} \in \alpha$, for any color $c \in C$,

for any source cell sid_c , if there exists an $i \in NF(\mathbf{x})$ such that $i \in CSC(\mathbf{x}, c)$, and for every state \mathbf{x} in α after a certain round, if cell i is nonempty, then eventually $signal_j$ becomes equal to i in some round of α , where j is a neighbor of i . Such conditions can be ensured if we suppose some oracle placing entities on source cells follows the same round-robin like scheme defined in the *Signal* subroutine in Figure 10. Scenarios where each of these cases can arise are illustrated in Figure 6.

A fault-free execution fragment α be a sequence of states starting from \mathbf{x} and along which no $fail(i)$ transitions occur. That is, a fault-free execution fragment is an execution fragment with no new failure actions, although there may be existing failures at the first state \mathbf{x} of α , so $F(\mathbf{x})$ need not be empty. Throughout the remainder of this section, we will consider fault-free executions that satisfy Assumptions 3 and 4.

Lemma 6. *Consider any reachable state \mathbf{x} of System, any color $c \in C$, and any $i \in TC(\mathbf{x}, c) \setminus \{tid_c\}$. Let $h = \rho_c(\mathbf{x}, i)$. Any fault-free execution fragment α starting from \mathbf{x} stabilizes within h rounds to a set of states S with all elements satisfying:*

$$\begin{aligned} dist_i[c] &= h, \text{ and} \\ next_i[c] &= i_n, \text{ where } \rho_c(\mathbf{x}, i_n) = h - 1. \end{aligned}$$

Proof. Fix an arbitrary state \mathbf{x} , a fault-free execution fragment α starting from \mathbf{x} , a color $c \in C$, and $i \in TC(\mathbf{x}, c) \setminus \{tid_c\}$. We have to show that (a) the set of states S is closed under update transitions and (b) after h rounds, the execution fragment α enters S .

First, by induction on h we show that S is stable. Consider any state $\mathbf{y} \in S$ and a state \mathbf{y}' that is obtained by applying an update transition to \mathbf{y} . We have to show that $\mathbf{y}' \in S$. For the base case, $h = 1$, so $\mathbf{y}.dist_i[c] = 1$ and $\mathbf{y}.next_i[c] = tid_c$. From lines 5 and 7 of the *Route* function in Figure 8, and that there is a unique tid_c for each color c , it follows that $\mathbf{y}'.dist_i[c]$ remains 1 and $\mathbf{y}'.next_i[c]$ remains tid_c . For the inductive step, the inductive hypothesis is, for any given h , if for any $j \in NF(\mathbf{x})$, $\mathbf{y}.dist_j[c] = h$ and $\mathbf{y}.next_j[c] = m$, for some $m \in ID$ with $\rho_c(\mathbf{x}, m) = h - 1$, then

$$\mathbf{y}'.dist_j[c] = h \text{ and } \mathbf{y}'.next_j[c] = m.$$

Now consider i such that $\rho_c(\mathbf{y}, i) = \rho_c(\mathbf{y}', i) = h + 1$. In order to show that S is closed, we have to assume that $\mathbf{y}.dist_i[c] = h + 1$ and $\mathbf{y}.next_i[c] = m$, and show that the same holds for \mathbf{y}' . Since $\rho_c(\mathbf{y}', i) = h + 1$, i does not have a neighbor with target distance smaller than h . The required result follows from applying the inductive hypothesis to m and from lines 5 and 7 of Figure 8.

Second, we have to show that starting from \mathbf{x} , α enters S within h rounds. Once again, this is established by induction on h , which is $\rho_c(\mathbf{x}, i)$. Consider any state \mathbf{y} such that $\rho_c(\mathbf{x}, i) = \rho_c(\mathbf{y}, i)$. The base case only includes the target distances satisfying $h = \rho_c(\mathbf{y}, i) = 1$ and follows by instantiating $i_n = tid_c$. For the inductive case, assume for the inductive hypothesis that at some state \mathbf{y} , $\mathbf{y}.dist_j[c] = h$ and $\mathbf{y}.next_j[c] = i_n$ such that $\rho_c(\mathbf{y}, i_n) = h - 1$, where i_n is

the minimum identifier among all such cells (since we used cell identifiers to break ties). Observe that there is one such $j \in \mathbf{y}.Nbrs_i$ by the definition of TC . Then at state \mathbf{y}' , by the inductive hypothesis and lines 5 and 7 of Figure 8, $\mathbf{y}'.dist_i[c] = \mathbf{y}'.dist_j[c] + 1 = h + 1$. \square

The following corollary of Lemma 6 states that, after new failures cease occurring, for all target-connected cells, the graph induced by the $next[c]$ variables stabilizes to the color c routing graph, $G_R(\mathbf{x}, c)$, within at most the diameter of the communication graph number of rounds, which is bounded by $\Delta(\mathbf{x})$.

Corollary 7. *Consider any execution α of System with an arbitrary but finite sequence of fail transitions. For any state $\mathbf{x} \in \alpha$ at least $2\Delta(\mathbf{x})$ rounds after the last fail transition, for any $c \in C$, every cell i target-connected to color c has $\mathbf{x}.next_i[c]$ equal to the identifier of the next cell along such a route.*

The following corollary of Lemma 6 and states that within $2\Delta(\mathbf{x})$ rounds after routes stabilize, for each color $c \in C$, the identifiers in the $path_i[c]$ variables equal the vertices of the color c entity graph $G_E(\mathbf{x}, c)$. The result follows since routes stabilize and that $Lock$ is a function of $next$ and $path$ variables only, and that $path_i$ variables are gossiped in Figure 9, line 6.

Corollary 8. *Consider any execution α of System with an arbitrary but finite sequence of fail transitions. For any state $\mathbf{x} \in \alpha$ at least $2\Delta(\mathbf{x})$ rounds after the last fail transition, for every $c \in C$, every cell i target-connected to color c has $path_i[c] = V_E(\mathbf{x}, c)$.*

The next corollary of Lemma 6 states that eventually the values of the $pint[c]$ variables equal the set of color-shared cells $CSC(\mathbf{x}, c)$ for any cell i and color c . This is important because the mutual exclusion algorithm is initiated between the cells in $pint[c]$ (Figure 9, line 13).

