
ON THE DEVELOPMENT OF THE
VORONOI-TOPOLOGICAL-PERCEPTION MODEL FOR
COLLECTIVE BEHAVIOR AND ITS GENERALIZATIONS

Jack Tisdell

Department of Mathematics and Statistics, McGill University, Montreal

November 2021

A thesis submitted to McGill University in partial fulfillment of the requirements of the
degree of Master of Science

© Jack Tisdell, 2021

Contents

1	Acknowledgements	3
2	Contribution to original knowledge	3
3	Contribution of authors	4
4	Introduction	5
5	Background	6
6	VTP model	13
6.1	Overview and motivation	13
6.2	Preliminaries	13
6.3	Governing equations	14
6.3.1	Repulsion/attraction	18
6.3.2	Alignment	19
6.3.3	Homing	20
6.3.4	Weighting coefficients	21
6.3.5	Personal-space speed scale	22
6.4	Effective parameters	25
6.5	Note on a related model	26
6.6	Observables	28
6.7	Results	30
6.7.1	Homing-free systems in domains without boundary	30
6.7.2	Pinwheels, rings, and cogs with point-target homing	36
6.8	Introducing sources and sinks	40
6.8.1	Source algorithm	40
6.8.2	The bi-directional corridor	41

7	The directed Delaunay network	43
7.1	Motivation: Obstacles and VTP in non-convex domains	43
7.2	Construction and computational methods	45
7.3	The DDN dual diagram and VTP implementation	52
7.4	Summary	55
A	Lemmas	56
B	Technicalities concerning vector translations in the torus	60
C	Alternate metrics	63
C.1	Ray-marching algorithm for computing approximate Voronoi diagrams in arbitrary metrics	63

Abstract

We motivate and develop a new model, Voronoi Topological Perception (VTP), for the dynamics of multi-agent biological systems such as pedestrian crowds, bird flocks, mammal herds, etc. All interactions predicted by the model arise naturally from consideration of the Voronoi tessellation and its dual topological structure. We investigate a wide variety of collective behaviors predicted by VTP in its simplest form, then after introducing targets, and further introducing sources and sinks. Additionally, we introduce the directed Delaunay network (DDN), a generalization of the Voronoi diagram which extends VTP to much wider contexts, and develop computational methods for this new structure.

Nous motivons et présentons un nouveau modèle, dit de Perception Topologique de Voronoi (VTP), appliqué à l'étude de la dynamique d'agents biologiques tels que: des foules piétonnes, des ensembles d'oiseaux, des troupeaux de mammifères, etc. Toutes les interactions prédites par notre modèle surviennent naturellement en considérant les pavages de Voronoi ainsi que leur structure duale; à savoir, la triangulation de Delaunay. Nous étudions une grande variété de comportements collectifs prédits par le modèle

VTP; dans un premier temps sous sa forme la plus simple et de façon subséquente nous y introduisons des cibles. Enfin nous considérons des scénarios d'entrée et de sortie d'agents. En sus nous développons le Réseau Dirigé de Delaunay (DDN), ceci est une généralisation du diagramme de Voronoi qui permet d'étendre le formalisme de VTP à des contextes plus généraux. Finalement nous présentons des méthodes computationnelles pour cette nouvelle structure.

1 Acknowledgements

I would like to thank my supervisors, Rustum Choksi and Jean-Christophe Nave for their constant support and encouragement. Also, I would like to thank Ivan Gonzalez, with whom I collaborated on much of the work in the first sections and whose earlier work in variational Voronoi problems inspired this project. Specifically credit is owed to Gonzalez for implementation of the simulations appearing in Section §6.

Ivan Gonzalez and I are first authors on a publication resulting from the work in Sections §6.2, §6.3, §6.4, and §6.7. The preprint is available here¹.

2 Contribution to original knowledge

The VTP model is original and is not a direct extension/generalization of an existing model. Moreover, the model incorporates the topological and geometric structure of the Voronoi diagram in a novel manner. The model's repulsion interaction is based on the Voronoi optimization techniques developed in [21] but its use of in the context of systems of self-directed biological agents is new. VTP predicts some collective behaviors not witnessed by other models such as co-rotating (anti-cog) mills, the spherical anti-cog analog, and mills of oscillating radius. We design a cheap, geometrically motivated computational method for modelling multi-agent systems with sources. Finally, we introduce the directed Delaunay

¹<https://arxiv.org/abs/2111.03448>

network, a generalization of the Voronoi diagram (well suited for the multi-agent systems application), and a method for its computation.

3 Contribution of authors

The author (Tisdell) conducted the research of Sections 5, 6, 7, and all appendices under the direction of Rustum Choksi and Jean-Christophe Nave. The MACN repulsion interaction described in Section 6.3.1 was inspired by the work of Ivan Gonzalez in [21] on optimal Voronoi tessellations. Credit is owed to Gonzalez as well for the MatLab implementation of the VTP model from which the simulations are generated. Gonzalez and the author together designed and implemented (the code which produced) Figures 4, 5, 6, 7, and 8. The work of Section 6 is featured in a publication² co-first-authored by Tisdell and Gonzalez.

²Preprint available at <https://arxiv.org/abs/2111.03448>

4 Introduction

The problem of modelling collective herding, schooling, and swarming behaviors in biological systems as well as traffic in pedestrian and animal crowds has seen myriad approaches over the last decades in many areas. The two- or three-dimensional problem of swarming presents vast challenges over the one-dimensional and quasi-one-dimensional settings of relatively well understood vehicular traffic patterns. The organization in multi-agent biological systems tends to emerge from intricate mutual interactions and individual observations, with decisions occurring simultaneously across multiple levels of control. Moreover, experimental tools for gathering data about real crowds and swarms tend to be highly specialized for certain species and environments. For these reasons and others, it is generally very difficult to produce models which exhibit universal features of multi-agent systems but which also are sufficiently complex so as to reproduce faithfully empirically observed behaviors.

Our task is to present a framework, dubbed *Voronoi Topological Perception* (VTP), capable of synthesizing the complex microscopic interactions governing multi-agent systems in a natural manner consistent with recent experimental findings. The ubiquity of Voronoi diagrams and Delaunay triangulations in biology, engineering, computing and many other sciences makes the structure an extremely compelling candidate for a universal approach to the crowd problem. Building atop these structures, the VTP framework uses their inherent geometric and topological properties whenever possible to capture the dynamics of multi-agent systems, giving rise to a class of models which is highly amenable to generalization and equally adaptable to specialization.

We first review existing crowd models and discuss important classifications among them. We then present the basic VTP model in detail. Equipped with the model, we devote much discussion to simulations which exhibit some of the variety of collective behaviors captured by VTP in its simplest form and describe the full phase space for its most bare-bones version. The rest of the main text is devoted to generalizing VTP to more sophisticated settings which arise naturally in the crowd problem. Firstly among these is to introduce sources and sinks,

moving from the paradigm of swarming properly into that of traffic (a distinction we will elaborate on later). Here too, we reinforce the idea that every aspect of a universal model should arise from the geometry and topology of the Voronoi diagram. Finally, we move to the crucial context of non-convex domains, introducing and motivating a new structure which generalizes the Voronoi diagram in a way suited for multi-agent systems.

5 Background

Throughout the literature, there are many approaches to the crowd modelling problem. To classify them, we must distinguish several salient features:

discrete/continuous in space and time Continuous-time models are characterized by evolution according to differential equations while discrete time models evolve in successive iterations. Several authors have introduced both discrete- and continuous-time versions of models [11, 29, 13]. Continuous-space models allow their dynamic entities to vary over all locations in a continuous domain while discrete-space models, often called cellular automata or lattice-gas models, restrict this to a fixed, discrete set of points. for more details, see [7, 16, 33] and more recently [8].

discrete/continuous in population Continuous-population models—or more succinctly, continuum models—describe the crowd itself as a continuously distributed object, i.e., a density, which is subject to evolution according to the model’s dynamical equations. Continuum models are well suited for very large crowds for which the fine-scale interactions are unimportant. One such fluid-dynamic continuum theory was put forward in [25, 26, 27]. Discrete-population or agent-based models, on the other hand, model each individual as a dynamic entity subject to its own evolution according to the dynamical equations. We will discuss several such models below. In the case of continuum models, “collective behavior” is intrinsic to the dynamics while in the agent-based case, it is an emergent phenomenon. The connection between discrete and continuum models

has been investigated in [24, 31, 20] and others.

deterministic/stochastic Completely deterministic models are those whose dynamical equations predict entirely the forward evolution of the system from its initial conditions. Models may incorporate stochasticity in many ways. Although it is natural to introduce probabilistic effects in pursuit of realistic dynamics, the exact role of stochastic effects can be unclear.

flocking/traffic There are two distinct phenomena which one might lump together under the umbrella of crowd modelling but whose interests conflict in some ways. The first is *flocking* or *swarming* where a particular fixed population evolves over time. The second, which we might call *traffic*, is where a particular environment is over time traversed by populations. Clearly, there is not a sharp distinction between these notions; a fixed population confined to a fixed environment properly belongs to both and unbounded environments suggests a grey area. Nonetheless, this distinction can greatly inform one's modelling choices. For instance, if one wishes to model migratory flocks of birds, it is essential that the flock diameter does not become arbitrarily large but if one wants to model automobile traffic at a large intersection, vehicles are completely forgotten the moment they exit the intersection area.

kinematic/dynamic Kinematic models describe motion without regard for its *physical* cause while dynamic models account for the physical constraints on the system. Important examples of kinematic models include [13, 37, 4]. Dynamic models often satisfy underlying physical principles like conservation laws but complex systems like crowds may be better understood in terms of *behavioral* or *social*, rather than physical, causes.

Vicsek et. al. put forward a seminal discrete-time model in [39] achieving a transition to coherent behavior. Agents in the Vicsek model move with constant speed at each step in directions obtained by averaging the directions of all agents in a disc of radius r . This

direction is then perturbed by some random noise.

$$x_i(t + \tau) = x_i(t) + \tau u_i(t), \quad (1)$$

$$\angle u_i(t + \tau) = \langle \angle u_j(t) \rangle_{j: \|r_{ij}\| < r} + \eta \xi_i(t), \quad \|u_i\| = \bar{u} \quad (2)$$

where $r_{ij} = x_j - x_i$ and the angle brackets denote an average. The speed \bar{u} and noise magnitude η are fixed and the random term ξ_i is drawn uniformly from $[-\pi, \pi]$. The original Vicsek paper, with over 6000 citations, is extremely influential and we will comment on some of its extensions momentarily. Analysis of Vicsek's model and other similar models appears in [28]. The Vicsek model demonstrated emergence of decentralized coordination of even an extremely simple multi-agent system.

Cucker and Smale developed in [11] both discrete- and continuous-time versions of a flocking model (C-S) and, like Vicsek, demonstrate a transition to coherent dynamics where agents converge to a common velocity. In C-S, every agent interacts with every other agent, that is, the communication graph is complete (and thus also fixed). The continuous C-S model has the form

$$\dot{x}_i = u_i \quad (3)$$

$$\dot{u}_i = \frac{1}{n} \sum_j \phi_{ij}(u_j - u_i) \quad (4)$$

where ϕ_{ij} is a pair-wise interaction function depending on the spatial separation between i and j . They show the long-term convergence of the flock depends only on initial conditions. The momentum of the C-S flock is conserved due to its symmetric pair-wise interactions and complete communication graph. This is easily seen, the velocities u_i evolve according to equations of the form $\dot{u}_i = \frac{1}{n} \sum_{j \leq n} \phi_{ij}(u_j - u_i)$ where crucially, $\phi_{ij} = \phi_{ji}$ for all i, j . So the total momentum satisfies $\frac{d}{dt}(\frac{1}{n} \sum_i u_i) = \frac{1}{n} \sum_i \dot{u}_i = 0$.

Among the models following C-S is that introduced by Motsch and Tadmor in [32],

hereafter referred to as M-T. M-T addresses a crucial drawback of C-S: because the influence on each agent in C-S is normalized by the total number of agents n , the dynamics of small subflocks in non-equilibrium configurations are essentially halted simply by the presence of other distant, large subflocks. To solve this, M-T leverages the following insight. Introducing a notion of *relative distance*, while breaking the symmetry of interactions, allows M-T to model inhomogeneous flocks. Their model takes the form

$$\dot{x}_i = u_i, \tag{5a}$$

$$\dot{u}_i = \sum_j a_{ij}(u_j - u_i), \quad a_{ij} = \frac{\phi_{ij}}{\sum_k \phi_{ik}} \tag{5b}$$

where again ϕ_{ij} is a pair-wise interaction depending on $\|x_j - x_i\|$. Through this modification, M-T removes the normalization by the number of agents n from the velocity-averaging interactions. In the homogeneously distributed case, this effectively reduces to C-S but in the inhomogeneous case, it allows rich dynamics in small subflocks. We will see later a family resemblance to the alignment interaction in the model introduced here. This insight is further supported experimentally by the work of [2] establishing the density-invariance of many flocking behaviors in birds.

Much work has been done to model biological behaviors beyond mere velocity consensus including, e.g., collision avoidance, aggregation, pattern formation, goal seeking, navigation, leadership effects, and predator/prey interactions. One collective formation of special interest is milling, or vortex, motion. Milling has been observed in many animal species across a wide range of scales (see [12] for a broad survey) and in many cases is poorly understood. One influential model exhibiting spontaneous milling is that of D'Orsogna et. al. [13]. Their equations of motion take the form

$$\dot{x}_i = u_i, \tag{6}$$

$$\dot{u}_i = \alpha u_i - \beta \|u_i\|^2 u_i - \nabla U_i, \quad U_i = \sum_{j \neq i} (-C_a \exp(-\|r_{ij}\|/l_a) + C_r \exp(-\|r_{ij}\|/l_r)) \tag{7}$$

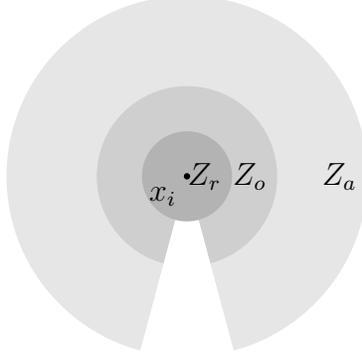


Figure 1: Schematic of the zones of interaction in the model [10]. From the inside, the interaction in each zone is repulsion, alignment (or orientation), and attraction.

where $r_{ij} = x_j - x_i$. Their model produces spontaneous milling formations of various kinds including double, interlocking mills and swarm vorticity of fixed angular velocity, as well as other types of collective formations.

The class of zone-based models shares a common approach to these more sophisticated behaviors. Specifically, the type of influence on agent i by an agents j is determined by which of several concentric annular zones about i 's position j falls in. The familiar boids model due to Reynolds [36] is one such zone-based model with repulsive interactions occuring at short range, (Vicsek-esque) alignment interactions at intermediate range, and attractive influences at long range. Bernardi and Scianna [3] use the zone approach in the context of predator/prey behavior. The influential model due to Couzin et. al. [10] incorporates overlapping zones, with the interaction types having a heirarchy within the overlap. In the discrete-time model, each agent steers in the direction d_r , d_o , or d_a accordingly as there are other agents within the concentric spherical shell zones Z_r , Z_o , and Z_a of repulsion, alignment (or orientation), and attraction, repectively, centered at x_i (shown schematically in Figure 1) where

$$d_r(t + \tau) = - \sum_{j:x_j \in Z_r} \frac{r_{ij}}{\|r_{ij}\|}, \quad d_o(t + \tau) = \sum_{j:x_j \in Z_o} \frac{u_j}{\|u_j\|}, \quad d_a(t + \tau) = \sum_{j:x_j \in Z_a} \frac{r_{ij}}{\|r_{ij}\|}. \quad (8)$$

The actual veocity updates are additionally subject to a maximum turning rate restriction

and a Vicsek-like Gaussian noise. All agents are blind to a cone behind them—in the sense that the volume of this cone is excluded from the zones Z_r , Z_o , and Z_a —whose interior angle is a parameter of the model. This model spontaneously produces four qualitative behaviors including milling through the volume of a torus (this occurs when the width of the zone of attraction is large compared to that of alignment). This model was used in [9] to study leadership effects and decision-making in mostly ignorant groups. In this setting, a subset of *informed* agents steer toward a preferred direction with some fixed weight. The Couzin model directly builds upon Vicsek’s in several ways. Like all zone-based models, the spherical shells generalize Vicsek’s disc of influence to incorporate a hierarchy of distinct types of interactions. Also in keeping with Vicsek, agents in Couzin’s model at constant speed and only update their heading at each time step.

One vital distinction among the above models and algorithms is whether the dynamical equations are first- or second-order in time. By second-order, we mean, specifically, in the sense of Newton’s second law where *accelerations* are determined by *forces*. The problem, in our view, is applying the Newtonian paradigm to systems of cognitive, decision-making agents when the relevant stimuli are not forces. Put another way, a cognitive agent may self-determine its immediate motion without consideration for *how* it will physically execute the motion. Thus, a model which seeks to describe the multi-agent system in terms of these individual, self-determined motions need not, and indeed should not, describe motions at a lower level of control. Of course, the high-level description must be physically feasible but this can be ensured through constraints such as maximum speeds, maximum turning angles (as in the Couzin model), sufficiently long time steps, etc. To drive this point home, a high-level, agent-based model of a biological system composed of self-determined individuals is much more closely analogous to a high-level control component of a robotics system with a strictly hierarchical control structure than to a physical description of interacting Newtonian bodies.

The use of Voronoi diagrams to derive dynamical influences in agent-based models has

been explored in various contexts. Voronoi diagrams may enter a dynamical model chiefly in two ways; (1) agents may in response to the *geometry* of their Voronoi cells or (2) the *topology* of the Voronoi diagram may determine the communication graph over the agents. Notably, the latter approach was shown by Ginelli and Chaté in [19]—also influenced by the aforementioned experiments of [2]—to produce qualitatively novel dynamics when applied straight-forwardly as a “topological” version of Vicsek’s model, with agents aligning with their Voronoi neighbors rather than those in a fixed-radius zone. They suggest, therefore, that, especially in light of recent experimental work, realistic modeling of biological multi-agent systems may necessarily incorporate *topological*, rather than metrical, interactions. This model is extended in [22] to include repulsive body forces.

By contrast, the vehicle coordination algorithm introduced by Lindhe et. al. in [30] only considers neighbors within a fixed radius, like Vicsek, but from this data derives a Voronoi region whose *geometry* is used to control the vehicle. They show that under this algorithm, every vehicle makes steady progress toward the (common) goal and the flock converges to a lattice formation. This particular scheme seems beyond the measurement and computational capacity of biological agents but nonetheless demonstrates the idea of deriving some individual behavior via Voronoi geometry.

Voronoi diagrams are commonplace in the literature of epithelia and soft tissues. In [6, 5], Bi et. al. use a Voronoi model to explain the collective solid/liquid phase transition in biological tissues. Often of interest in this context is the so-called shape index of the cells/configuration. This geometric property is closely related to quantization error which we introduce for our model as a system observable.

6 VTP model

6.1 Overview and motivation

The dynamics are governed by a system of delay difference equations. However, the delay represents a kind of inertia in the dynamics rather than memory in the agents. This is worth emphasizing: despite the delay in the model, physical agents do not require any memory whatsoever under plausible assumptions. In our model, agents only need access to the *direction* information of other agents. We assume that the species of agents can infer accurate directional information from the body geometry of other agents in a computationally trivial amount of time. We further assume that agents’ direction of motion reliably corresponds to this body geometry—i.e., agents face forward when moving—so that for the featureless point particles representing agents in the model, their “facing direction” is given exactly by their normalized velocities. Thus, employing velocities in this way introduces delay into the equations without requiring memory of the “real” agents. To put it another way, say our agents are experimental physicists, carrying with them the usual assortment of light clocks, meter sticks, and so on as they are wont to do and imagine each physicist is holding a compass pointing in the direction she’s moving. At all times, the physicists can see nearby physicist’s compasses and obtain direction information *instantaneously* without ever measuring each other’s velocities.

6.2 Preliminaries

Definition 6.2.1. Let Ω be a compact, connected, metric space with (extended) metric d . For a finite set of points $X = \{x_1, \dots, x_n\}$ in Ω , the *Voronoi cell* V_i generated by x_i (with respect to d) is the set

$$V_i = \{x \in \Omega : d(x, x_i) < \infty \text{ and } d(x, x_i) \leq d(x, x_j) \text{ for all } 1 \leq j \leq n\}.$$

The d -Voronoi diagram \mathcal{V} on Ω generated by X is the set of Voronoi cells $\mathcal{V} = \{V_1, \dots, V_n\}$. Usually, the metric d is clear from context and we write simply the *Voronoi cell/diagram*. Two generators x_i and x_j whose Voronoi cells have nonempty intersection are said to be *Voronoi neighbors* (or simply neighbors, if clear from context). For every pair of neighbors x_i and x_j , each connected component of $V_i \cap V_j$ is a *Voronoi edge*. Every point in Ω belonging to three or more Voronoi cells is a *Voronoi vertex*.

Certain properties of Voronoi diagrams are readily apparent. If d is real-valued (i.e., never taking value ∞), then $\bigcup \mathcal{V} = \Omega$ for every possible set of generators. Also apparent is the fact that for any set of generators and any extended metric, the interiors of Voronoi cells are pairwise disjoint. There is a dual structure to the Voronoi diagram.

Definition 6.2.2. If \mathcal{V} is the Voronoi diagram in Ω generated by X , its *dual* \mathcal{V}^* is the pseudograph over X which has an edge $\{x_i, x_j\}$ for every Voronoi edge in \mathcal{V} adjacent to V_i and V_j . If d is the Euclidean metric, \mathcal{V}^* is called the *Delaunay triangulation*.

Note that in general, \mathcal{V}^* may include multiple-edges and loops (usually as a consequence of periodic geometry of the domain Ω). There is a geometric notion of the Delaunay triangulation but only its graph structure is relevant to the work that follows.

Generalizing the notions of the Voronoi diagram and its dual in a manner suitable for multi-agent systems in and more general context is the focus of Section 7. For now, we will restrict our attention to convex domains with and without boundary where the standard notions serve us well.

6.3 Governing equations

We set out to design an model for pedestrian traffic using an agent-based approach, each agent having only local information about the state of the system. However, by local we mean with respect to the structure of the Voronoi diagram generated by the positions of all agents rather than metric distances. We refer to this framework as *topological perception*.

The other crucial feature of our model is that it is *first-order in time*. This is for two reasons. The system of delay difference equations which govern the dynamics relate, on one hand, the factors which influence “intellegent” agents’ *descisions* about their motion and, on the other, their achieved update. Pedestrians do not think in terms of accelerations (their own or others’) and do not make descisions based on this kind of second order information about the state of the system. Moreover, unlike say automobiles, human pedestrians can break and accellerate almost instantaneously [1, pp. 409], so it is perfectly sensisble for agents achieve arbitrary accellerations provided they are never moving implausibly fast.

Our model includes

1. a compact connected domain Ω ,
2. a finite index set Λ of agents,
3. positions $x_i = x_i(t) \in \Omega$ for each $i \in \Lambda$ such that $x_i(t) \neq x_j(t)$ for all t and for all $j \neq i$,
and
4. closed (possibly empty) targets $T_i \subset \Omega$ for each $i \in \Lambda$.

We write $u_i(t) = x_i(t) - x_i(t - 1)$ for the velocity of agent i , implicitly setting the timestep and characteric speed to 1. (We will discuss these implicit parameters later.) For initial conditions $x_i(0)$ and $u_i(0)$ for all $i \in \Lambda$ the system evolves according to the coupled equations

$$x_i(t + 1) = x_i(t) + \lambda_i f_i(X(t), U(t)) \quad \text{for all } i \in \Lambda \quad (9)$$

for functions $f_i : (T\Omega)^n \rightarrow T_{x_i}\Omega$, where X and U are shorthand for $X(t) = (x_i(t) : i \in \Lambda)$ and $U(t) = (u_i(t) : i \in \Lambda)$ and $n = \#\Lambda$ and $\lambda_i \in [0, 1]$ is maximal such that $x_i(t + 1) \in \Omega$. The vector sum above should be understood as a geodesic translation. I.e., $f_i : (T\Omega)^n \rightarrow T_{x_i}\Omega$ takes n position-direction pairs in the tangent bundle and returns a vector $u_i \in T_{x_i}\Omega$. This determines a geodesic with isometric parametrization $\Gamma(s) \subset \Omega$ with $\Gamma(0) = x_i$ and $\Gamma'(0) = \alpha u_i$ for some $\alpha > 0$. Then $x_i(t + 1) = \Gamma(\lambda_i \|u_i\|) = \Gamma(\lambda_i \|f_i(U, X)\|)$. For clarity, we have

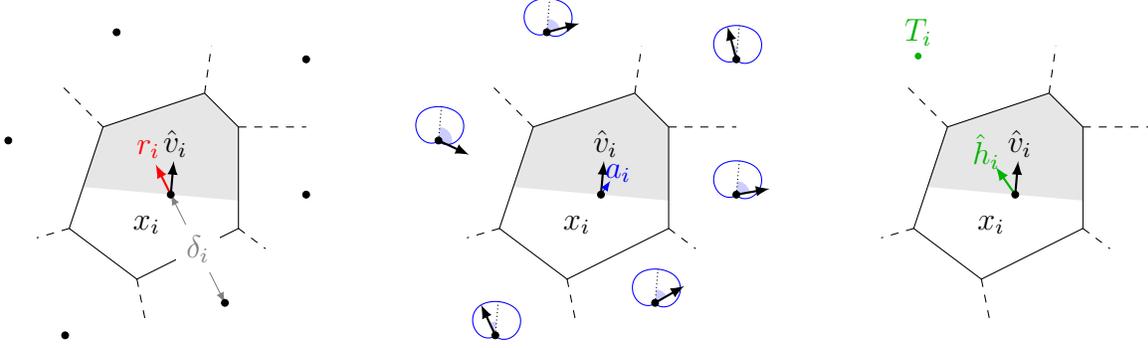
implicitly set two parameters to unity. Namely the time step τ and characteristic speed \bar{u} . Including these explicitly, we have $x_i(t + \tau) = x_i(t) + \tau\bar{u}\lambda_i f_i(X(t), U(t))$. Most of the time, we will ignore this subtlety but we will need to consider their relation when transforming parameters regimes between different domains. We consider here highly symmetric domains where tangent spaces and geodesics are nice. Namely, the flat torus and the 2-sphere. The geometry of the torus is that of a square with periodic (orientable) boundary conditions. The technicalities of working with vectors in this space as if it is \mathbb{R}^2 are discussed in appendix B. On the sphere, the geodesic translations are merely rotations and the displacements are suitably scaled angular displacements on the surface.

At any one time, each function f_i depends only on agent i 's Voronoi neighbors as we shall see. Since their identities will change in general (that is, this is a switched system), we formally write $f_i(X, U)$ as depending on all agents. The functions f_i are given by

$$f_i(X, U) = \rho_i d_i, \quad d_i = \frac{\sigma_i r_i + \nu a_i + \bar{\sigma}_i h_i}{\sigma_i + \nu + b_i \bar{\sigma}_i}. \quad (10)$$

The vector d_i is a weighted sum of *repulsion* r_i , *alignment* a_i , and *homing* h_i , with nonnegative coefficients σ_i , ν , and $\bar{\sigma}_i$. (The factor $b_i \in \{0, 1\}$ is 1 whenever homing is well-defined and 0 otherwise.) The terms r_i , a_i , and h_i are defined in Equations (11,13,14) and the weights σ_i , $\bar{\sigma}_i$ in (15). In general, $r_i, a_i, h_i \in T_{x_i}\Omega$ and are well defined provided there are unique smooth geodesics joining i to each Voronoi neighbor. We scale by ρ_i which depends on i 's personal space (i.e. the geometry of i 's Voronoi cell) and is defined in (16). The three central components of repulsion, alignment, and homing are illustrated schematically in Figure 2.

In general, d_i is not a unit vector so (10) is not a magnitude/direction decomposition. Rather, the external influences of i are captured by d_i while i 's maximum speed given its local environment is captured by ρ_i . Notably, f_i might be small for two very different reasons. If repulsion, alignment, and homing are competing, then d_i will be small. On the other hand, if i is confined to a small region, then ρ_i is small. Nonetheless, $\|d_i\|$ is bounded above by



(a) **Repulsion.** Unit repulsion vector r_i always points away from nearest neighbor or domain boundary. The distance δ_i to this nearest neighbor determines the relative weight of r_i and \hat{h}_i . (b) **Alignment.** Alignment a_i is given by a weighted average of the velocities of Voronoi-neighbors. The circularly-wrapped weighting functions are indicated by the blue curves where the relative angle marked with light blue sectors is the argument. (c) **Homing.** Unit homing vector \hat{h}_i points toward target T_i , if it is nonempty and does not contain x_i . (Here the target is shown as a dot but may be any region, in general.)

Figure 2: **Schematic of the influences on a generic agent.** Here we show one agent i at position x_i as well as its Voronoi cell and Voronoi-neighbors whose positions are marked with black dots. We illustrate the three components which influence i 's motion in the triptych above. Repulsion r_i and homing \hat{h}_i are weighted with coefficients $\sigma(\delta_i/L)$ and convex complement $1 - \sigma(\delta_i/L)$, respectively, where δ_i is the distance to i 's nearest neighbor, as shown in (a) above. The relative weight of alignment a_i is given by the parameter ν . We also show shaded in light gray, the region within i 's Voronoi cell in front of i whose area (appropriately nondimensionalized) contributes to i 's next step size.

$1 + 1/\nu$ on average so we are justified in thinking of ρ_i as limiting i 's maximum speed (at least on the macro scale of the system as a whole).

Throughout, we denote by \mathcal{V} the Voronoi diagram in the Ω generated by $X = (x_i : i \in \Lambda)$ the agents' positions. We write $i \sim_{\mathcal{V}} j$ for $i, j \in \Lambda$ to mean that x_i and x_j are Voronoi neighbors in \mathcal{V} in the sense described in Section 6.2, that is, that the the Voronoi cells V_i and V_j in \mathcal{V} generated by x_i and x_j , respectively, have nonempty intersection. Speaking in terms of the indices i, j, \dots of Voronoi neighbors rather than their positions (which generate \mathcal{V}) is often convenient.

We now describe in detail the components of the model.

6.3.1 Repulsion/attraction

Agents attempt to avoid collisions with nearby agents in the following way. Motivated by the *move-away-from-closest-neighbor* dynamics (MACN) used in the context of Voronoi energy minimization in [21], we define the repulsion term r_i , illustrated in Figure 2a. for positions $X = (x_1, \dots, x_n) \in \Omega^n$, the *repulsion vector* r_i is the average of unit vectors

$$r_i(X) = \frac{1}{\#K_i} \sum_{y \in K_i} \frac{x_i - y}{\|x_i - y\|} \quad (11)$$

where K_i is the set of “obstacles” that minimize proximity to x_i . By “obstacles” we mean all other agents’ positions and possibly the domain boundary. Precisely, $K_i = K_i(X) = \{y \in \Omega : \|x_i - y\| = \inf_w \|x_i - w\|\}$ where the infimum is taken over $w \in \{x_j : j \neq i\} \cup \partial\Omega$, the other points in the tuple X and boundary points of Ω . Equivalently, we need only consider $w \in \{x_j : j \sim_{\mathcal{V}} i\} \cup (\partial\Omega \cap V_i)$ among (the positions of) i ’s Voronoi neighbors and boundary points within i ’s Voronoi cell. So in fact, r_i does not depend on distant effects (in the sense of the Voronoi diagram).