Corollary 9. *Consider any execution α of System with an arbitrary but finite sequence of fail transitions. For any state $\mathbf{x} \in \alpha$ at least $2\Delta(\mathbf{x})$ rounds after the last fail transition, for every $c \in C$, every cell i target-connected to color c has $\mathbf{x}.pint[c] = CSC(\mathbf{x}, c)$.*

4.3. Scheduling Entities through Color-Shared Cells

In this section, we show that there is at most a single color on the set of color-shared cells if there are no failures. We then show that any cell that requests a lock eventually gets one, under an additional assumption that failures do not cause entities of more than one color to reside on the set of color-shared cells. Because failures cause the routing graphs and entity graphs to change, the color-shared cells that could previously be scheduled may now be deadlocked. Additionally, because we separately lock each disjoint set of color-shared cells to allow entities of some color to flow toward their target, it could be the case that the intermediate states between when the failure occurred and

when routes have stabilized allowed entities to move in such a way that deadlocks the system. Such deadlocks could be avoided if a centralized coordinator informs every non-faulty cell to disable their signals when a failure is detected. The assumption states that with failures, the color-shared cells either all have the same-colored entities, or have no entities (and combinations thereof).

Assumption 5. Feasibility of Locking after Failures: *For any reachable state \mathbf{x} , for any color $c \in C$, consider the color-shared cells $CSC(\mathbf{x}, c)$. For all distinct cells $i, j \in CSC(\mathbf{x}, c)$ either $\mathbf{x}.color_i = \mathbf{x}.color_j$ or $\mathbf{x}.color_i = \perp$.*

The next lemma states that without failures, there are entities of at most a single color on the set of color-shared cells. The result is not an invariant because failures may cause the set of color-shared cells to change, resulting in deadlocks, which is why we need Assumption 5. By Invariant 3, we know that there are entities of at most a single color in each cell, so the following invariant is stated in terms of the color $color_i$ of each cell. We emphasize that Assumption 5 is unnecessary if there are no failures, as the algorithm ensures there are entities of at most a single color on the color-shared cells by the following lemma.

Lemma 10. *If there are no failures, for any reachable state \mathbf{x} , for any $c \in C$, for any $i \in CSC(\mathbf{x}, c)$, if $\neg \mathbf{x}.lock_i[c]$, then for all $j \in CSC(\mathbf{x}, c)$, we have $\mathbf{x}.color_j \neq c$.*

Proof. The proof is showing an inductive invariant, supposing no failures occur. For the initial state, all cells are empty, so we have $\mathbf{x}.color_i = \perp$ for any $i \in ID$. For the inductive step, we are only considering update actions by assumption. In the pre-state, we have $\neg \mathbf{x}.lock_i[c]$ and $\forall j \in CSC(\mathbf{x}, c)$, we have $\mathbf{x}.color_j \neq c$. Fix some $c \in C$ and some $i \in CSC(\mathbf{x}, c)$. For any subsequent state \mathbf{x}' , if $\mathbf{x}'.lock_i[c]$, the result follows vacuously. If $\neg \mathbf{x}'.lock_i[c]$, we must show $\forall j \in CSC(\mathbf{x}, c)$ that $\mathbf{x}.color_j \neq c$, so fix some $j \in CSC(\mathbf{x}, c)$. If $j \in CSC(\mathbf{x}', c)$, the result follows, since by the inductive hypothesis, $\mathbf{x}.color_j = \mathbf{x}'.color_j \neq c$. If $j \notin CSC(\mathbf{x}', c)$, the condition in *Signal* (Figure 10, line 17) cannot be satisfied since $\neg \mathbf{x}'.lock_i[c]$. Thus, no cell with entities of color c could move toward any cell in $CSC(\mathbf{x}', c)$, and we have $\mathbf{x}'.color_j \neq c$. \square

The next lemma states that without failures, or with “nice” failures as described by Assumption 5, that any cell requesting a lock of some color will eventually get it, and thus it may move entities onto the color-shared cells.

Lemma 11. *For any reachable state \mathbf{x} satisfying Assumption 5, for any $c \in C$, for any $i \in NF(\mathbf{x})$, if $i \in \mathbf{x}.pint[c]$ and all cells in $CSC(\mathbf{x}, c)$ are empty, then eventually a state \mathbf{x}' is reached where $\mathbf{x}'.lock_i[c]$.*

Proof. By correctness of the mutual exclusion algorithm, eventually a color $d \in SC(\mathbf{x}', c)$ is returned and $\mathbf{x}'.lock_i[d] = true$ (Figure 9, line 13). If $c = d$, then the result follows. If $c \neq d$, by Lemma 10 and Assumption 5, we know that no other color aside from c has entities on any cell $j \in CSC(\mathbf{x}', c)$. The next time the mutual exclusion algorithm is initiated, d is excluded from the input set to the mutual exclusion algorithm (Figure 9, line 19), and by repeated argument, eventually $lock_i[c]$. \square

4.4. Progress of Entities towards their Targets

Using the results from the previous sections, we show that once new failures cease occurring, for every color $c \in C$, every entity of color c on a cell that is target-connected eventually gets to the target of color c . The result (Theorem 2) uses two lemmas which establish that, along every infinite execution with a finite number of failures, every nonempty target-connected cell gets permission to move infinitely often (Lemma 13), and a permission to move allows the entities on a cell to make progress towards the target (Lemma 12).

For the remainder of this section, we fix an arbitrary infinite execution α of System with a finite number of failures, satisfying Assumption 5. Let \mathbf{x}_f be any state of System at least $2\Delta(\mathbf{x})$ rounds after the last failure, and α' be the infinite failure-free execution fragment $\mathbf{x}_f, \mathbf{x}_{f+1}, \dots$ of α starting from \mathbf{x}_f . For any $c \in C$, observe that the number of target-connected cells remains constant starting from \mathbf{x}_f for the remainder of the execution. That is, $TC(\mathbf{x}_f, c) = TC(\mathbf{x}_{f+1}, c) = TC(\dots, c)$, so we fix $TC(c) = TC(\mathbf{x}_f, c)$.

Lemma 12. *For any $c \in C$, for any $i \in TC(c)$, for some $j \in \mathbf{x}_f.Nbrs_i$, if $k > f$, $\mathbf{x}_k.signal_j = i$, and $\mathbf{x}_k.next_i[c] = j$, for any entity $p \in \mathbf{x}_k.Entities_i$, let the distance function be defined by the lexicographically ordered tuple*

$$R(\mathbf{x}, p) = \langle \rho_c(\mathbf{x}, i), ds - \bar{p} \rangle,$$

where ds is the point on the shared side $Side(i, j)$ defined by the line passing through \bar{p} with direction $u(i, j)$. Then, $R(\mathbf{x}_{k+1}, p) < R(\mathbf{x}_k, p)$.