Typically, K_i is a singleton containing the unique nearest obstacle to x_i . Then r_i is just the unit vector pointing x_i away from the obstacle. Naturally, we have defined r_i above to handle in the obvious way cases when x_i is equidistant from several obstacles.

By assigning negative weight to r_i , this term may function as well as an attractive effect and we will return to this possibility below in our discussion of the weighting coefficients.

We also define

$$\delta_i := \text{dist}(x_i, K_i) = \|x_i - y\|$$

for all $y \in K_i$ to be the unique distance from x_i to its nearest obstacle(s), as indicated in Figure 2a.

As the connectivity of \mathcal{V}^* changes, we see in many parameter ranges oscillatory structure in r_i over short time scales. (See [21] for more detailed discussion.) Consequently, in these parameter regimes, agents move frenetically on short time scales but this frenetic motion

averages out of larger spatiotemporal scales and collective motion emerges despite these small effects. While these frenetic motions are not random, the resulting dynamics when in parameter regimes where they are pronounced are qualitatively similar to those produced by models with a stochastic noise effect in agents' velocities or headings.

6.3.2 Alignment

Alignment is illustrated schematically in Figure 2b and is defined as follows. For positions $X = (x_1, \dots, x_n) \in \Omega^n$ and headings $U = (u_1, \dots, u_n) \in \mathbb{R}^{2n}$, we define the *alignment vector* \tilde{a}_i by the *weighted* average

$$\tilde{a}_i = \tilde{a}_i(X, U) = \frac{1}{\#\{j \sim_{\mathcal{V}} i\}} \sum_{j \sim_{\mathcal{V}} i} g(\theta_{ij}) \hat{u}_j \quad (12)$$

where $\hat{u}_j = u_j / \|u_j\|$ and $\theta_{ij} = \arccos(\hat{u}_i \cdot \hat{u}_j)$ is the angle between u_i and u_j . Here $g : [0, \pi] \rightarrow [0, 1]$ is a continuous non-increasing function with $g(0) = 1$ and $g(\pi) = 0$. Thus, agent i considers the *direction of motion* of each of its neighbors and averages these, favoring those whose motion is consistent with its own (θ_{ij} near 0) and virtually ignoring those whose motion is opposed (θ_{ij} close to π). By weighting in this way, we allow the model to tolerate very high shear in the flow despite the averaging of velocities. Because opposing agents have little aligning effect on each other, we observe subsystem behaviors that would not otherwise occur, like co-rotating vortices (we call this an anti-cog), and other phenomena like lane formation which do not arise easily under direction-neutral alignment.

The *alignment coefficient* φ_i is defined by $\varphi_i(X) = \frac{1}{6} \#\{j \sim_{\mathcal{V}} i\}$. To motivate this definition, we note that in any Voronoi diagram (in the torus, sphere, plane, or planar region), a typical cell has at most six neighboring cells (see the Lemma A.0.2 in the appendix). So φ_i captures the ‘‘surroundedness’’ of agent i in the \mathcal{V}^* network. The alignment term is then given by

$$a_i = \varphi_i \tilde{a}_i = \frac{1}{6} \sum_{j \sim_{\mathcal{V}} i} g(\theta_{ij}) \hat{u}_j. \quad (13)$$

Due to φ_i , those agents with relatively few neighbors are subject to a weaker alignment effect. The need for such a scaling is especially evident if one considers the dynamics of systems where most agents have very few neighbors, e.g. only a handful of agents in total or agents confined to a very narrow passage. In such situations, with out φ_i , the agents become tethered to each other, crippling their navigation and dynamics. Introducing φ_i allows these separated agents to navigate more independently than those in the middle of a herd, say. The improvement represented by φ_i is similar in spirit to that the *relative distance* notion brought by [32] to improve upon [11].

One may object that the connectivity of \mathcal{V}^* and therefore the communication topology may have changed in the last time step. If i has only just gained a neighbor j , how can i be expected to have velocity information u_j of this neighbor? The key is that i only needs to know j 's *direction* of motion, not j 's speed. Therefore, we take for granted that the species of agent under consideration can infer orientation of their neighbors from their body geometry (in a computationally trivial amount of time). Because we model agents as point-particles with no oriented geometry, we assume, moreover, that normalized velocity is a sound proxy for orientation, i.e., agents face forward when moving.

6.3.3 Homing

The homing term is shown for a point-target in Figure 2c but the target may be any convex subset of Ω . The homing vector points from i 's position $x_i \in \Omega$ toward its target T_i . Say the target point $x_i^* \in T_i$ satisfies $\|x_i^* - x_i\| = \text{dist}(x_i, T_i)$. There is in general an issue of uniqueness here but in practice, this ambiguity is inconsequential. The *unit homing vector* \hat{h}_i is given by

$$\hat{h}_i(X) = \frac{x_i^* - x_i}{\|x_i^* - x_i\|} \quad \text{for } x_i \notin T_i \quad (14)$$

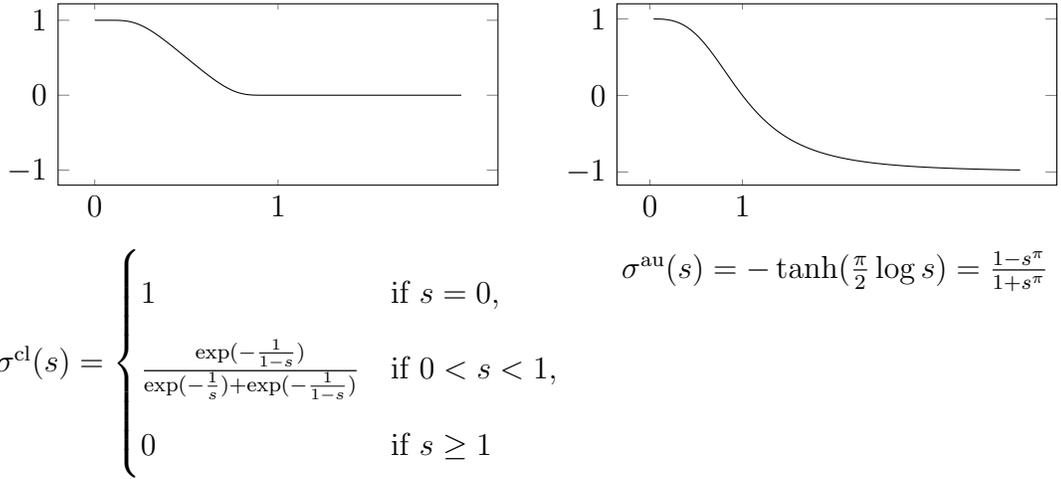
Of course, i may not have a target ($T_i = \emptyset$) or i may have already reached its target ($x_i \in T_i$). So we define the dominating factor b_i to be 0 in these cases and 1 otherwise and write $h_i = b_i \hat{h}_i$ so that $h_i = 0$ when \hat{h}_i is ill-defined.

6.3.4 Weighting coefficients

The weighting coefficients in (10) are defined as follows. The repulsion r_i and homing h_i appear with weights σ_i and $\bar{\sigma}_i$ defined respectively by

$$\sigma_i = \sigma(\delta_i/L) \quad \text{and} \quad \bar{\sigma}_i = 1 - |\sigma_i| = 1 - |\sigma(\delta_i/L)| \quad (15)$$

where $\sigma : [0, \infty) \rightarrow [-1, 1]$ is continuous at 0, non-increasing, and satisfies $\sigma(0) = 1$ and $\sigma(1) = 0$. Here $L > 0$ is a parameter of our model we refer to as the *repulsive falloff distance*. If σ is nonnegative, we see that σ_i and $\bar{\sigma}_i$ are convexly related. Otherwise, they are convexly related in absolute value and the sign of σ_i determines whether r_i acts repulsively or attractively. Canonical choices in each case are *claustrophobic* $\sigma = \sigma^{\text{cl}}$ and *autophobic* $\sigma = \sigma^{\text{au}}$ profiles, plotted below.



Recall δ_i gives the distance between i 's position and the nearest obstacle. So the combination $\sigma_i r_i + \bar{\sigma}_i h_i$ facilitates the following behavior. In the claustrophobic case (i.e. $\sigma(s) = 0$ for $s \geq 1$), if x_i is at least a distance L from all obstacles, then full priority is given to target-seeking via h_i . On the other hand, as obstacles encroach on x_i at distances less than L , collision avoidance via r_i takes priority over target seeking. Additionally, in the autophobic case (i.e. σ negative eventually), if x_i is further than $L \inf\{s : \sigma(s) < 0\}$ from all neighbors, than i will try to close the gap to its nearest neighbor. Strictly speaking, when allowing

$\sigma < 0$, we should also distinguish between cases when repulsion/attraction is due to another agent or to the domain boundary (i.e., environmental obstacles). It usually will not make sense for walls to attract agents so σ_i should generally be a function not only of the distance ratio δ_i/L but also of the *type* of i 's nearest obstacle, for instance taking $\sigma = \sigma^{\text{au}}$ if the witnessing obstacle is another agent but taking $\sigma = \sigma^{\text{cl}}$ if it is the domain boundary.

The remaining coefficient $\nu \geq 0$ in (10) determines the overall strength of the alignment effect compared to the combined homing-repulsion effect.

6.3.5 Personal-space speed scale

Until now, we have relied only on the dual \mathcal{V}^* of the Voronoi diagram. That is, only on the topological structure of \mathcal{V} . The final piece is to consider the geometry of the Voronoi cells—the agents' personal areas—to scale their speed. The scalar ρ_i in (10) rescales the step size accordingly as i 's Voronoi cell allows more space. Precisely, for $x, u \in \mathbb{R}^2$ we define

$$H(x, u) = \{x + w \in \mathbb{R}^2 : u \cdot w \geq 0\}$$

to be the half plane whose boundary contains x and inward normal is (parallel to) u . Then for locally flat spaces Ω , for $X = (x_1, \dots, x_n) \in \Omega^n$ and $U = (u_1, \dots, u_n) \in \mathbb{R}^{2n}$ and Voronoi diagram $\mathcal{V} = \{V_1, \dots, V_n\}$ generated by $\{x_1, \dots, x_n\}$, define

$$A_i^{\text{forward}} = A_i^{\text{forward}}(X, U) = \begin{cases} \text{Area}(V_i \cap H(x_i, u_i)) & \text{if } u_i \neq 0, \\ \frac{1}{2} \text{Area}(V_i) & \text{if } u_i = 0. \end{cases}$$

The $u_i = 0$ case is understood probabilistically as the expected value of $V \cap H(x_i, u_i)$ for arbitrary x_i and measurable set V over u_i from a radially symmetric distribution. The proof is given in the appendix (Lemma A.0.1).

Interpreting $\text{Area}(V_i)$ as agent i 's personal space, then A_i^{forward} gives a natural measure of how much of the space available to i in its facing direction. Note that in two dimensions (or

in three under the obvious generalization), A_i^{forward} is of full dimension in the space. There is an alternative, strictly one-dimensional measure we will discuss shortly.

We nondimensionalize A_i^{forward} via the length scale L we have already introduced, the repulsive falloff distance, taking $\frac{A_i^{\text{forward}}}{\omega_2(L)/2}$, where $\omega_2(L)$ is the area of the disc of radius L .

On the sphere $\Omega = RS^2$, we define A_i^{forward} analogously. For $(x, u) \in T\Omega$ in the tangent bundle, we take $H(x, u)$ to be the *hemisphere* whose bounding great circle γ includes x and whose tangent at x is orthogonal to u . Then $V_i \cap H(x_i, u_i)$ describes the desired region provided V_i itself is contained in a hemisphere, which is guaranteed for at least two generators. We now rescale A_i^{forward} by $\omega_2(L)/2$ where now $\omega_2(L)$ is the area of a spherical cap of geodesic radius $L \in [0, \pi R]$.

Finally, to ensure a physically reasonable, bounded characteristic speed, we enclose $\frac{A_i^{\text{forward}}}{\omega_2(L)/2}$ in an increasing function η that behaves like the identity near zero and goes to unity asymptotically so that agents attain maximum speed when unobstructed. The coefficient ρ_i is given by

$$\rho_i = \rho_i(X, U) = \eta \left(\frac{A_i^{\text{forward}}}{\pi L^2/2} \right). \quad (16)$$

We make the canonical choice $\eta = \tanh$.

As an aside, for general manifolds, we need to be more careful to define the “intersection” $V_i \cap H(x_i, u_i)$. Even on the flat torus, one can concoct Voronoi digrams where it is not clear what the region $V_i \cap H(x_i, u_i)$ should be. The problem, of course, is that V_i lives in Ω while $H(x_i, u_i)$ naturally lives in the tangent plane $T_{x_i}\Omega$. (Notice, we have been careful to circumvent this issue on the sphere.) To capture the same idea in full generality, consider the following. Fix $(x, u) \in T\Omega$ and a closed subset $V \subseteq \Omega$ containing x . Let $D = \{w \in T_x\Omega : \|w\| = 1, w \cdot u \geq 0\}$ be the unit semicircle (or d -hemisphere) in $T_x\Omega$ which is symmetric about the ray extending x through u . For each $w \in D$, let $\Gamma_w : I_w \rightarrow \Omega$ continuously parametrize a geodesic where I is the largest closed interval whose minimum is 0 and satisfies

- $\Gamma_w(0) = x$ and $\Gamma_w'(0) = w$,

- For all $s \in I$, we have $d(\Gamma_w(s), x) = s$ and $\Gamma(s) \in V$

Lastly, for all $s \geq 0$, let $\gamma_s = \bigcup_{\{w \in D: s \in I_w\}} \Gamma_w(s)$. To understand this, imagine sending a unit speed (with respect to the metric d) pulse in Ω from x in all the directions D . Say the boundary ∂V of the region absorbs the pulse. Then γ_s is the pulse front at time s . Now we are equipped to properly define $H(x, u)$ in general. We have

$$\tilde{H}(x, u, V) = \bigcup_{s \geq 0} \gamma_s.$$

In other word, $\tilde{H}(x, u, V)$ is the subset of V seen from x in all the directions D .

There is an alternative (inequivalent) definition of ρ_i which will be preferable in some contexts. It uses the one-dimensional analog of the quantity A_i^{forward} defined above. Let

$$R(x, u) = \{x + lu \in \mathbb{R}^2 : l \geq 0\}$$

be the ray with vertex x_i in the direction u_i . Then define

$$s_i^{\text{forward}} = \bar{s}_i = \begin{cases} \text{length}(V_i \cap R(x_i, u_i)) & \text{if } u_i \neq 0 \\ \bar{s}_i & \text{if } u_i = 0, \end{cases}$$

where \bar{s}_i is chosen suitably. (The analogous probabilistic argument does not yield freindly results³.) For example, one reasonable choice³ of \bar{s}_i is the average of Voronoi neighbors' speeds along direction d_i . Now we may define

$$\rho_i = \rho_i(X, U) = \eta\left(\frac{s_i^{\text{forward}}}{L}\right), \quad (17)$$

where here L is naturally understood as the non-dimensionalizing constant as the 1-measure of the one-dimensional half-ball of radius L .

³Employing the same argument, we find $\frac{1}{2\pi} \int_0^{2\pi} \text{length}(V \cap R(0, \phi)) d\phi = \frac{1}{\pi} \text{Area}(\sqrt{V})$ where \sqrt{V} is obtained from V as follows. If ∂V is given by the polar parametrization $r = w(\theta)$ for function $w : [0, 2\pi) \rightarrow [0, \infty)$, then \sqrt{V} is the region enclosed by $\{\sqrt{w(\theta)} : 0 \leq \theta < 2\pi\}$.

6.4 Effective parameters

parameter	takes values in	dimensions	interpretation
Ω	convex subsets of \mathbb{R}^2	–	environment
τ	positive reals	time	time step
\bar{u}	positive reals	velocity	characteristic speed
n	positive integers	1	number of agents
g	monotonic $C([0, \pi]; [0, 1])$ functions with b.v. $g(0) = 1$ and $g(\pi) = 0$	1	alignment transition function
σ	monotonic $C([0, 1]; [0, 1])$ (or $C([0, 1]; [-1, 1])$) functions with b.v. $\sigma(0) = 1$ and $\sigma(1) = 0$ (or $\sigma(\infty) = -1$)	1	repulsion transition function
L	positive reals	length	repulsive length scale
ν	positive reals	1	alignment strength
η	monotonic $C^1([0, \infty); [0, 1])$ functions with b.v. $\eta(0) = 0$, $\eta(\infty) = 1$, $\eta'(0) = 1$, $\eta'(\infty) = 0$	1	speed scaler
$T_i (\forall i \leq n)$	convex subsets of Ω	–	target regions

Table 1: Degrees of freedom for VTP with fixed number of agents. Note that τ and \bar{u} are suppressed in most equations. They need only be considered when establishing equivalences of the effective parameters between manifolds or if attempting to model a real-world species from experimental observation.

All degrees of freedom are summarized in Table 1. The number of effective parameters however is much smaller. Firstly, we think of the domain Ω and the targets T_i as fixtures of the setup rather than control parameters. Secondly, all dynamics discussed hereafter are robust under small changes to the transition functions σ , g , and η , for which we make canonical choices, so we need not consider these as parameters. Lastly, the time step τ , the characteristic speed \bar{u} , and the domain scale are not independent degrees of freedom and we can assume $\tau = \bar{u} = 1$ (as we have been), rescaling Ω if necessary.

So including the freedom to rescale Ω , we are left with four parameters: the alignment coefficient ν , the repulsive falloff distance L , the number of agents n , and the domain size $|\Omega|$. All the other “weights” are directly determined by the local Voronoi geometry.

However, in the model’s simplest applications these can be further reduced to only two

effective dimensionless parameters. The first is $\nu > 0$ which determines the strength of alignment compared to the combined repulsion-homing effect. The second effective parameter $\mu > 0$ combines n , L , and $|\Omega|$ and represents the ratio of the repulsive length scale to the average inter-agent distance. It is defined to be

$$\mu = \frac{L}{(|\Omega|/n)^{1/2}}.$$

Thus, on the torus $\mathbb{R}^2/l\mathbb{Z}^2$ of primary domain $[0, l)^2$, we have $\mu = L\sqrt{n/l^2} = \sqrt{n}L/l$ and on the sphere of radius R , we have $\mu = L\sqrt{n/(4\pi R^2)} = \sqrt{\frac{n}{\pi}}\frac{L}{2R}$.

The parameter μ is applicable in the absence of homing (i.e., $T_i = \emptyset$ for all $i \in \Lambda$). For the sake of universality, we will also use μ for certain situations with homing but we should note that specific configurations of targets introduce new characteristic length scales particular to those configurations. (Will we see examples of this.) Nonetheless, where μ does apply, we can study the μ - ν phase plane to understand the dynamical regimes. First, we need to define and motivate several *observables* relevant for these systems

6.5 Note on a related model

The model presented by Grégoire, Chaté, and Tu (G-C-T) in [22] bears some cosmetic similarity to ours deserving comment. The motivation behind C-G-T is to demonstrate a “minimal” extension of the model due to Vicsek et. al. [39] which achieves collective motion in highly unfavorable circumstances, as characterized by a handful of criteria. In particular, they seek to achieve aggregation in of agents. The goal of VTP, on the other hand, is to demonstrate a simple (if not “minimally” so) and highly generalizable model which achieves a wide *variety* of collective behaviors including some proper to explicitly “realistic” models. Moreover, VTP mostly disregards sponaneous aggregation instead opting to individual homing as a means to achieve consensus among sub-crowds. Thus, the two models serve very different purposes. Still, given the cosmetic similarities, it is worth a closer

look at exactly where they differ.

For one, the alignment interaction, which is taken to be fundamental by VTP, C-G-T, and of course Vicsek upon which C-G-T is based, occurs strictly among Voronoi neighbors in VTP. Meanwhile, C-G-T imposes simultaneously Voronoi and metrical restrictions on this interaction. Imposing the Voronoi structure on top of the radial zone structure in this way seems somewhat ad hoc but, on the other hand, ensures consistency between the alignment neighborhood and the body force neighborhood. VTP resolves this consistency issue by restricting the body-force-type interaction to only the nearest neighbor. This allows us to consider alignment over all Voronoi neighbors rather than only those also within a fixed radius. Whereas the Voronoi structure is fundamental to VTP’s dynamics, it is merely an additional constraint in C-G-T over Vicsek to ensure that the body-force term only acts between physically adjacent agents.

Secondly, G-C-T is stochastic whereas VTP is completely deterministic. Stochasticity is necessary in G-C-T since one of their unfavorability criteria is “strongly noisy environment and/or communication”. The determinism of VTP is a modelling choice for which we can offer a philosophical justification. Under VTP, agents must measure several quantities of the local environment and their neighbors. Once these quantities are known, we take the error in the subsequent calculation to be insignificant. It certainly makes sense to adopt the noisiness criterion of G-C-T; however, in the VTP context, this noisiness should affect not the final update, but rather the measured quantities that determine the update (e.g. neighbors’ orientations). But deterministic VTP is robust to small changes in these quantities as they are arguments of continuous functions. Thus, the qualitative features of VTP are unaffected if the agents have imperfect knowledge of the world.

As an exception to the above claim. VTP’s *move-away-from-closest-neighbor* (MACN) repulsion component is highly discontinuous in an agent’s neighbors’ positions. However, this component’s frenetic nature is averaged out over large temporal and spatial scales. So since the system’s qualitative behavior depends only on MACN’s *typical* behavior over many

agent-neighbor interactions, introducing noise to the neighbor positions in each instance of its calculation, provided this noise is suitably symmetric, will not influence the qualitative behavior. (An asymmetric noise scheme might reflect higher aptitude for estimating distances ahead than to the side, for example.) The MACN repulsion itself represents another significant distinction from G-C-T whose repulsion averages over all neighbors. VTP prioritizes computational economy when it can be afforded. The fact that MACN works is, in that sense, enough justification. See [21] for more on MACN in the context centroidal Voronoi tessellation energy landscapes.

Also apparently similar in form are the alignment terms in the two models. However, VTP’s alignment interaction is more sophisticated, not merely averaging neighbor’s direction but also weighting them in such a way that neighbors moving against an agent i will be ignored by i . Consequently, VTP produces stable behaviors where parts of the flow maintain high shear such as adjacent, co-rotating vortices and emergent lanes amidst opposing streams. Although the examples just given incorporate homing, let this not eclipse the role of the weighting in the alignment component; these behaviors do not emerge without the more sophisticated *weighted* alignment interaction.

Finally, G-C-T is constant speed while this is variable in VTP. Moreover, speed in VTP depends on the geometry of the Voronoi diagram. This emphasizes the primary tenet of VTP: the Voronoi diagram provides a natural means to synthesize several sophisticated features of agent-based crowd models that are otherwise ad hoc.

6.6 Observables

Before we can discuss the model, we must introduce the observable quantities we use to distinguish distinct qualitative regimes. Throughout, we use three global observables to qualify and distinguish different behaviors. The simplest observable, polarization, concerns only the velocities of all agents. It is given by $\mathcal{P}(U) = \frac{\|\sum_i u_i\|}{\sum_i \|u_i\|}$ for $U = (u_1, \dots, u_n) \in \mathbb{R}^{2n}$. Obviously, this functional is bounded above by 1 and is continuous in U . As the term

suggests, maximal polarization of 1 occurs only when all velocities are in the same direction.

The second observable, angular momentum (with respect to the origin), takes into account velocities and positions and is given by $\mathcal{L}(X, U) = \frac{\|\sum_i x_i \times u_i\|}{\sum_i \|x_i\| \|u_i\|}$ for $X = (x_1, \dots, x_n) \in \Omega^n$ and $U = (u_1, \dots, u_n) \in \mathbb{R}^{2n}$. Recalling that the 2D cross product is given by $a \times b = \|a\| \|b\| \sin \theta$ where θ is the (signed) angle between a and b , we see that \mathcal{L} is maximized when the angle between x_i and u_i is $\pm\pi/2$ for all i , that is, circular motion (in one direction) about the origin. Closely related is the absolute angular momentum \mathcal{L}_{abs} given by $\mathcal{L}_{\text{abs}}(X, U) = \frac{\sum_i \|x_i \times u_i\|}{\sum_i \|x_i\| \|u_i\|}$. Angular momentum about an arbitrary center is obtained by shifting the X coordinates. One can easily convince oneself that, for a given configuration (with at least two agents), \mathcal{P} and \mathcal{L} cannot both be large. As such, only one of these observables will be relevant in each setup but it will always be clear which we are considering.

The third observable is defined in terms of only agents' positions. It may be less familiar than polarization or angular momentum from a particle systems viewpoint, however, it arises frequently in Voronoi literature—especially in the context of *centroidal* Voronoi tessellations, see [14, 15]. Given $X = (x_1, \dots, x_n) \in \Omega^n$, consider the functional

$$E(X) = \int_{\Omega} \text{dist}(x, X)^2 dx = \sum_j \int_{V_j} \|x - x_j\|^2 dx. \quad (18)$$

where V_j is the Voronoi cell containing x_j in the Voronoi diagram generated by X . The quantity $E(X)$ is called the quantization energy of the configuration in variational contexts. Hopefully, the first form above makes it clear that this is a natural to consider if we want a global quantity which captures the variance of the spatial distribution of Ω with respect to the distinguished set X . The second form, on the other hand, should make it clear why Voronoi diagrams are intimately linked to this functional. Gershgorin's conjecture [17], proven in two dimensions, states that for large n , the minimum value of E approaches the value obtained by organizing the generators nearly on a triangular lattice (wherein V_j is close to a

regular hexagon centered at x_j for all j). Precisely,

$$\inf_{|X| \leq n} E(X) \sim n \int_{H_n} x^2 dx =: E_n \quad \text{as } n \rightarrow \infty$$

where H_n is a regular hexagon of area $|\Omega|/n$ centred at the origin. We refer the reader to [38, 34, 23] for the subtleties of Gershgorin's conjecture. For our purposes, it shall suffice to note that this gives us a means to scale E in such a way that its value can be compared for different values of n . Our third observable \mathcal{E} is then

$$\mathcal{E}(X) = \frac{E(X)}{E_n}$$

where $E_n = n \int_{H_n} x^2 dx = \frac{5|\Omega|^2}{n \cdot 18\sqrt{3}}$ here⁴. (This value can be verified by direct calculation). In this context, this quantity should be understood as follows. Large values of $E(X)/E_n$ correspond to *clustering* in the agents' positions X and values close to 1 indicate a lack of clustering. In some figures, we plot $\log \mathcal{E}$ rather than \mathcal{E} to save space and superimpose it with other data on one scale. Each figure specifies which version is plotted. This quantity is highly correlated with the mean *shape index*, given for each cell by the square root of the perimeter-to-area ratio. The shape index of Voronoi cells is of particular interest in epithelia and biological tissue literature.

6.7 Results

6.7.1 Homing-free systems in domains without boundary

We call a system *homing-free* under VTP if targets $T_i = \emptyset$ are empty for all i . In this circumstance, the homing coefficient $b_i = 0$ vanishes for all i and (10) reduces to

$$f_i(X, U) = \rho_i \frac{\sigma_i r_i + \nu a_i}{\sigma_i + \nu} = \rho_i \frac{\sigma(\delta_i/L) r_i + \nu a_i}{\sigma(\delta_i/L) + \nu}. \quad (19)$$

⁴On the sphere, E_n is computed differently but has the same scaling properties.

and repulsion and alignment alone determine the dynamics.

First we consider the square torus $\Omega = \mathbb{R}^2/l\mathbb{Z}^2$ with fundamental domain $[0, l)$. As

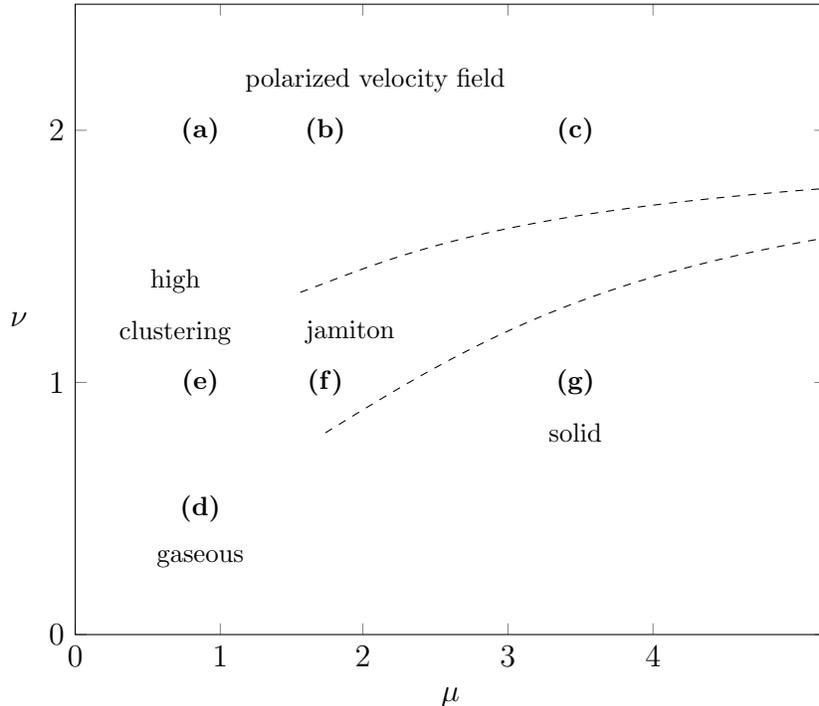


Figure 3: Phase portrait sketch where $\Omega = [0, l)^2$ (with periodic boundaries). The dimensionless parameter μ is given by $\mu = L(n/|\Omega|)^{1/2} = \sqrt{n}L/l$. The dashed lines are merely conceptual delineations, not sharp bifurcation loci. Their rough shape is based on coarse probing of phase space with simulations. The indicated features of the phase landscape are merely suggestive; the precise features of the phase diagram of course also depend on the particular choices of transition functions and a finer parameter sweep is appropriate only following tuning of those transition functions for specific application. At this level of generality, we only highlight the coarse structure of phase space.

mentioned above, the parameter μ is vital for homing-free systems. Figure 3 shows a sketch of the μ - ν phase plane with key values indicated. The labels correspond to the regimes shown in Figure 4. Here, we wish to conceptually distinguish the dynamics at the top and bottom of the diagram from those between the dotted curves. We think of this intermediate zone as a transitional region between the upper and lower parts of the diagram.

At the top of the diagram, where ν is large, we observe high polarization with agents tending to align globally. Regimes at the bottom of the diagram by contrast do not spontaneously achieve any velocity consensus. The dashed curves are meant to suggest the

transitional region where the consensus behavior is harder to characterize.