Proof. The first case is when no entity transfers from i to j in the $k+1^{th}$ round: if $p' \in \mathbf{x}_{k+1}.Entities_i$ such that $p' = p$, then $\|ds - \bar{p}'\| < \|ds - \bar{p}\|$. In this case, the result follows since a velocity $v > 0$ is applied towards cell j by *Move* in Figure 11, line 5. The second case is when some entity p transfers from i to j , so $p' \in \mathbf{x}_{k+1}.Entities_j$ such that $p' = p$. In this case, we have $\rho_c(\mathbf{x}_k, j) < \rho_c(\mathbf{x}_k, i)$, since the distance between j and tid_c is smaller than the distance between i and tid_c since routes have stabilized by Lemma 6. In either case, $R(\mathbf{x}_{k+1}, p) < R(\mathbf{x}_k, p)$, so entity p is closer to the appropriate target. \square

The following lemma states that all cells with a path to the target receive a signal to move infinitely often, so Lemma 12 applies infinitely often.

Lemma 13. *For any $c \in C$, consider any $i \in TC(c) \setminus tid_c$, such that for all $k > f$, if $\mathbf{x}_k.Entities_i \neq \emptyset$, then $\exists k' > k$ such that $\mathbf{x}_{k'}.signal_{next_i[c]} = i$.*

Proof. Fix some $c \in C$. Since $i \in TC(c)$, there exists $h < \infty$ such that for all $k > f$, $\rho_c(\mathbf{x}_k, i) = h$. We prove the lemma by inducting on h . The base case is $h = 1$. Fix i and instantiate $k' = f + ns(tid_c)$. By Lemma 6, for any $t \in ID_T$, for all non-faulty $i \in Nbrs_t$, $\mathbf{x}_f.next_i[c] = t$ since $k > f$. For all $k > f$, if $\mathbf{x}_k.Entities_i \neq \emptyset$, then $signal_{tid_c}$ changes to a different neighbor with entities every round. It is thus the case that $|\mathbf{x}_k.NEPrev_{tid_c}| \leq ns(tid_c)$ and since $Entities_{tid_c} = \emptyset$ always, exactly one neighbor satisfies the conditional of Figure 10, line 6 in any round, then within $ns(tid_c)$ rounds, $signal_{tid_c} = i$.

For the inductive case, let $k_s = k + h$ be the step in α after which all non-faulty $a \in Nbrs_i$ have $\mathbf{x}_{k_s}.next_a[c] = i$ by Lemma 6. Also by Lemma 6, $\exists m \in Nbrs_i$ such that $\mathbf{x}_{k_s}.dist_m < \mathbf{x}_{k_s}.dist_i$, implying that after k_s , $|\mathbf{x}_{k_s}.NEPrev_i| \leq ns(i)$ since $\mathbf{x}_{k_s}.next_i = m$ and $\mathbf{x}_{k_s}.next_m \neq i$. By the inductive hypothesis, $\mathbf{x}_{k_s}.signal_{next_i[c]} = i$ infinitely often. If $i \in ID_S$, then entity initialization does not prevent $\mathbf{x}_k.signal_i = a$ from being satisfied infinitely often by the second assumption introduced in Subsection 4.2. It remains to be established that $signal_i = a$ infinitely often. Let $a \in \mathbf{x}_{k_s}.NEPrev_i$ where $\rho_c(\mathbf{x}_{k_s}, a) = h + 1$.

In any of the following cases, if $i \in \mathbf{x}_{k_s}.pint[c]$ and all cells $j \in CSC(\mathbf{x}_{k_s}, c)$ are empty, then by Lemma 11, eventually $lock_i[c]$. If $|\mathbf{x}_{k_s}.NEPrev_i| = 1$, then since the inductive hypothesis satisfies $signal_{next_i[c]} = i$ infinitely often, then Lemma 12 applies infinitely often, and thus $Entities_i = \emptyset$ infinitely often, finally implying that $signal_i = a$ infinitely often.

If $|\mathbf{x}_{k_s}.NEPrev_i| > 1$, there are two sub-cases. The first sub-case is when no entity enters i from some $d \neq a \in \mathbf{x}_{k_s}.NEPrev_i$, which follows by the same reasoning used in the $|\mathbf{x}_{k_s}.NEPrev_i| = 1$ case. The second sub-case is when an entity enters i from d , in which case it must be established that $signal_i = a$ infinitely often. This follows since if $\mathbf{x}_{k'}.token_i = a$ where $k' > k_t > k_s$ and k_t is the round at which an entity entered i from d , and the appropriate case of Lemma 4 is not satisfied, then $\mathbf{x}_{k'+1}.signal_i = \perp$ and $\mathbf{x}_{k'+1}.token_i = a$ by Figure 10, line 25. This implies that no more entities enter i from either cell d satisfying $d \neq a$. Thus $token_i = a$ infinitely often follows by the same reasoning $|\mathbf{x}_{k_s}.NEPrev_i| = 1$ case. \square

The final theorem establishes that entities on any cell in $TC(c)$ eventually reach the target in α' .

Theorem 2. *For any $c \in C$, consider any $i \in TC(c)$, $\forall k > f$, $\forall p \in \mathbf{x}_k.Entities_i$, $\exists k' > k$ such that $p \in \mathbf{x}_{k'}.Entities_{next_i[c]}$.*

Proof. Fix $c \in C$, $i \in TC(c)$, a round $k > f$ and $p \in \mathbf{x}_k.Entities_i$. Let $h = \max_{i \in TC(c)} \rho_c(\mathbf{x}_f, i)$ which is finite. By Lemma 6, at every round after $k_s = k + h$ for any $i \in TC(c)$, the sequence of identifiers $\beta = i, \mathbf{x}_{k_s}.next_i[c], \mathbf{x}_{k_s}.next_{next_i[c]}[c], \dots$ forms a fixed path to tid_c . Applying Lemma 13 to $i \in TC(c)$ shows that there exists $k_m \geq k_s$ such that $\mathbf{x}_{k_m}.signal_{next_i[c]} = i$. Now applying Lemma 12 to \mathbf{x}_{k_m} establishes movement of p towards $\mathbf{x}_{k_s}.next_i[c]$, which is also $\mathbf{x}_{k_m}.next_i[c]$. Lemma 13 further establishes that this occurs infinitely often, thus there is a round $k' > k_m$ such that p gets transferred to $\mathbf{x}_{k_m}.Entities_{next_i[c]}$. \square

By an induction on the sequence of identifiers in the path β , it follows that entities on any cell in $TC(c)$ eventually get consumed by the target.

Summary of Results

In this section, we established several invariant properties culminating in proving safety: entities never collide, in spite of failures. Next, we proved that the routing algorithm used to construct paths to the destinations is self-stabilizing in spite of arbitrary crash failures. Next, we showed—under an




Figure 12: Throughput versus safety spacing r_s for several values of v , for $K = 2500$, $l = 0.25$ for System with an 8×8 unit square tessellation.