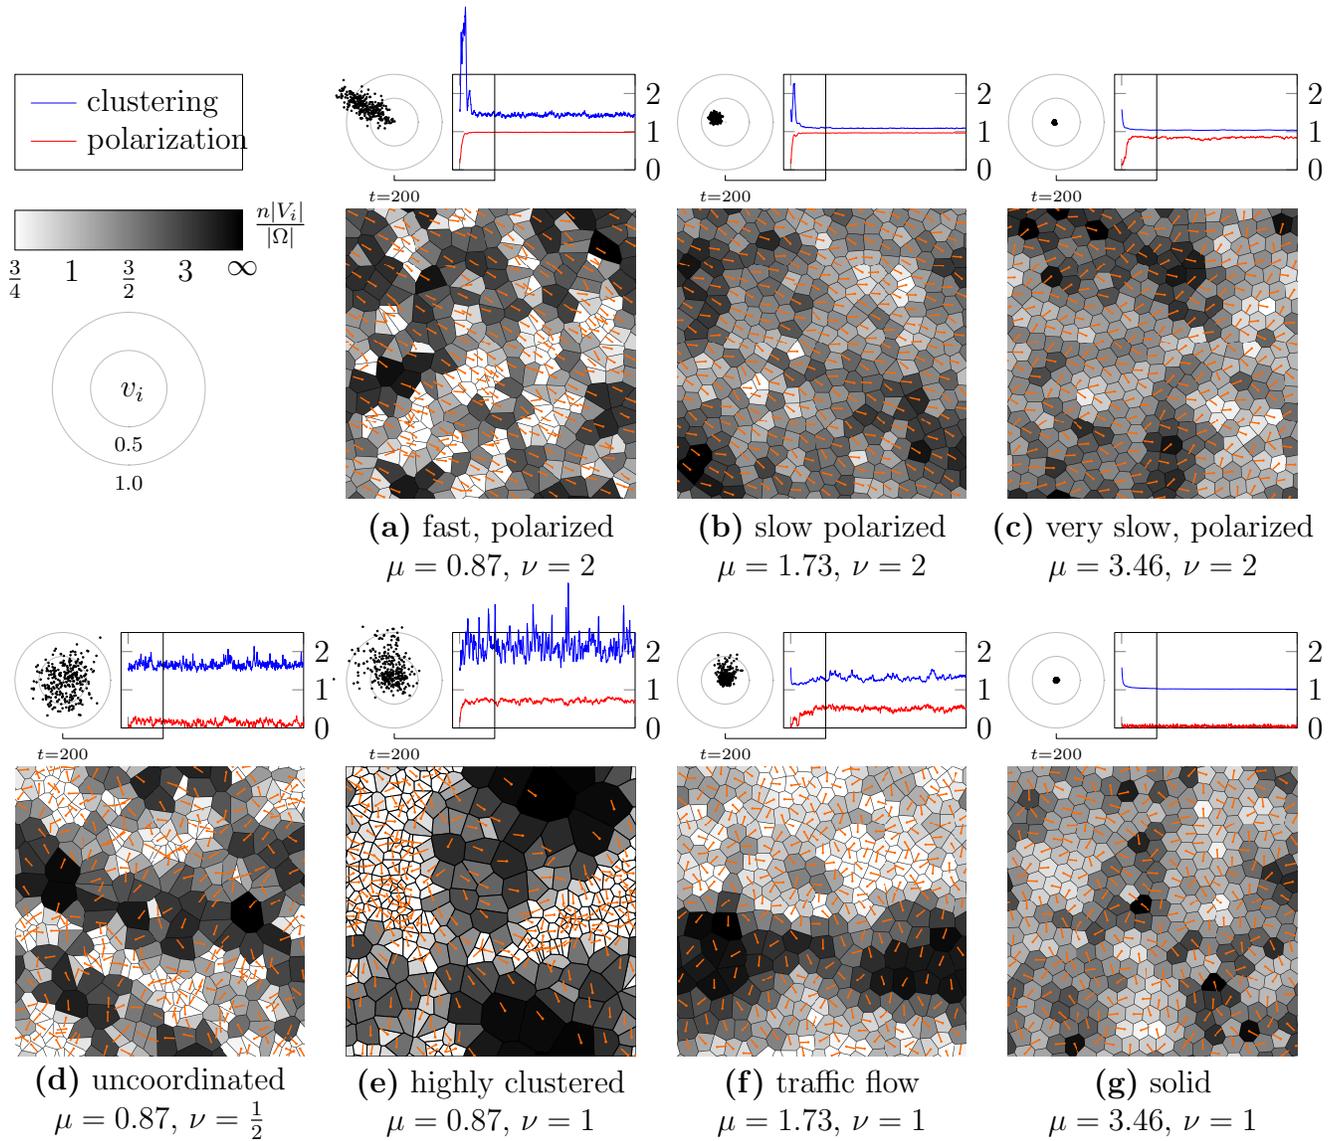


Figure 4: Different qualitative long-term behaviors of homing-free systems in the torus. In the digital version of this document, click on an image to view the corresponding animation online. Each image above is a still (at time step 200) representative of the *long term behavior* for the indicated parameter values with $n = 300$ agents. The qualitative features of each image are persistent into the future for that neighborhood of parameter values. The obscure values of μ are due to the factor of $n^{1/2} = \sqrt{300}$ but L/l is a simple ratio in all cases. Shown also are time series plots over the first 1000 iterations for the polarization \mathcal{P} in red and quantization error (or clustering measure) \mathcal{E} in blue and scatter plots of all $n = 300$ velocities superimposed at time $t = 200$. Each radial increment is a speed increment of $\frac{1}{2}$.

We argue that all qualitative behaviors are captured within this bounded region of the phase plane. As ν becomes very large and we move far upward, we see from (19) that $f_i \approx \rho_i a_i$ and the dynamics cease to depend on ν . Moving the far right of the phase space, we see the following. Keeping n fixed, we have the asymptotic behavior $\rho_i = O(|\Omega|/L^2) = O(\mu^{-2})$ and so agents move ever more slowly as L grows. The long-term consensus behavior is nonetheless unchanged. Let $s > 0$ such that no two neighboring agents are separated by more than a fixed multiple $s|\Omega|^{1/2}$ at any point in the evolution. (The very crude choice $s = \text{diam}(\Omega) = \sup_{x,y \in \Omega} \|x - y\|$ suffices for this argument.) Then, $\sigma(s/L) \leq \sigma(\delta_i/L) \leq 1$. By the continuity of σ at 0, if $L \gg s$, then $\sigma(\delta_i/L) \approx 1$ for all i and $f_i \approx \rho_i \frac{r_i + \nu a_i}{1 + \nu}$ depends on L only through ρ_i .

For regimes near the left of the diagram where μ is small, we tend to find high clustering, corresponding to large values of \mathcal{E} , and the states have Voronoi diagrams with more irregular cells of varying area. As μ increases, repulsion strengthens and agents spread out and achieving more regular Voronoi diagrams with less eccentric cells of nearly uniform area, corresponding to low values of \mathcal{E} . Near the bottom of the diagram, no spontaneous velocity consensus emerges and directions vary frenetically, reflecting low values of \mathcal{P} . Increasing ν , the velocity field attains a very high polarization \mathcal{P} as the system spontaneously gains a stable nonzero drift velocity. Across this highly polarized region of the phase space, there is a tradeoff between extremely high polarization $\mathcal{P} \approx 1$ at the left and high regularity of the Voronoi diagram at the right, $\mathcal{E} \approx 1$. Hence, the (scaled) quantization error \mathcal{E} and polarization \mathcal{P} , naturally mirror the μ - ν phase plane. Figure 4 shows representative simulations for each of the marked points **(a)**–**(f)** on the phase plane sketch for $n = 300$.

Now let's characterize the distinct regimes observed here. The regimes of (spontaneous) high polarization are represented in 4**(a)**, 4**(b)**, and 4**(c)**. Here, ν is large and the polarization \mathcal{P} quickly rises to a stable value, as can be seen in the plots above each subfigure. In the regimes of 4**(c)** and 4**(g)**, where L/l (and hence μ) is large, agents tend to crystalize, locally organizing into a triangular lattice. The “solid” regime is shown in Figure 4**(g)**. Here,

the system is pushed toward constant density by the large repulsion and the weak alignment cannot dominate the frenetic nearest-neighbor repulsion. So the system remains stationary macroscopically with only small jittering at the scale of individual particles. Even in this, strong repulsion range, increasing ν slightly is enough to cause a spontaneous overall drift, as shown in Figure 4(c).

Where both parameters are small, as in 4(d), no spontaneous consensus is achieved and the long-term behavior depends highly on initial conditions. These small-parameter states are characterized by very weak alignment and a sufficiently sparse agents that, typically, an agent is within L of at most one of its neighbors. Because both alignment and repulsion are too weak to alter the qualitative dynamics significantly, the long-term behavior tends to roughly preserve polarization and energy of the initial conditions. The other regimes are independent of initial conditions.

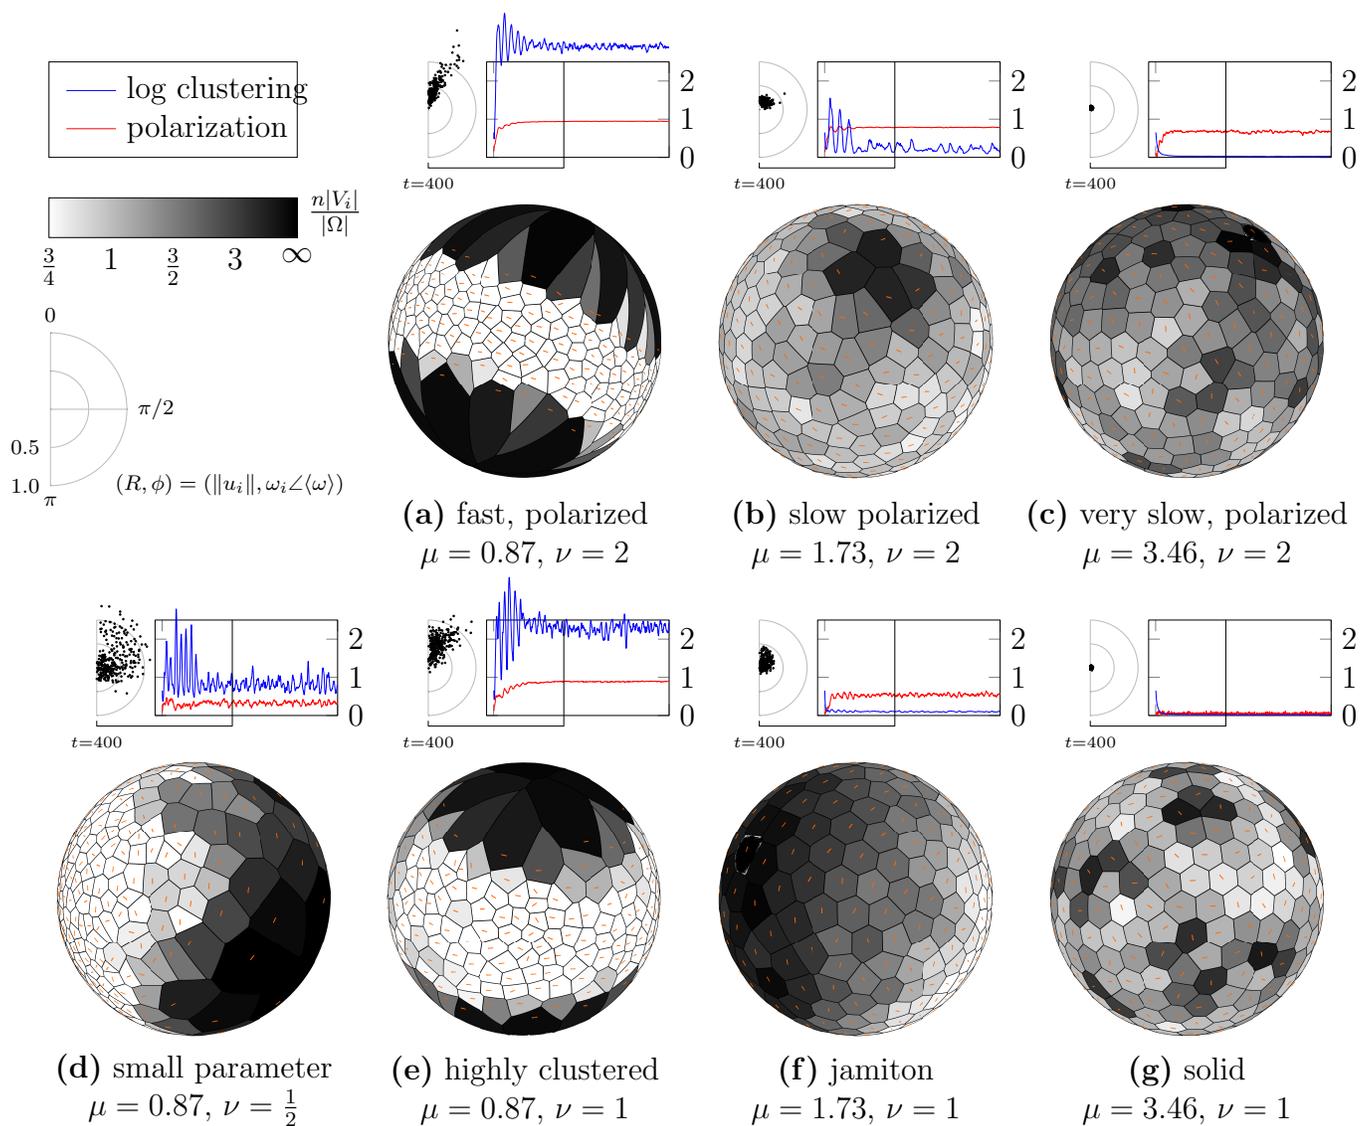


Figure 5: **Different qualitative long-term behaviours in the homing-free sphere.** In the digital version of this document, click on an image to view the corresponding animation online. Each image above is a still (at time step 400) representing the *long-term behavior* for the indicated parameter values with $n = 300$ agents. The qualitative features of each image are persistent into the future and robust under small changes to μ and ν . Shown also are time series plots over the first 1000 iterations for the polarization \mathcal{P} in red and the logarithm clustering measure $\log \mathcal{E}$ in blue. The scatter plots represent the magnitude and angular deviation from the mean of all $n = 300$ angular velocities at time $t = 400$. Each radial increment is a speed increment of $\frac{1}{2}$.

In Figures 4(e) and 4(f), we see examples of regimes in the transitional region of the phase plane. In both, we see that clusters of agents spontaneously align for a time and then fall out of alignment. We also see fluctuations in the density which do not disappear in the future. In 4(e), this transient density fluctuation is especially pronounced, giving rise to the high and \mathcal{E} values with large deviation from the mean. In this regime, alignment is simply too weak to ever eradicate the tight clusters where repulsion is strongest and even out the density. Increasing μ within the transition region, the stronger repulsion does begin to even out the density but the spatial variation in local polarization persists. In Figure 4(f), we see a sensitive regime where \mathcal{E} is low and \mathcal{P} stabilizes on an intermediate value. Here, the alignment information propagates through the system but the alignment interaction is not strong enough for its influence to spread through the entire system.

We do not observe sustained regimes wherein agents form (multiple) groups which move coherently but mostly independently of each other, i.e., emergence of subflocks. Introducing an attractive effect through σ as described in the previous section allows for behaviors like this but because the attraction only acts between nearest neighbors, the subflocks will not have any mutual attraction to each other once sufficiently separated. Emergent subflocks driven by this mechanism will continue to form, dissipate, and interact in bounded domains but will do so less moving toward the zero-density limit.

Under a suitable transformation we implement the same parameter values on the sphere $\Omega = S^2$ as on the torus. We find extremely good agreement of qualitative behaviors on the two manifolds in all ((a)–(c), (f), (g)) but the smallest parameter cases ((d), (e)) where spontaneous emergence effects are very weak. Figure 5 shows each regime, again for $n = 300$ agents, with the corresponding parameter values from Figure 4. Here, μ is defined as $\mu = L\sqrt{n/|\Omega|} = \frac{\sqrt{n}L}{\sqrt{\pi}2R}$. As reasoned above, very large values of the parameters do not produce new behaviors (the argument applies verbatim on any bounded domain). In Figure 5, we depict iteration $t = 400$ rather than $t = 200$ as in Figure 4 only because the initial transient tends to be slightly longer on the sphere (as can be seen in the plots). Also we

plot here $\log \mathcal{E}$ rather than \mathcal{E} because its values are typically much higher on the sphere (the agreement with the torus here is in the shape of the the curve).