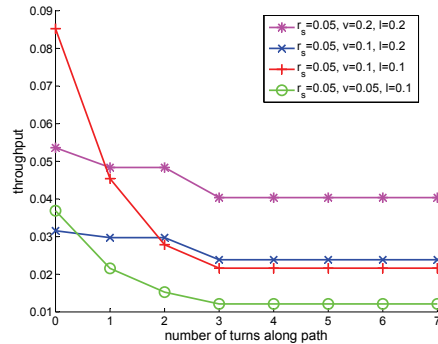


Figure 13: Throughput versus number of turns along a path, for a path of length 8, where $K = 2500$, $r_s = 0.05$, and each of l and v are varied for System with an 8×8 unit square tessellation.

assumption that failures do not introduce deadlock scenarios—that the locking algorithm allows multi-color flows to mutual-exclusively take control of intersections (color-shared cells). Finally, under a fairness assumption, we established the main progress property through two results, that any cell gets permission to move infinitely often, and that any cell with a permission to move decreases the distance of any entities on it from its destination.

5. Simulation Experiments

We have performed several simulation studies of the algorithm for evaluating its throughput performance. In this section, we discuss the main findings with illustrative examples taken from the simulation results. We implemented the simulator in Matlab, and all the partition figures displayed in the paper are created using it. The algorithm and simulator are entirely deterministic, except for (a) the choice of the initial positions of new entities in source cells, and (b) failure and recovery rates in scenarios that consider random failures and recoveries, as detailed below.

Let the K -round *throughput* of System be the total number of entities arriving at the target over K rounds, divided by K . We define the *average throughput* (henceforth *throughput*) as the limit of K -round throughput for large K . All simulations start at a state where all cells are empty and subsequently entities are added to the source cells.

Single-color throughput without failures as a function of r_s , l , v . Rough calculations show that throughput should be proportional to cell velocity v , and inversely proportional to safety distance r_s and entity radius l . Figure 12 shows throughput versus r_s for several choices of v for an 8×8 unit square tessellation instance of System with a single entity color. The parameters are set to

$l = 0.25$ and $K = 2500$. The entities move along a line path where the source is the bottom left corner cell and the target is the top left corner cell. For the most part, the inverse relationship with v holds as expected: all other factors remaining the same, a lower velocity makes each entity take longer to move away from the boundary, which causes the predecessor cell to be blocked more frequently, and thus fewer entities reach tid from any element of ID_S in the same number of rounds. In cases with low velocity (for example $v = 0.1$) and for very small r_s , however, the throughput can actually be greater than that at a slightly higher velocity. We conjecture that this somewhat surprising effect appears because at very small safety spacing, the potential for safety violation is higher with faster speeds, and therefore there are many more blocked cells per round. We also observe that the throughput saturates at a certain value of r_s (≈ 0.55). This situation arises when there is roughly only one entity in each cell.

Single-color throughput without failures as a function of the path. For a sufficiently large number of rounds K , throughput is independent of the length of the path. This of course varies based on the particular path and instance of System considered, but all other variables fixed, this relationship is observed. More interesting however, is the relationship between throughput and path complexity, measured in the number of turns along a path. Figure 13 shows throughput versus the number of turns along paths of length 8. This illustrates that throughput decreases as the number of turns increases, up to a point at which the decrease in throughput saturates. This saturation is due to signaling and indicates that there is only one entity per cell.

Single-color throughput under failure and recovery of cells. Finally, we considered a random failure and recovery model in which at each round each non-faulty cell fails with some probability p_f and each faulty cell recovers with some probability p_r [33]. A recovery sets $failed_i = false$ and in the case of tid also resets $dist_{tid} = 0$, so that eventually $Route$ will correct $next_j$ and $dist_j$ for any $j \in TC$. Intuitively, we expect that throughput will decrease as p_f increases and increase as p_r increases. Figure 14 demonstrates this result for $0.01 \leq p_f \leq 0.05$ and $0.05 \leq p_r \leq 0.2$. There is a diminishing return on increasing p_r for a fixed p_f , in that for a fixed p_f increasing p_r results in smaller throughput gains.

Multi-color throughput as a function of the number of intersecting cells. Now we discuss the influence of multi-color throughput. In the case where the paths between different sources and targets do not overlap, all the results from the single-color simulation results apply. In the case where the paths do overlap, the mutual exclusion algorithm runs to ensure no deadlocks occur. This additional control logic will have an influence on the throughput. For the multi-color cases, we consider the summed throughput, which is the sum of the throughput for each color.

Figure 16 shows the roughly exponential decrease in throughput as the fraction of overlapping paths increases for two colors with path length 8 and no




Figure 14: Throughput versus failure rate p_f for several recovery rates p_r with an initial path of length 8, where $K = 20000$, $r_s = 0.05$, $l = 0.2$, and $v = 0.2$ for System with an 8×8 unit square tessellation.

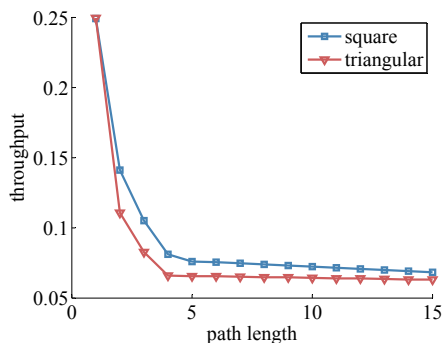


Figure 15: Throughput versus increasing path length of square (blue) and equilateral triangular (red) partitions.

turns. The fraction of overlapping paths is defined as the number of vertices in the color-shared cells $CSC(\mathbf{x}, c)$. As the fraction increases, the paths lie completely on top of one another, so in this case with path length 8, we have no overlap, 1 cell overlap, etc.

Multi-color throughput as a function of the number of intersecting colors. Intersections (that is, having at least one color-shared cell) have a fixed cost on throughput. Specifically, the summed throughput of there being two overlapping colors on a cell is the same as the summed throughput of three or more. Figure 17 shows this fixed decrease in summed throughput as the number of overlapping colors increases for a fixed path of length 3 with 3 color-shared cells, where the decrease in summed throughput from having no overlaps to having one color overlapping is about 4.5 times. Once there are two colors, all additional colors do not decrease the summed throughput. This observation agrees with intuition—the decrease in summed throughput due to an intersection is independent of the number of destinations for the entities that must pass through that intersection.

6. Related Work

There is a large amount of work on traffic control in transportation systems (see, e.g., [4, 34]) and robotics (see, e.g., [35]). We briefly summarize some of the more related work, but highlight that we are presenting a formal model of an example of such systems. Distributed air and automotive traffic control have been studied in many contexts. Human-factors issues are considered in [36, 37] to ensure collision avoidance between the coordination of numerous pilots and a supervisory controller modeling the semi-centralized air traffic control components. The Small Aircraft Transportation Protocol (SATS) is semi-distributed




Figure 16: Throughput versus fraction of path overlap for two colors on a 1×16 unit square tessellation.

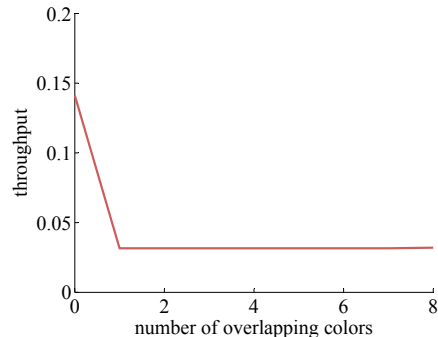


Figure 17: Throughput versus number of overlapping colors on a 1×3 unit square tessellation.

air traffic control protocol designed for small airports without radar, so pilots and their aircraft coordinate among themselves to land after being assigned a landing sequence order by an automated system at the airport [16]. SATS has been formally modeled and analyzed using a combination of model checking and automated theorem proving [38]. SATS and this paper share an abstraction: the physical environment is a priori partitioned into a set of regions of interest, and properties about the whole system are proved using compositional analysis. Safe conflict resolution maneuvers for distributed air traffic control are designed in [39]. A formal model of the traffic collision avoidance system (TCAS) is developed and analyzed for safety in [40]. TCAS is a system deployed on aircraft that alerts pilots when other aircraft are in close proximity and guides them along safe trajectories.

A distributed algorithm (executed by entities, vehicles in this case) for controlling automotive intersections without any stop signs is presented in [18]. Some methods for ensuring liveness for automotive intersections are presented in [41]. A method to detect the mode of a hybrid system control model of an autonomous vehicle in intersections is developed in [42], and is used to reduce conservatism of the maximally controlled invariant set (the set of collision-free controls). Efficient distributed intersection control algorithms are developed in [43]. There is a large amount of work on flocking [44] and platooning [45, 46, 47, 48]. Only a few works consider failures in such systems, like the arbitrary failures considered in [49, 50], the actuator failures considered in [48], or in synchronization of swarm robot systems in [51].

Distributed robot coordination on discrete abstractions like [52, 23, 53, 54, 55, 56, 57] can be viewed as traffic control. For instance, [23] establishes a formal connection between the continuous and the discrete parts of these protocols, and also presents a self-stabilizing algorithm with similar analysis to the analysis in this paper. These works also decompose the continuous problem into a discrete abstraction by partitioning the environment, but all these works allow at most a single entity (robot) in each partition, while our framework

allows numerous entities in each partition. If several entities are to visit some destination in [53, 56, 57], like our targets here, that destination is represented as the union of a set of partitions and each entity must reside in one of these partitions.

The Kiva Systems robotic warehouse [52] is a robotic traffic control system on square partitions, and can be described in our framework by allowing a single entity per cell. In these warehouse systems, there is a central coordinator scheduling tasks, but the robots are responsible for path planning using an A*-like search algorithm [52]. However, several deadlock scenarios are identified when performing such path planning [54]. The *Adaptive Highways Algorithm* presented in [54] for scheduling entities relies on using the tentative trajectories of other robots collected by the central controller. Deadlocks are also observed in other distributed robotics path-planning algorithms on discrete partitions in [58]. Deadlock scenarios can also arise without a discrete abstraction, such as in the doorways considered in [59], the path formation algorithms of [60], or the warehouse automation system of [61].

Lastly, we mention that most of these works on traffic control from aviation, automotive, swarm robotics, and warehouse automation applications can be modeled within the framework of spatial computing [62, 63, 64].

7. Discussion

In this section, we discuss some ways to generalize assumptions used in the paper and some alternative methods. In this paper, we presented a distributed traffic control algorithm for the partitioned plane, which moves entities without collision to their destinations, in spite of failures. While our algorithm is presented for two-dimensional partitions, an extension to some three-dimensional partitions (e.g., cubes and tetrahedra) follows in an obvious way. An extension to the more general case where there are multiple sources and multiple targets of each color—and entities of each color move toward the nearest target of that color—is straightforward, but complicates notation.

Self-Stabilizing Mutual Exclusion and Distributed Snapshot Algorithms. There are a variety of mutual exclusion algorithms that could be used to determine locks (Figure 9, line 13). For this paper, we require the overall system to be stabilizing and therefore the locking algorithm itself should be stabilizing. To this end, any of the following algorithms could be adapted to our framework: the token circulation algorithm [65], mutual exclusion [66], group mutual exclusion [67], snap-stabilizing propagation of information with feedback (PIF) algorithm [68], or k -out-of- l mutual exclusion [69]. A self-stabilizing distributed snapshot algorithm (see [27, Ch. 5]) can be used to determine if all c color-shared cells are empty, after having had some entity of color c (Figure 9, line 19). If all cells are empty, then another round of mutual exclusion commences, excluding color c from the input set.

General Triangulations and Affine Dynamics. We assumed in Section 2 that the partitions satisfy several geometric assumptions for feasibility of entity transfers. We considered using vector fields generated by a discrete abstraction like those presented in [70, 71, 72, 73, 74]. The affine vector fields generated on simplices in [70, 73] can be used to move an entity (with potentially non-holonomic or nonlinear dynamics) through any side of a cell in a triangulation (simplex) [70, 72] or rectangle [71]. However, it turns out that it is impossible to maintain our notion of safety for such vector fields without additional collision avoidance mechanisms implemented on each entity. This is due to a simple geometric observation—moving entities through a shorter side than the side they entered through may require the entities to come closer together. For example, if a cell in the triangulation has an obtuse angle, then the vector field generated by [70] flowing from the longest edge to the shortest edge has negative divergence. Furthermore, a vector field having negative divergence implies the flow corresponding to any two distinct points starting in that field come closer together, hence safety cannot be maintained. The distributed problems using these discrete abstractions [53, 56, 57] avoid this by requiring at most one entity in any (triangular) partition at a time.

We also mention a simple condition to ensure that triangulations have the required geometric partition properties (Assumptions 1 and 2). If all the triangles in the triangulation are non-obtuse, then the triangulation satisfies these assumptions. We also note that restricting allowable triangulations of an environment to ones without obtuse angles is not restrictive, since any polygon can be efficiently partitioned into a triangulation with non-obtuse [75, 76] or acute [77] angles.