Since the velocities on the sphere do not lie in one plane, we cannot superimpose them as we did in Figure 4. Instead, we calculate the mean angular velocity $\langle \omega \rangle$ at each time step, and plot the deviation from the mean of i 's angular velocity ω_i . Contrast Figure 5(a) and Figure 5(f). In the former, we find $\omega_i \angle \langle \omega \rangle$ near 0 for all i , corresponding to high \mathcal{P} , and that the angular deviation increases with speed. So the fastest agents deviate the most from the overall tendency. In (f), the opposite is true, the fastest agents determine the trend in the angular momenta while the slower agents have uncorrelated angular momenta.

6.7.2 Pinwheels, rings, and cogs with point-target homing

In the previous section, in the homing-free setting, we saw a number of emergent qualitative behaviors from dynamics governed only by alignment and repulsion. Introducing targets, the model predicts a variety of interesting behaviors. Returning to the torus $\Omega = \mathbb{T}^2$, we consider point-targets, $T_i = T$ for some discrete T , which are the same for all agents. For very large ν or μ , targets have little effect on the dynamics. The argument follows basically as above. If ν is very large, alignment dominates homing and $f_i \approx \rho_i a_i$ does not depend on the targets. For L very large $\bar{\sigma}_i = 1 - \sigma(\delta_i/L) \approx 0$, and repulsion dominates homing.

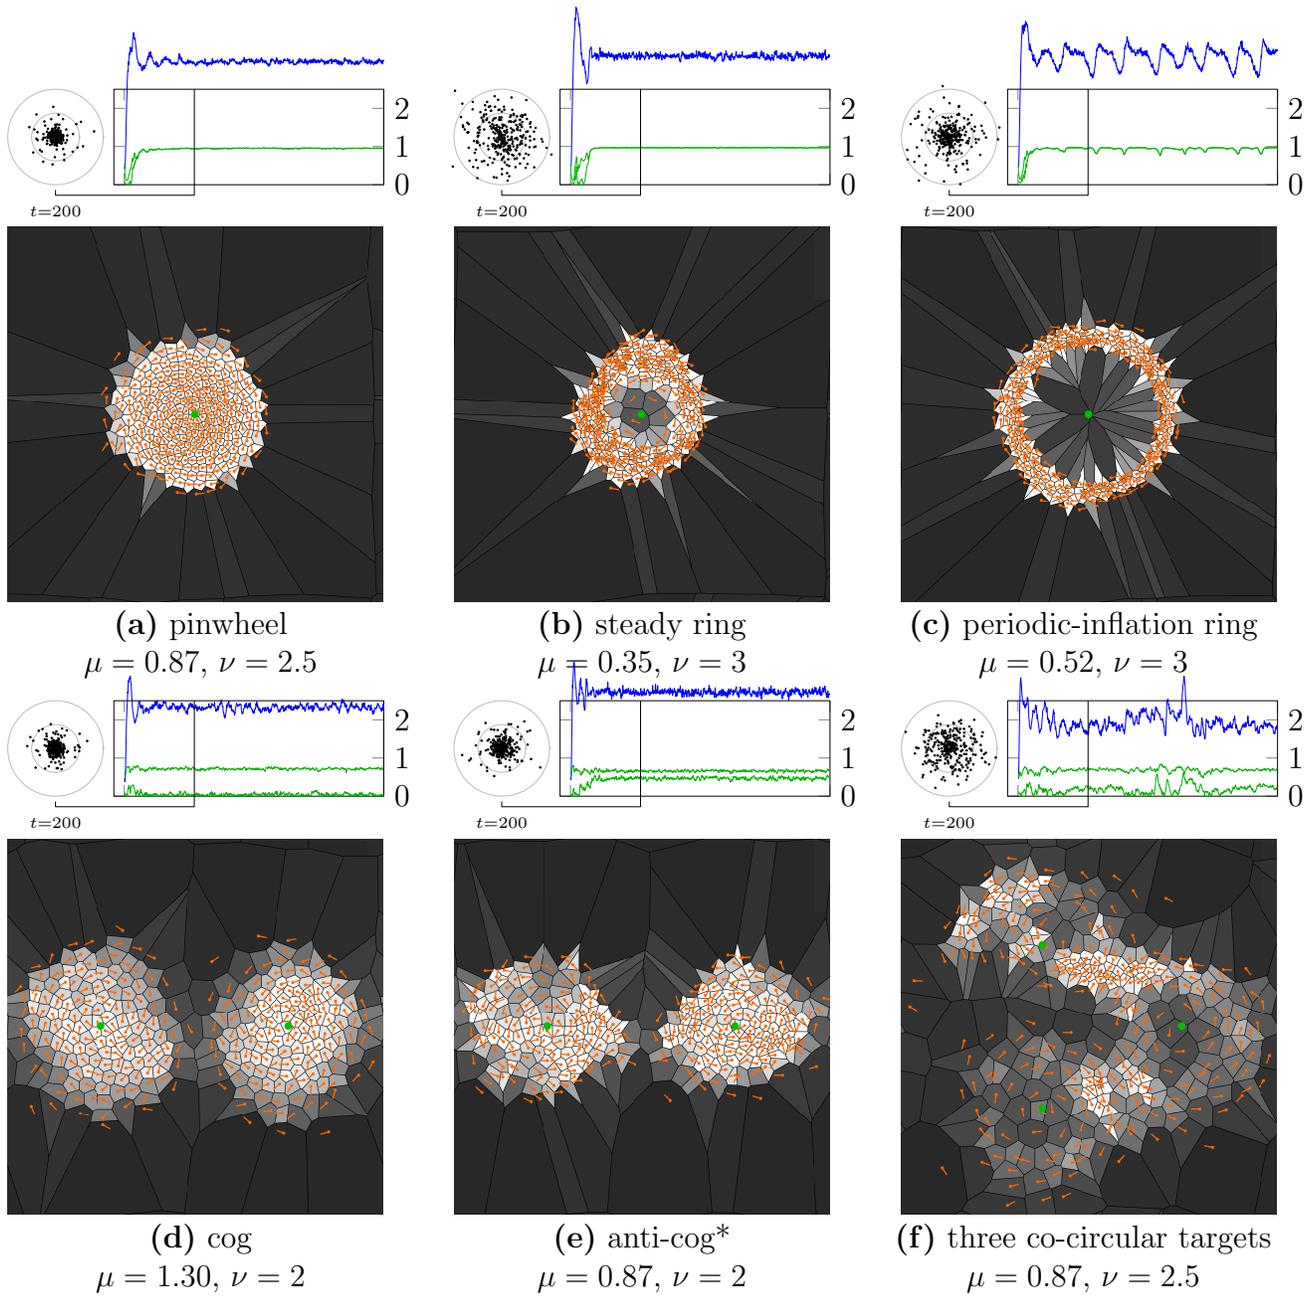


Figure 6: **Example behaviors in the torus with one, two, and three point-targets.** In the digital version of this document, click on an image to view the corresponding animation online. We show here configurations with one, two, and three point-targets, marked by green dots, at the indicated parameter values after an initial transient. Above each image is a plot over the first 1000 iterations of the log-energy $\log \mathcal{E}$ in blue and both the angular momentum \mathcal{L} and absolute angular momentum \mathcal{L}_{abs} of the system in green with respect to the point at the image center.

Figure 6 shows representative simulations for one-, two-, and three-point targets. As for the homing-free systems, we catalogue the regimes based on μ and ν . This is mainly for reproducibility. But we emphasize again, μ is not really the right parameter here and indeed, we expect there is no natural one-size-fits-all parameter for different configurations of targets because targets introduce new characteristic length scales.

We find three distinct milling regimes with a single point target. They are shown in Figures 6(a), 6(b), and 6(c). Of course, all points are created equal in the torus. The target point p is shown at the center of each image. With $T_i = \{p\}$ for all i , every agent's homing vector h_i points from its position x_i (along the shortest geodesic) to p always. The milling regimes emerge for larger values of ν than considered in the last section. At $\nu = 2.5$, we obtain a “pinwheel” (Figure 6(a)), a dense disc of agents orbiting a point near the target point p with roughly constant angular velocity akin to a rigid-body rotation. In the right range of μ , as ν increases to 3, the center of the pinwheel becomes unstable and the system settles into a rotating ring (Figure 6(b)). Then increasing μ and the repulsion scale, the ring radius is no longer stable, instead it oscillates. The ring repeatedly expands in radius slowly until reaching a critical diameter whence it becomes unstable and spirals in to again. This periodicity is mirrored in \mathcal{E} and is apparent in the plot.

Panels (d) and (e) have a two-point target. As one might expect from the single-point target cases, two vortices form after a brief transient. However, we observe both counter- and co-rotating (resp., (d) and (e)) vortices. This manifests strongly in the angular momentum measurements. In both cases, we see high *absolute* angular momentum but the counter-rotating case has near zero angular momentum, i.e., perfect cancellation. The anti-cog is stable for μ sufficiently small but is not as robust as the cog. In particular, for certain distributions of random initial conditions, either state may be reached but there exist more restrictive initial conditions for which the anti-cog appears consistently.

In Figure 6(f), we have three-point target $T = \{p_1, p_2, p_3\}$ where p_1, p_2 , and p_3 are equispaced on a circle of radius one quarter the length of a minimal geodesic so that the minimal

geodesic through p_1 and p_2 is a minimal geodesic of the torus. As expected, we see three vortices under suitable parameter values. Here, a cluster forms at the interface of the two co-rotating vortices where inflow of agents exceeds outflow. After enough agents accumulate here, the whole configuration becomes unstable and we have a transient regime before three new vortices emerge, restarting the process. One such recurring transient happens in the pictured simulation near $t = 600$, corresponding to the spikes in the angular momentum and quantization error.

We examine next the two-point target $T = \{p, q\}$ where p and q are equi-spaced along a minimal geodesic in the torus. This choice is arbitrary but exhibits the interesting regimes we wish to discuss, shown in Figures 6(d) and 6(e). Recall that each agent seeks the nearer of p and q at each time. For the right parameter range, we see a vortex form spontaneously about each target point. What is more, within this parameter range, both counter-rotating (Figure 6(d), “cog”) and co-rotating (Figure 6(e), “anti-cog”) mills are dynamically stable. Readers are encouraged to view the animations (linked from the images) to appreciate the difference but it is also clear from the angular momenta, plotted in green above each subfigure. (Arbitrarily, the angular momenta are computed with respect to the point at the image center.) For both the cog and anti-cog, \mathcal{L}_{abs} is high, but \mathcal{L} is near zero for the cog, indicating near cancellation. Near the indicated parameter values, the cog is extremely robust. The anti-cog, on the other hand, is very stable but its spontaneous emergence is sensitive to initial conditions. In the extreme case, at indicated parameter values allowing for the anti-cog, the cog is also stable and either may emerge from uniformly random initial positions and direction. However, we can design initial conditions which do settle into the anti-cog consistently. On the torus, if the targets lie on an equator, we can start with all agents in the “northern” cylinder moving “westward”. This configuration has built-in relatively high counterclockwise circulation about each target but is (perhaps) not too contrived (depending of course on the precise species and environment one wishes to model). Indeed, this settles into the anticog and is robust under the addition of substantial noise to the initial directions.

On the sphere, the analogs of the cog and anti-cog for antipodal targets are the rolling and anti-rolling regimes, respectively. In each case, there is a vortex about each target point but in the former, their angular momenta align while in the latter, they cancel. The anti-rolling behavior never results from random initial condition within the explored parameter ranges, however, the analogous tailored initial conditions (here a single vortex about a point on the equator between the targets with noise) consistently produces the anti-rolling regime.

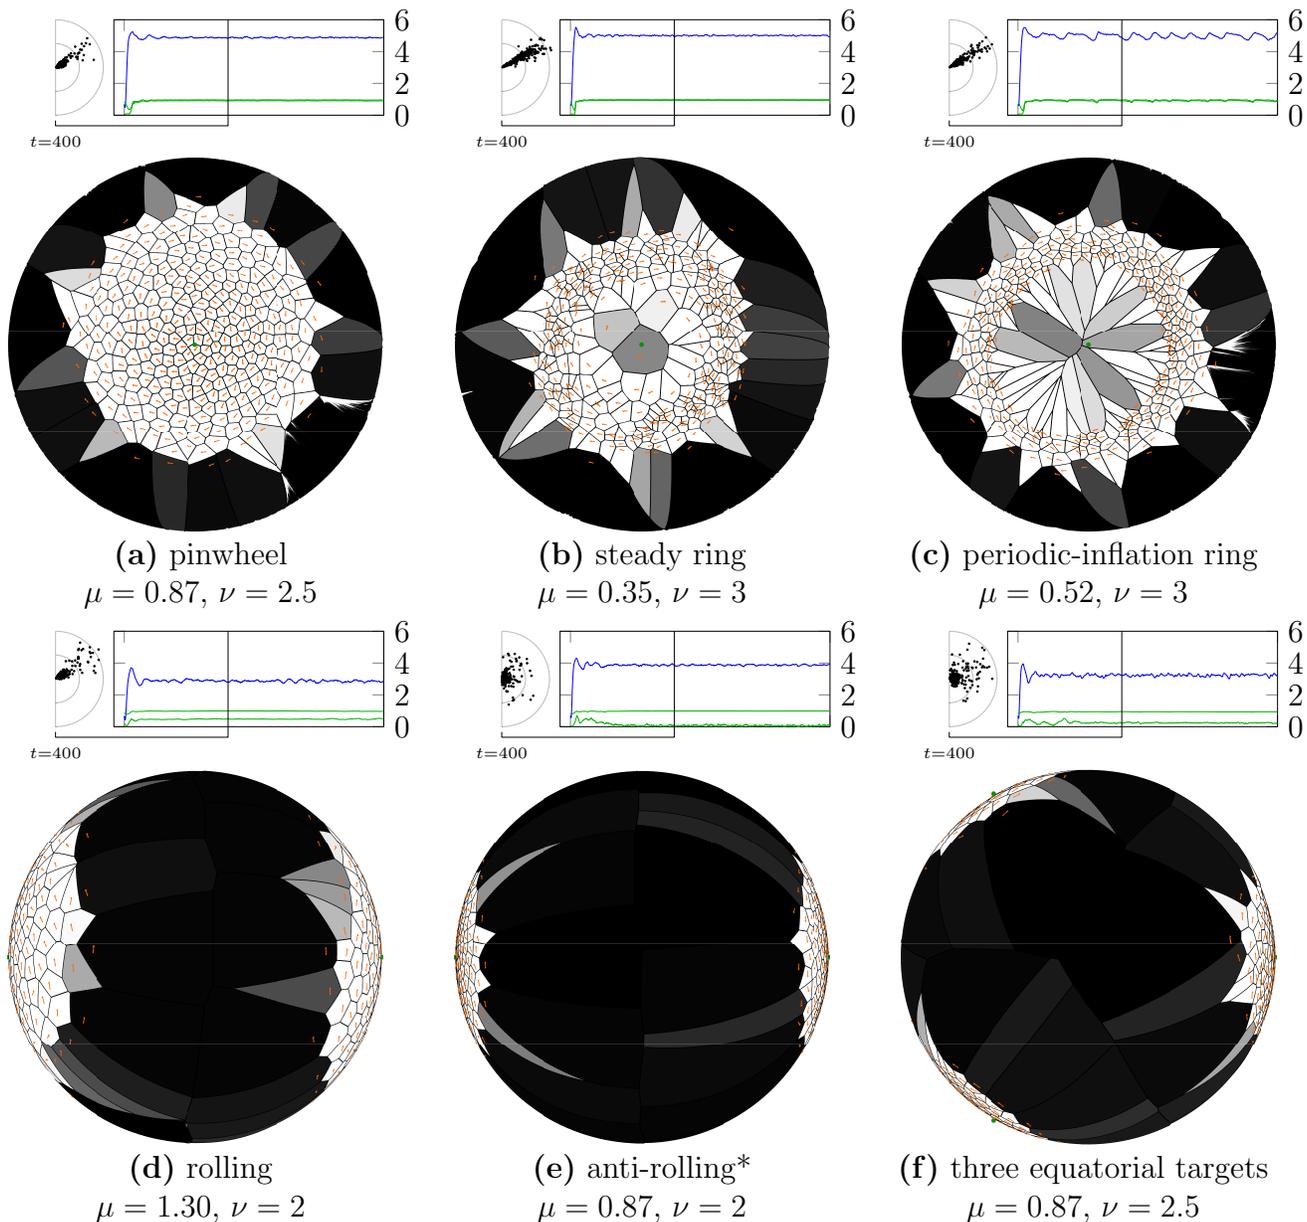


Figure 7: **Example behaviors in the sphere with one, two, and three point-targets.** In the digital version of this document, click on an image to view the corresponding animation online. Here we consider configurations with one, two, and three point-targets, marked by green dots. Above each image is a plot over the first 1000 iterations of the log-energy $\log \mathcal{E}$ in blue and both the angular momentum \mathcal{L} and absolute angular momentum \mathcal{L}_{abs} of the system in green.