Insufficiency of Disjoint Paths. Finding disjoint paths, such as by using the algorithms from [78, 79, 80, 81], could be another approach to solving the multi-color problem, but the locking mechanism used here solves a more general problem. Even without failures, there are many environments and choices of sources and targets for which there are no disjoint paths between sources and targets. One such environment is shown in Figure 4, where for two distinct colors c and d , the paths between the respective sources and targets necessarily overlap, so an algorithm for finding disjoint paths cannot be used as there are no disjoint paths between sources and targets. However, there are disjoint paths in some cases, so no scheduling would be necessary if these are found, but our routing algorithm does not necessarily find these, as the disjoint paths may not be shortest distance. A self-stabilizing algorithm for finding disjoint paths on planar graphs would be an enhancement to our algorithm, as it would increase throughput in the case that paths need not overlap.

Back-Pressure and Wormhole Routing. Back-pressure routing [82, 83] is an algorithm for dynamically routing traffic over an underlying graph using congestion gradients. If we view the color of each entity as its intended address and consider this problem from the perspective of queuing theory, one might think back-pressure routing could provide a throughput-optimal solution for




Figure 18: Hexagonal partition that does not satisfy the projection property (Assumption 1). An extension to allow such partitions would require enlarging the transfer region and receiving a signal from all of the potential next neighbors, which would require cells 3 and 7 both to signal cell 4 to move.




Figure 19: Example system on a parallelogram partition with failed cells in black. The several turns along the path from the source to the target cause a saturation of entities on cells 6, 10, and 14. The movement vector $u(i, j)$ is defined as the unit vector parallel to the x axis for movement between horizontal neighbors, and the unit vector parallel to the vertical sides of the parallelograms between vertical neighbors.

the problem. However, our physical motion model is incompatible with back-pressure routing. For a given cell, our model does not allow arbitrary choice of the next neighbor for each entity on that cell. In particular, when one cell moves its entities toward a neighboring cell, all entities sufficiently near the shared side between the two neighbors would transfer.

Wormhole routing [84] is a flow control policy over a fixed underlying graph for determining when packets move to the node on the graph. Addresses in wormhole routing are very short and come at the beginning of a packet, so a packet can be subdivided into pieces or *flits* and begin being forwarded after the address is received, yielding a snake-like sequence of flits in transfer. One could also view the sequence of entities on a path toward the appropriately-colored target (see Figure 19) sequence of flits flowing to a destination in wormhole routing. While similar deadlock scenarios can arise in our system and wormhole routing, wormhole routing is incompatible with our system due to the motion model just like back-pressure routing.

8. Conclusion

We presented a self-stabilizing distributed traffic control protocol for the partitioned plane, where each partition controls the motion of all entities within that partition. The algorithm guarantees separation between entities in the face of crash failures of the software controlling a partition. Once new failures cease occurring, it guarantees progress of all entities that are neither isolated by (a) failed partitions, nor (b) cells with entities of other colors that become deadlocked due to failures, to the respective targets. Through simulations, we presented estimates of throughput as a function of velocity, minimum separation, single-target path complexity, failure-recovery rates, and multi-target

path complexity.

It would be interesting to develop strategies allowing entities of different colors on a single cell. Our strategy of preventing entities of different colors from residing on a single cell simplified some analysis, but it also complicated some analysis, by making it harder to prove progress because deadlock scenarios may frequently arise. It would be interesting to develop algorithms allowing mixing and sorting of colors using different types of motion coupling. It would also be interesting to design algorithms that can allow relaxing the assumption on what failures may occur to ensure liveness. We believe this would require a more complex routing algorithm to temporarily move entities of some colors off the color shared cells, thus allowing some other color on the color shared cells to make progress.

9. Acknowledgments

The authors thank Zhongdong Zhu for helping develop the current version of the simulator, Karthik Manamcheri for helping develop an earlier version of the simulator, and Nitin Vaidya for helpful feedback. We also thank the anonymous reviewers who helped improve the earlier version of this paper.

References

- [1] D. Helbing, M. Treiber, Jams, waves, and clusters, *Science*
- [2] B. S. Kerner, Experimental features of self-organization in traffic flow, *Phys. Rev. Lett.*
- [3] C. Daganzo, M. Cassidy, R. Bertini, Possible explanations of phase transitions in highway traffic, *Transportation Research A*
- [4] M. Nolan, Fundamentals of air traffic control,
- [5] F. Borgonovo, L. Campelli, M. Cesana, L. Coletti, Mac for ad hoc inter-vehicle network: services and performance, in: *IEEE Vehicular Technology Conf.*, Vol. 5, 2003,
- [6] M. Karpiriski, A. Senart, V. Cahill, Sensor networks for smart roads, in: *Pervasive Computing and Communications Workshops, 2006. PerCom Workshops 2006. Fourth Annual IEEE International Conference on*, 2006,
- [7] S. S. Manvi, M. S. Kakkasageri, J. Pitt, Multiagent based information dissemination in vehicular ad hoc networks, *Mob. Inf. Syst.*
- [8] S. R. Azimi, G. Bhatia, R. R. Rajkumar, P. Mudalige, Vehicular networks for collision avoidance at intersections, *SAE International Journal of Passenger Cars - Mechanical Systems*

- [9] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L.-E. Jendrossek, C. Koelen, C. Markey, C. Rummel, J. van Niekerk, E. Jensen, P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, P. Mahoney, Stanley: The Robot That Won the DARPA Grand Challenge, in: M. Buehler, K. Iagnemma, S. Singh (Eds.), *The 2005 DARPA Grand Challenge*, Vol. 36 of Springer Tracts in Advanced Robotics, Springer Berlin / Heidelberg, 2007,
- [10] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y.-W. Seo, S. Singh, J. Snider, A. Stentz, W. R. Whittaker, Z. Wolkowicki, J. Ziglar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, D. Ferguson, Autonomous driving in urban environments: Boss and the urban challenge, *Journal of Field Robotics*
- [11] X. Yang, L. Liu, N. Vaidya, F. Zhao, A vehicle-to-vehicle communication protocol for cooperative collision warning, in: *Mobile and Ubiquitous Systems: Networking and Services. MOBIQUITOUS. The First Annual International Conference on*, 2004,
- [12] J. Misener, R. Sengupta, H. Krishnan, Cooperative collision warning: Enabling crash avoidance with wireless technology, in: *12th World Congress on Intelligent Transportation Systems*, 2005,
- [13] B. Hoh, M. Gruteser, R. Herring, J. Ban, D. Work, J.-C. Herrera, A. M. Bayen, M. Annavaram, Q. Jacobson, Virtual trip lines for distributed privacy-preserving traffic monitoring, in: *MobiSys '08: Proceeding of the 6th International Conference on Mobile Systems, Applications, and Services*, ACM, New York, NY, USA, 2008,
- [14] A. Girard, J. de Sousa, J. Misener, J. Hedrick, A control architecture for integrated cooperative cruise control and collision warning systems, in: *Decision and Control. Proceedings of the 40th IEEE Conference on*, Vol. 2, 2001,
- [15] M. Mamei, F. Zambonelli, L. Leonardi, Distributed motion coordination with co-fields: a case study in urban traffic management, in: *Autonomous Decentralized Systems. ISADS. The Sixth International Symposium on*, 2003,
- [16] T. S. Abbott, K. M. Jones, M. C. Consiglio, D. M. Williams, C. A. Adams, Small aircraft transportation system, higher volume operations concept: Normal operations, Tech. Rep. NASA/TM-2004-213022, NASA

- [17] M. Kelly, G. Di Marzo Serugendo, A decentralised car traffic control system simulation using local message propagation optimised with a genetic algorithm, in: S. Brueckner, S. Hassas, M. Jelasity, D. Yamins (Eds.), *Engineering Self-Organising Systems*, Vol. 4335 of *Lecture Notes in Computer Science*, Springer, 2007,
- [18] H. Kowshik, D. Caveney, P. R. Kumar, Safety and liveness in intelligent intersections, in: *Hybrid Systems: Computation and Control (HSCC)*, 11th International Workshop, Vol. 4981 of *LNCS*, 2008,
- [19] K. Dresner, P. Stone, A multiagent approach to autonomous intersection management, *Journal of Artificial Intelligence Research*
- [20] P. Weiss, Stop-and-go science, *Science News*
- [21] Kornylak, [Omniwheel brochure](http://www.kornylak.com/images/pdf/omni-wheel.pdf)
URL <http://www.kornylak.com/images/pdf/omni-wheel.pdf>.
- [22] K. An, A. Trewyn, A. Gokhale, S. Sastry, Model-driven performance analysis of reconfigurable conveyor systems used in material handling applications, in: *Cyber-Physical Systems (ICCPs)*, 2011 IEEE/ACM International Conference on, Vol. 2, IEEE, 2011,
- [23] S. Gilbert, N. Lynch, S. Mitra, T. Nolte, Self-stabilizing robot formations over unreliable networks, *ACM Trans. Auton. Adapt. Syst.*
- [24] S. Dolev, L. Lahiani, S. Gilbert, N. Lynch, T. Nolte, Virtual stationary automata for mobile networks, in: *PODC '05: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, ACM, New York, NY, USA, 2005,
- [25] T. Nolte, N. Lynch, A virtual node-based tracking algorithm for mobile networks, in: *Distributed Computing Systems, International Conference on (ICDCS)*, IEEE Computer Society, Los Alamitos, CA, USA, 2007,
- [26] A. Arora, M. Gouda, Closure and convergence: A foundation of fault-tolerant computing, *IEEE Trans. Softw. Eng.*
- [27] S. Dolev, *Self-stabilization*,
- [28] S. M. Loos, A. Platzer, L. Nistor, Adaptive cruise control: Hybrid, distributed, and now formally verified, in: M. Butler, W. Schulte (Eds.), *Formal Methods*, LNCS,
- [29] A. Platzer, Quantified differential invariants, in: *Proc. of the 14th ACM Intl. Conf. on Hybrid Systems: Computation and Control*, ACM, 2011,
- [30] T. T. Johnson, S. Mitra, Parameterized verification of distributed cyber-physical systems: An aircraft landing protocol case study, in: *ACM/IEEE 3rd International Conference on Cyber-Physical Systems*,

- [31] T. T. Johnson, S. Mitra, A small model theorem for rectangular hybrid automata networks, in: Proceedings of the IFIP International Conference on Formal Techniques for Distributed Systems, Joint 14th Formal Methods for Open Object-Based Distributed Systems and 32nd Formal Techniques for Networked and Distributed Systems (FORTE-FMOODS), Vol. 7273 of LNCS,
- [32] T. T. Johnson, S. Mitra, K. Manamcheri, Safe and stabilizing distributed cellular flows, in: Proceedings of the 30th IEEE International Conference on Distributed Computing Systems (ICDCS),
- [33] R. E. L. DeVille, S. Mitra, Stability of distributed algorithms in the face of incessant faults, in: Proceedings of 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS), Springer, 2009,
- [34] P. A. Ioannou, Automated Highway Systems,
- [35] F. Bullo, J. Cortés, S. Martínez, Distributed Control of Robotic Networks, Applied Mathematics Series, Princeton University Press, 2009,
- [36] N. Leveson, M. de Villepin, J. Srinivasan, M. Daouk, N. Neogi, E. Bachelder, J. Bellingham, N. Pilon, G. Flynn, A safety and human-centered approach to developing new air traffic management tools, in: Proceedings Fourth USA/Europe Air Traffic Management R&D Seminar, 2001,
- [37] T. Prevot, Exploring the many perspectives of distributed air traffic management: The multi aircraft control system (macs), in: Proceedings of the HCI-Aero, 2002,
- [38] C. Muñoz, V. Carreño, G. Dowek, Formal analysis of the operational concept for the small aircraft transportation system, in: M. Butler, C. Jones, A. Romanovsky, E. Troubitsyna (Eds.), Rigorous Development of Complex Fault-Tolerant Systems, Vol. 4157 of LNCS, Springer Berlin / Heidelberg, 2006,
- [39] C. Tomlin, G. Pappas, S. Sastry, Conflict resolution for air traffic management: a study in multiagent hybrid systems, IEEE Trans. Autom. Control
- [40] C. Livadas, J. Lygeros, N. A. Lynch, High-level modeling and analysis of TCAS, in: Proceedings of the 20th IEEE Real-Time Systems Symposium (RTSS'99), 1999,
- [41] T.-C. Au, N. Shahidi, P. Stone, Enforcing liveness in autonomous traffic management, in: Proceedings of the Twenty-Fifth Conference on Artificial Intelligence,
- [42] R. Verma, D. Vecchio, Semiautonomous multivehicle safety, Robotics Automation Magazine, IEEE