6.8 Introducing sources and sinks

6.8.1 Source algorithm

Introducing sources and sinks allows us to study, within the VTP framework, collective behaviors—like lane formation or queuing, efficiency of multidirectional flow through different environments—which are not characterized by dynamical stability of a fixed collection of agents but rather by emergent features which manifest continually in transient collections of agents.

We introduce here a computationally cheap scheme for implementing sources under VTP. This procedure is motivated chiefly by pedestrian crowds but is very simple and relies only on the Voronoi geometry and the motion of agents.

For each entrance region $B \subseteq \partial\Omega$, agents may enter at some point in B according to the following rules. Suppose at time t , the agents in Ω are indexed by Λ_t and have positions $X = \{x_i\}_{i \in \Lambda_t}$ and velocities $\{u_i\}_{i \in \Lambda_t}$. New agents may now enter in B according to a probability distribution. Let $\mathcal{V} = \{V_i\}_{i \in \Lambda_t}$ be the Voronoi diagram generated by X so that V_i is the Voronoi cell with generator x_i . Notice that B is partitioned into Voronoi edges as $B = \bigcup_{i \in \Lambda_t} B \cap V_i$. Assume that B is C^1 so that the outward normal to $\partial\Omega$ is well-defined on B . Then for almost every $y \in B$, there is a unique $i_y \in \Lambda_t$ such that $y \in V_{i_y}$, i.e., $y \in B$ almost surely belongs to exactly one Voronoi cell in V_{i_y} in \mathcal{V} . Denote the outward unit normal to B at y by \hat{n}_y . Let

$$Q(y) = \begin{cases} 1 & \text{if } \Lambda_t = \emptyset, \\ 1 - \sigma\left(\frac{\|y - x_{i_y}\|}{L}\right)g(\arccos(\hat{n}_y \cdot \hat{u}_{i_y})) & \text{if } \|y - x_{i_y}\| \geq L_c, \\ 0 & \text{if } \|y - x_{i_y}\| < L_c, \end{cases} \quad (20)$$

where L_c is a new parameter which in essence limits the maximum flow. (Here σ and g are the same functions as previously introduced.) A brief aside, it may seem unruly to introduce this new parameter L_c . However, in some sense this does represent a simplification in that

this whole scheme subsumes initial conditions, a space with many more degrees of freedom.

Pseudocode for the VTP algorithm including sources is given below. The inner loop will always terminate as long as L_c is positive and each B_i has finite length. This is because at each iteration (of the inner loop), some point x_i is specified in the entrance region B which has the effect of forcing Q to zero on an interval (inside the curve B) centered at x_i of width L_c . Thus, on the next iteration, the new point x_i cannot appear in this interval and after $\lfloor \text{length}(B)/L_c \rfloor + 1$ iterations, Q vanishes everywhere on B with probability 1.

In practice, one need not compute the Voronoi diagram in every iteration of the inner loop. Rather, we need to keep track of only the distances between generators (including those introduced during the loop), significantly improving the runtime of the source algorithm. The Voronoi diagram need only be computed at the end of each loop in preparation for the next step of the dynamics subroutine. Practically, the algorithm with sources is nearly as fast as without, the brunt of the computation being due to Voronoi methods which the sources do not invoke.

6.8.2 The bi-directional corridor

Here we discuss the motivating example for VTP with sources and sinks, the bi-directional corridor. The corridor Ω is a rectangle $[0, l] \times [0, 1]$ of fixed aspect ratio l . The left wall acts as a source for a class of Λ^R of rightward-moving agents. All agents in Λ^R take the entire right wall as a target and this wall acts as a sink for Λ^R , those agents disappearing if they reach the wall. Similarly, there is a class Λ^L of leftward-moving agent for whom the left and right walls have reversed roles to those of Λ^R . The corridor is initialized with no agents and the source distribution density parameter L_c is fixed. Representative stills for three different values of L_c are shown in Figure 8. Beneath a critical value of L_c , we observe steady percolation of the rightward- and leftward-moving agents past each other and we see emergent single-file lanes. Above this critical value, we see opposing fronts form while agents continue to pour in from the sources. By design of the entrance algorithm, this

Algorithm 1: VTP with domain entrances

Data: Domain Ω , finite set of finite-length C^1 entrance regions $B_1, \dots, B_k \subset \partial\Omega$,

Dynamics subroutine

begin

$\Lambda \leftarrow \emptyset$;

$i \leftarrow 0$;

while *true* **do**

foreach entrance region $B \in \{B_1, \dots, B_k\}$ **do**

 define $Q : B \rightarrow [0, 1]$ as given by Equation (20);

$Z \leftarrow \int_B Q(y) dy$;

while $Z > 0$ **do**

 let $\tilde{Q} = \frac{1}{Z}Q$;

$i \leftarrow i + 1$;

$\Lambda \leftarrow \Lambda \cup \{i\}$;

 choose $x_i \in B$ according to the probability density \tilde{Q} ;

$u_i \leftarrow$ inward unit normal to B at x_i ;

$\mathcal{V} \leftarrow$ Voronoi diagram in Ω generated by $\{x_j\}_{j \in \Lambda}$;

 update Q ; // for new Λ and \mathcal{V}

$Z \leftarrow \int_B Q(y) dy$;

end

end

 run *Dynamics* on $(x_j, u_j)_{j \in \Lambda}$ one step; // in general, modifies Λ and \mathcal{V}

end

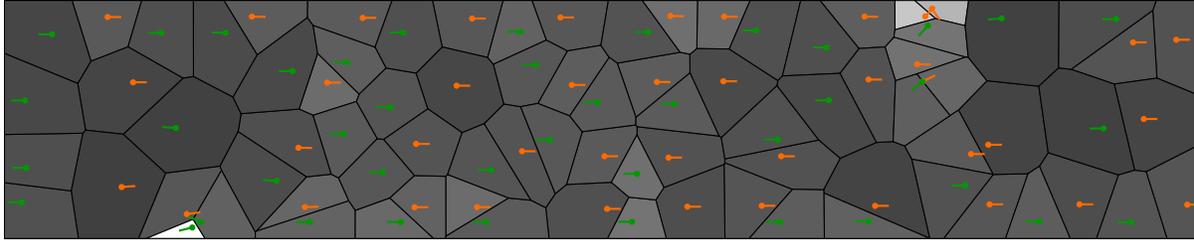
end

cannot continue indefinitely. The hallway become saturated and no more agents enter. It is unclear whether this stalled state lasts indefinitely or if Λ^L and Λ^R eventually escape past each other. Interestingly, intermediate values of ν facilitate easier flow through the hallway for fixed values of L_c . Moderate alignment destabilizes a stalled front perpendicular to the length of the corridor by causing it to precess, allowing Λ^R and Λ^L to slip past each other. For large ν , this rotation effect become very pronounced, even leading groups of agents away from their targets. Detailed study of the bi-directional corridor and related domains is the subject of future work.

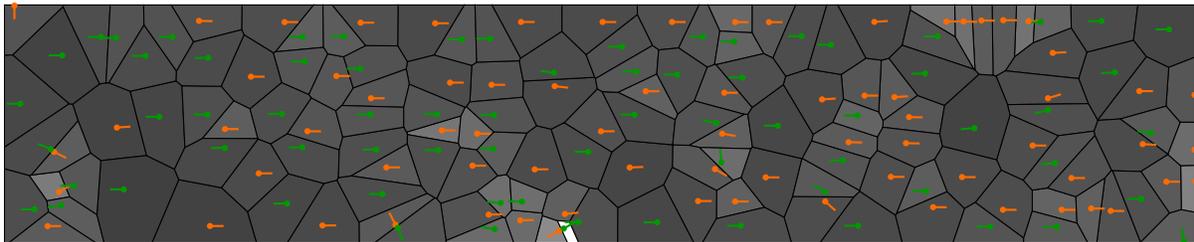
7 The directed Delaunay network

7.1 Motivation: Obstacles and VTP in non-convex domains

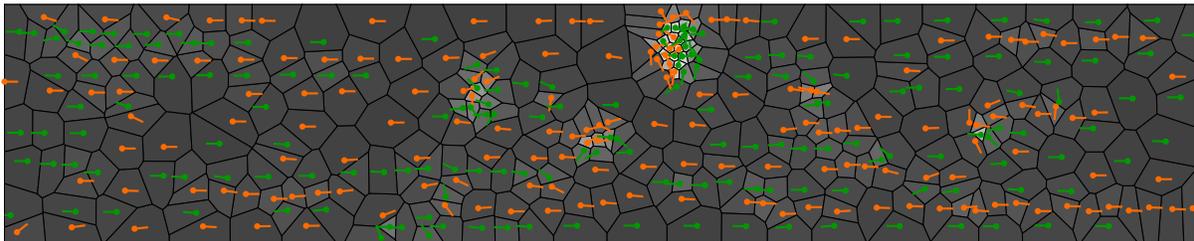
We would like to model the behavior of multi-agent systems in complex environments which may include internal obstacles like walls, pillars and chokepoints or in environments like hallways with corners whose large-scale geometry is essentially non-convex. Throughout this section, we consider only planar domains, that is, pre-compact, connected subspaces of \mathbb{R}^2 . Most of what follows generalizes straight-forwardly to higher dimensional Euclidean spaces, but not as easily to general manifolds and non-Euclidean spaces. In the context of VTP, we take the domain boundary to represent an *impassable* and *opaque* barrier. Moreover, we assume that agents' perception of their environment and other agents is mediated along straight lines. These assumptions are motivated by human crowds whose primary perceptive apparatus is sight. We wish therefore to clearly delineate the domain of applicability of the notions introduced here. For a robotic fleet, for example, a given environment may present impassable but radio-transparent barriers and these constructions will not apply. Contrastingly, humans whose environment is partitioned by various walls and closed doors (opaque but passable) are not well described by these constructions. Nonetheless, a wide variety of human and animal settings fall within this section's relevance.



(a) Light influx: $L/w = 1/12$ and $L_c/w = 1/3$



(b) Moderate influx: $L/w = 1/12$ and $L_c/w = 1/4$



(c) Heavy influx: $L/w = 1/12$ and $L_c/w = 1/6$

Figure 8: Bi-directional corridor. Agents marked in orange emerge from the left edge according to a probability distribution determined at each time step by the local configuration of the system and they target the right edge, exiting if they reach it. Agents marked in green have the reciprocal instruction. Readers are encouraged to view the animated version of these simulations. Among the three pictures scenarios, only L_c varies. Notice how nearby agents with the same objective tends to file to avoid counter-moving agents.

Firstly, let us see why the standard Voronoi notions are inadequate in this setting. Under the assumptions listed above, multi-agents systems in non-convex domains are characterized by *incomplete perceptual information*. In general an agent cannot perceive the entire domain and all other agents within it. The standard Voronoi diagram, even in metrics which somehow respect the non-convexity, generally requires of agents information they do not have access to. Consider the “horseshoe” domain Ω shown in Figure 9. Here we consider only two generators. The horseshoe domain represents the simplest class of (non-convex) domains we would like to consider: simply-connected polygonal regions. What follows applies equally well to more complicated domains but even this simplest class exhibits most of the difficulties we will need to overcome. The horseshoe domain with the inherited Euclidean metric yields a Voronoi diagram which is totally incongruous with the VTP assumptions. As Figure 9a shows, cells need not even be connected, not at all in keeping with the local personal space interpretation VTP assumes valid. To resolve the connectedness issue, one might propose using a metric which respects the domain’s non-convexity like the geodesic, or shortest path, metric. The geodesic Voronoi diagram for the same configuration is sketched in Figure 9b. While this solves the connectedness problem, it yields cells with complicated, non-polygonal boundaries. More importantly though, both metrics badly violate the assumption that every agent can compute its own Voronoi cell because, to drive home the obvious, the two agents in each example cannot perceive each other. This problem cannot be remedied by designing a clever metric. Even by restricting the metric to have finite values only when points are mutually visible (in the sense that the segment connecting them is contained in Ω), as in Figure 9c, one does not circumvent the fact that agents need inaccessible information to construct their own Voronoi regions. Indeed, if Ω is non-convex, then for any non-trivial (\mathbb{R} - or $(\mathbb{R} \cup \{+\infty\})$ -valued) metric d on Ω , if there exist three points $p, q, r \in \Omega$ such that the segment $\overline{pq} \not\subseteq \Omega$ while $\overline{pr} \subseteq \Omega$ and $\overline{qr} \subseteq \Omega$ and $d(p, r)$ and $d(q, r)$ are finite, then we can reproduce this dilemma by simply placing generators at p and q because to determine whether r belongs to each of their Voronoi cells, p and q must each know $d(p, r)$ and $d(q, r)$. To extend

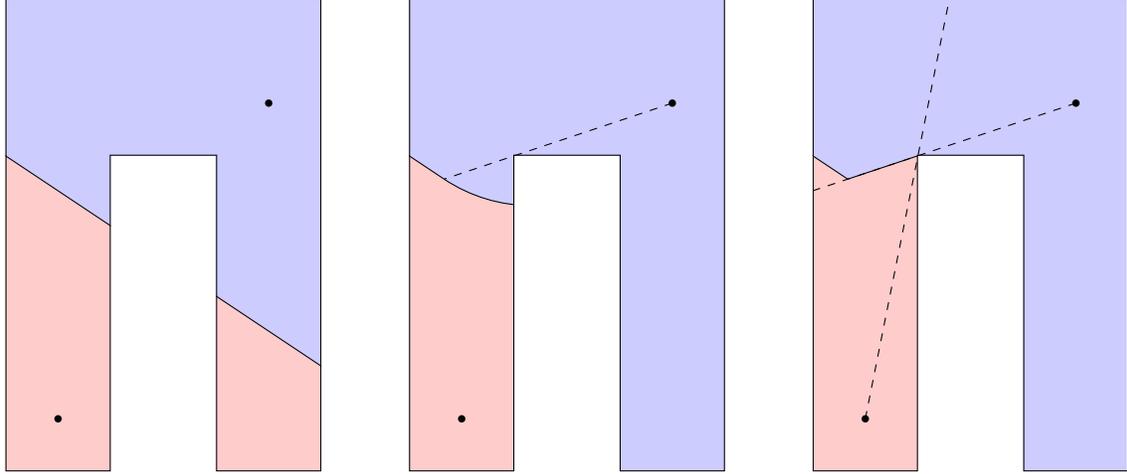
VTP to non-convex domains, we require a new generalization of the Voronoi diagram. We will argue that the right generalization is given by the *directed Delaunay network*.