- [43] A. Colombo, D. D. Vecchio, Efficient algorithms for collision avoidance at intersections, in: *Hybrid Systems: Computation and Control (HSCC)*,
- [44] R. Olfati-Saber, Flocking for multi-agent dynamic systems: algorithms and theory, *IEEE Trans. Autom. Control*
- [45] P. Varaiya, Smart cars on smart roads: Problems of control, *IEEE Trans. Autom. Control*
- [46] E. Dolginova, N. Lynch, Safety verification for automated platoon maneuvers: A case study, in: *HART'97 (International Workshop on Hybrid and Real-Time Systems)*, Vol. 1201 of LNCS,
- [47] D. Swaroop, J. K. Hedrick, Constant spacing strategies for platooning in automated highway systems, *Journal of Dynamic Systems, Measurement, and Control*
- [48] T. T. Johnson, S. Mitra, Safe flocking in spite of actuator faults using directional failure detectors, *Journal of Nonlinear Systems and Applications*
- [49] V. Gupta, C. Langbort, R. Murray, On the robustness of distributed algorithms, in: *Decision and Control. 45th IEEE Conference on*, 2006,
- [50] M. Franceschelli, M. Egerstedt, A. Giua, Motion probes for fault detection and recovery in networked control systems, in: *American Control Conference*, 2008, 2008,
- [51] A. Christensen, R. O'Grady, M. Dorigo, From fireflies to fault tolerant swarms of robots, *IEEE Transactions on Evolutionary Computation*
- [52] P. R. Wurman, R. D'Andrea, M. Mountz, Coordinating hundreds of cooperative, autonomous vehicles in warehouses, *AI Magazine*
- [53] M. Kloetzer, C. Belta, Automatic deployment of distributed teams of robots from temporal logic motion specifications, *Robotics, IEEE Transactions on*
- [54] H. Roozbehani, R. D'Andrea, Adaptive highways on a grid, in: C. Pradalier, R. Siegwart, G. Hirzinger (Eds.), *Robotics Research*, Vol. 70 of *Springer Tracts in Advanced Robotics*, Springer, 2011,
- [55] J. W. Durham, R. Carli, P. Frasca, F. Bullo, Discrete partitioning and coverage control for gossiping robots, *Robotics, IEEE Transactions on*
- [56] X. C. Ding, M. Kloetzer, Y. Chen, C. Belta, Automatic deployment of robotic teams, *Robotics Automation Magazine, IEEE*
- [57] Y. Chen, X. C. Ding, A. Stefanescu, C. Belta, Formal approach to the deployment of distributed robotic teams, *Robotics, IEEE Transactions on*

- [58] R. Luna, K. Bekris, Network-guided multi-robot path planning in discrete representations, in: Intelligent Robots and Systems (IROS). IEEE/RSJ International Conference on, 2010,
- [59] D. Herrero-Perez, H. Matinez-Barbera, Decentralized coordination of autonomous agvs in flexible manufacturing systems, in: Intelligent Robots and Systems. IROS. IEEE/RSJ International Conference on, 2008,
- [60] S. Nouyan, A. Campo, M. Dorigo, Path formation in a robot swarm: Self-organized strategies to find your way home, *Swarm Intelligence*
- [61] A. Kamagaew, J. Stenzel, A. Nettstrater, M. ten Hompel, Concept of cellular transport systems in facility logistics, in: Automation, Robotics and Applications (ICARA). 5th International Conference on, 2011,
- [62] F. Zambonelli, M. Mamei, Spatial computing: An emerging paradigm for autonomic computing and communication, in: M. Smirnov (Ed.), *Autonomic Communication*, Vol. 3457 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2005,
- [63] J. Beal, J. Bachrach, Infrastructure for engineered emergence on sensor/actuator networks, *Intelligent Systems*, IEEE
- [64] J. Bachrach, J. Beal, J. McLurkin, Composable continuous-space programs for robotic swarms, *Neural Computing & Applications*
- [65] C. Johnen, G. Alari, J. Beauquier, A. Datta, Self-stabilizing depth-first token passing on rooted networks, in: M. Mavronicolas, P. Tsigas (Eds.), *Distributed Algorithms*, Vol. 1320 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 1997, pp. 260–274,
- [66] A. K. Datta, C. Johnen, F. Petit, V. Villain, Self-stabilizing depth-first token circulation in arbitrary rooted networks, *Distributed Computing*
- [67] J. Beauquier, S. Cantarell, A. Datta, F. Petit, Group mutual exclusion in tree networks, in: *Parallel and Distributed Systems*, 2002. Proceedings. Ninth International Conference on, IEEE Computer Society, 2002,
- [68] A. Bui, A. Datta, F. Petit, V. Villain, Snap-stabilization and pif in tree networks, *Distributed Computing*
- [69] A. Datta, S. Devismes, F. Horn, L. Larmore, Self-stabilizing k-out-of-1 exclusion on tree networks, in: *Parallel Distributed Processing*, 2009. IPDPS 2009. IEEE International Symposium on, 2009,
- [70] C. Belta, V. Isler, G. Pappas, Discrete abstractions for robot motion planning and control in polygonal environments, *Robotics*, IEEE Transactions on
- [71] C. Belta, L. Habets, Controlling a class of nonlinear systems on rectangles, *Automatic Control*, IEEE Transactions on

- [72] L. Habets, P. Collins, J. van Schuppen, Reachability and control synthesis for piecewise-affine hybrid systems on simplices, *Automatic Control, IEEE Transactions on*
- [73] M. Kloetzer, C. Belta, A fully automated framework for control of linear systems from temporal logic specifications, *Automatic Control, IEEE Transactions on*
- [74] G. E. Fainekos, A. Girard, H. Kress-Gazit, G. J. Pappas, Temporal logic motion planning for dynamic robots, *Automatica*
- [75] B. Baker, E. Grosse, C. Rafferty, Nonobtuse triangulation of polygons, *Discrete & Computational Geometry*
- [76] M. Bern, S. Michell, J. Ruppert, Linear-size nonobtuse triangulation of polygons, *Discrete & Computational Geometry*
- [77] H. Maehara, Acute triangulations of polygons, *European Journal of Combinatorics*
- [78] H. Mohanty, G. P. Bhattacharjee, A distributed algorithm for edge-disjoint path problem, in: *Proceedings of the Sixth Conference on Foundations of Software Technology and Theoretical Computer Science, Springer-Verlag, London, UK, UK, 1986,*
- [79] R. Ogier, V. Rutenburg, N. Shacham, Distributed algorithms for computing shortest pairs of disjoint paths, *Information Theory, IEEE Transactions on*
- [80] S.-J. Lee, M. Gerla, Split multipath routing with maximally disjoint paths in ad hoc networks, in: *Communications. ICC. IEEE International Conference on, Vol. 10, 2001,*
- [81] M. Marina, S. Das, On-demand multipath distance vector routing in ad hoc networks, in: *Network Protocols. Ninth International Conference on, 2001,*
- [82] L. Tassiulas, A. Ephremides, Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks, *Automatic Control, IEEE Transactions on*
- [83] B. Awerbuch, T. Leighton, A simple local-control approximation algorithm for multicommodity flow, in: *Foundations of Computer Science. Proceedings., 34th Annual Symposium on, IEEE, 1993,*
- [84] L. Ni, P. McKinley, A survey of wormhole routing techniques in direct networks, *Computer*