Before we describe the directed Delaunay network in detail, let us address one other possible approach to the non-convexity problem. One could partition the domain into convex pieces and take the Voronoi diagram to be the union of the Voronoi diagrams in each piece, adopting suitable rules for how agents interact across pieces. This approach has the advantage that its implementation is essentially no more costly than standard Voronoi methods and this certainly may provide an adequate solution sometimes. However, two problems are evident. Firstly, it is not hard to imagine situations for which it is not at all clear how to choose the convex partitioning; consider for example a forest, modelled as a (non-simply-connected) region with many, irregularly placed “holes” at tree trunk sites. Secondly, this approach might be intolerable depending on what *behaviors* one wishes to study. For instance, if one wants to apply the VTP framework to study predator-prey dynamics, then correctly accounting for occlusion and visibility of agents within the domain may be crucial. Finally, even if one adopts a simpler solution than what follows, consideration of the pros and cons of the more sophisticated approach may be informative nonetheless.

7.2 Construction and computational methods

The guiding principle behind the directed Delaunay network is that *each agent should behave as if the visible part of the domain and only those agents it contains are the entire configuration*. In particular, the analog of the Voronoi cells we shall construct will correspond exactly to the usual Voronoi cells each agent would construct if using only the accessible information. We expect to lose some nice properties of standard Voronoi diagrams like symmetry in the notion of *neighboring* cells, partitioning of the domain, and convexity of the cells. However, we will—as a matter of principle—be rid of the ails described above. Chiefly, we will respect the hidden information characteristic of non-convex domains.

Here we outline the construction of the directed Delaunay network. This notion general-



(a) Voronoi diagram with the (inherited) Euclidean metric. (b) Voronoi diagram sketch with the geodesic (shortest path) metric. (c) Voronoi diagram under the visibility geodesic metric.

Figure 9: Horseshoe domain demonstrates failure of standard Voronoi diagram in non-convex domains regardless of metric.

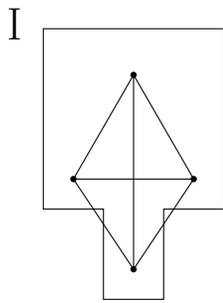
izes the Delaunay triangulation in the sense that for convex domains, they are isomorphic. However, in general, the directed network is very far from being a triangulation. The construction proceeds as follows. Beginning with a domain pre-compact $\Omega \subset \mathbb{R}^2$ and a finite set $P \subset \Omega$,

- I. Let G be the complete undirected graph over P (where here the notion of complete excludes loops).
- II. Consider the map $\eta : G \rightarrow \mathbb{R}^2$ which preserves P and maps each (abstract) edge $\{p, q\}$ in G to the segment $\eta\{p, q\} = \overline{pq} \subset \mathbb{R}^2$. Next, remove from G all edges $\{p, q\}$ for which $\eta\{p, q\}$ intersects the domain boundary $\partial\Omega$.
- III. Now, for every vertex $p \in P$, we obtain an undirected graph G_p with vertices $G_p^0 = \{p\} \cup \{q \in P : \{q, p\} \in G^1\}$, i.e. p and all q adjacent to p in G , and edges $G_p^1 = \{\{p, q\} : q \in G_p^0 \setminus \{p\}\}$ —all edges incident to p in G .
- IV. We compute the (usual) delaunay triangulation T_p over each G_p^0 by a conventional algorithm.

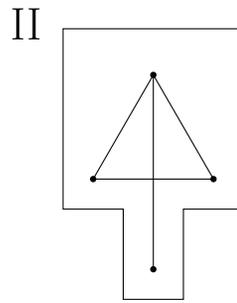
- Va. For each $p \in P$, construct the directed graph G'_p by taking the intersection of G_p and T_p (in which all edges are incident to p) and set every edge *exiting* p and no other (directed) edges.
- Vb. We then obtain a single directed graph via the union $\bigcup_{p \in P} G'_p = (\bigcup_{p \in P} G_p^0, \bigcup_{p \in P} G_p^1) = (P, \bigcup_p G_p^1)$. We denote this directed graph by $\text{DDN}(P, \Omega)$ or simply $\text{DDN}(P)$ if the domain of discourse is clear.

Consider the example of the DDN construction illustrated in Figure 10. This simple example shows the emergence of the asymmetry and the constructive procedure. Here, three of the four intermediate Delaunay triangulations are trivial—each having at most three vertices—but the first is not and even in a situation this simple, we can see its consequence on the final network, giving an asymmetry between the top and bottom generator. Hopefully, the example makes clear the sense in which the directed Delaunay network generalizes the standard Delaunay triangulation. Each generator in the DDN, upon considering its out-neighbors, “thinks” it sees the planar Delaunay triangulation generated by itself and all other generators visible to it. The proof is the construction itself: for each generator p , we effectively disregard those other generators occluded from p by the (non-convex) domain and compute the Delaunay triangulation T_p on those remaining. We then retain only the information of T_p local to (i.e., incident to) p . This information is not generally symmetric so a directed graph is the natural structure.

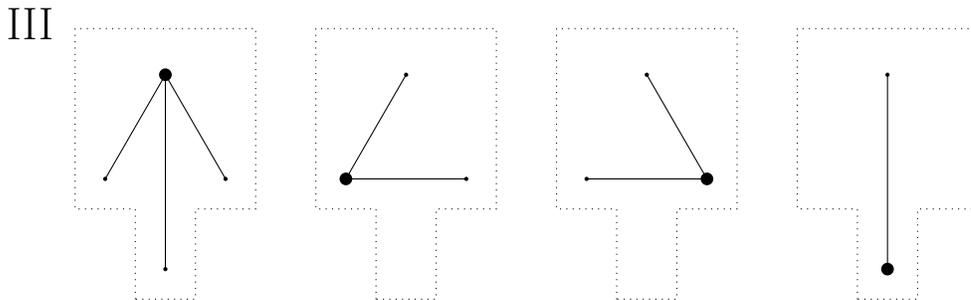
Let’s make a few key observations about the directed Delaunay network and then we will consider optimization of the construction. The first obvious observation is that if the convex hull of P is contained in Ω , then the directed Delaunay network is isomorphic to the standard Delaunay triangulation. Clearly, as suggested by Figure 10, it would be wise to parallelize steps III and IV in the worst case. However, the previous observation and the given example suggest that there is a middle ground, somewhere between the cases when every generator p gives a unique triangulation T_p and the case $\text{conv}(P) \subseteq \Omega$ where they all



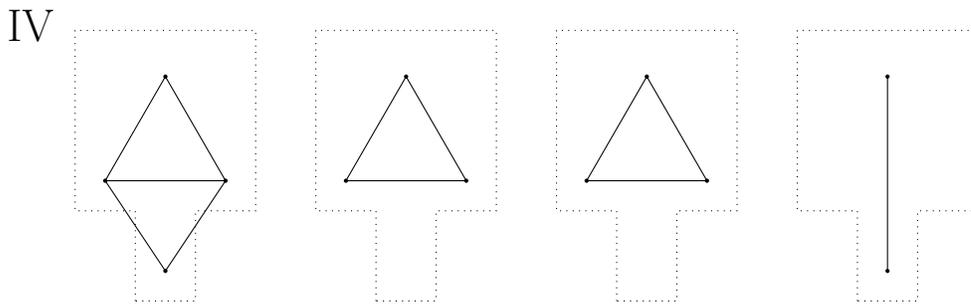
Initialize complete graph over P .



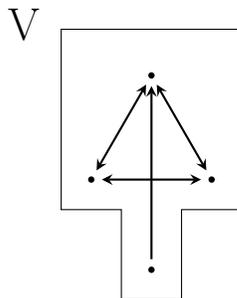
Remove edge $\{p, q\}$ if $\eta\{p, q\} \cap \partial\Omega \neq \emptyset$.



For each p construct the graph G_p from the neighbors of p in II.



Construct the standard Delaunay triangulation T_p over each vertex set G_p^0 ,



For each $p \in P$, take $G_p \cap T_p$ with all (directed) edges exiting p and take the union of digraphs over all p .

Figure 10: Directed Delaunay network construction

give the same triangulation.

Assuming feasibility of step II, we can be more discriminating using only the adjacency matrix of the (undirected) graph at that stage. Precisely, if G_{II} is the graph after step II and A is its adjacency matrix, we can define an equivalence relation \sim on P by $p \sim q$ if and only if the p and q -th rows of $A + I$ are identical where I is the identity matrix. We call an equivalence class under this relation a *coterie*⁵. This works because identical rows of $A + I$ correspond exactly to p and q neighboring each other and having the same neighbors among $P \setminus \{p, q\}$. Put another way, over each class $[p]$, the (sub)graph is complete and if $r \notin [p]$ is connected to some $q \in [p]$ by an edge, then r is connected by an edge to every point in $[p]$. Then, if $p \sim q$, they will generate the same triangulation in step IV (simply because they determine identical subsets of P). So we need only compute one triangulation for each coterie. The necessary modification to step III is to construct a graph $G_{[p]}$ for each coterie rather than for each p and indeed $G_{[p]} = \bigcup_{q \in [p]} G_p$ is just the union. That is $G_{[p]}^0 = \{r \in P : \{q, r\} \in G^1 \text{ for some } q \in [p]\}$ and $G_{[p]}^1 = \{\{q, r\} : q \in [p], r \in G_{[p]}^0 \setminus \{q\}\}$. The modified algorithm is illustrated for the same example in Figure 11.

In all, the algorithm described above is not cheap, requiring $O(n^3)$ or $O(n^2 \log n)$ time depending on the intermediate Delaunay algorithm. If Ω is an m -gon⁶ (or a difference of polygons with m -sides total) and $\#P = n$, then the first step involves determining (the existence of) intersections between a family of $\binom{n}{2}$ segments (the edges of the initial complete graph) and a family of m segments (the sides of Ω). We must then pairwise compare the rows of the $n \times n$ matrix $A + I$ which in the worst case is $O(n^3)$ binary comparisons ($\frac{1}{2}n(n+1)$ pairs of rows, each with n entries). We proceed to compute (in parallel or otherwise) a Delaunay triangulation for every equivalence class. It is in this step where we see the salient feature: the number of coterie is the most significant factor in the complexity of the computation

⁵The word *coterie* is a seldom used synonym of *clique*, which of course has a distinct meaning in the context of graphs. Unlike cliques, coterie are disjoint, very elite!

⁶We can in fact weaken the conditions on Ω further. We need only the $\partial\Omega - \partial\text{conv}(\Omega)$ is a union of polygonal curves.

(provided the network is connected⁷).

We make the following general observations about coteries.

Proposition 7.2.1. Given an undirected graph $G = (G^0, G^1)$ with no loops (i.e., an edge whose vertices are the same), let $C \subseteq G^0$ be a subset of vertices of G . Then the following are equivalent.

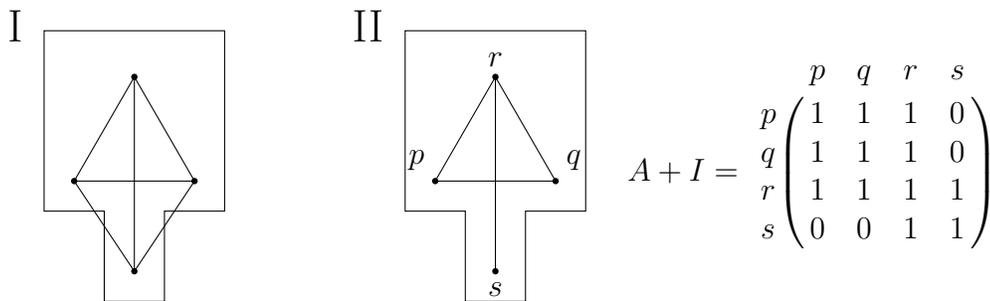
1. C is a coterie.
2. If $q \in G^0$ is connected to C by an edge, then q is connected by an edge to every point of C (except itself if $q \in C$).
3. For all $p, q \in C$, we have $\{p, q\} \in G^1$ and if $\{p, r\} \in G^1$ for some $r \in G^0 - \{p, q\}$, then $\{q, r\} \in G^1$ as well.

Proof. Recall that coteries are equivalence classes of the relation \sim given by $p \sim q$ if and only if the p -th and q -th rows of $A + I$ are identical where A is the adjacency matrix of G . Let's show first that 1 implies 2. Assume C is a coterie and say $a_{pr} = 1$ for some $p \in C$ and $r \in G^0$. Since C is a coterie, it must be that also $a_{qr} = 1$ for all $q \in C$. If $r \notin C$, this shows that r is connected to every point of C . If $r \in C$, then this shows that r is connected to every other point of C by an edge (although not itself because $a_{rr} = 1$ only due to I fixing the diagonal). That 2 implies 3 is obvious. Finally, assume 3. Let $p, q \in C$. We need to verify that $a_{pr} = a_{qr}$ for all $r \in G^0$. But this exactly our assumption if $r \notin \{p, q\}$. Otherwise, since G has no loops, $a_{pp} = a_{qq} = 1$. By assumption $a_{pq} = a_{qp} = 1$. □

Note that the above is equally true if one allows looped edges in G provided we define our equivalence relation over $A \vee I$ where \vee is the bitwise OR operation. (Clearly, $A \vee I = A + I$ is the loop-free case because the diagonal of A is zero.)

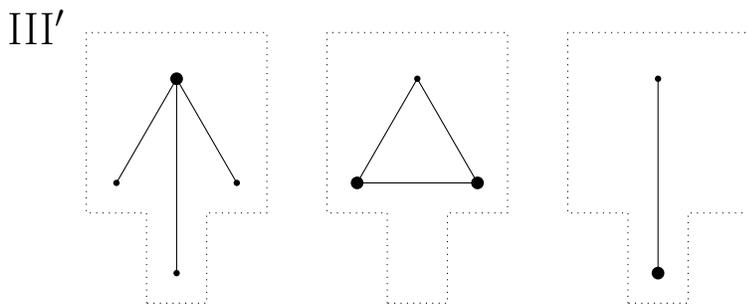
Corrolary. Every coterie is a clique.

⁷it is not hard to construct examples where the number of equivalence classes is maximal but where also every vertex is isolated so the triangulations are trivial.

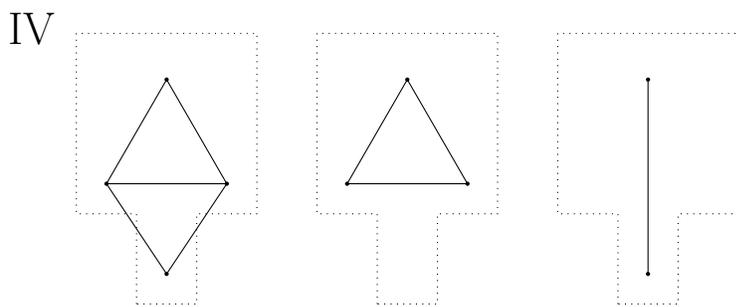


Complete graph over P .

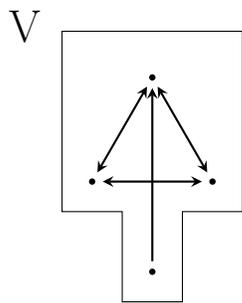
Remove edge $\{p, q\}$ if $\eta\{p, q\} \cap \partial\Omega \neq \emptyset$.



For each class $[p]$ construct the graph $G_{[p]}$ from the neighbors of $[p]$ in II.



Construct the standard Delaunay triangulation $T_{[p]}$ over each vertex set $G_{[p]}^0$,



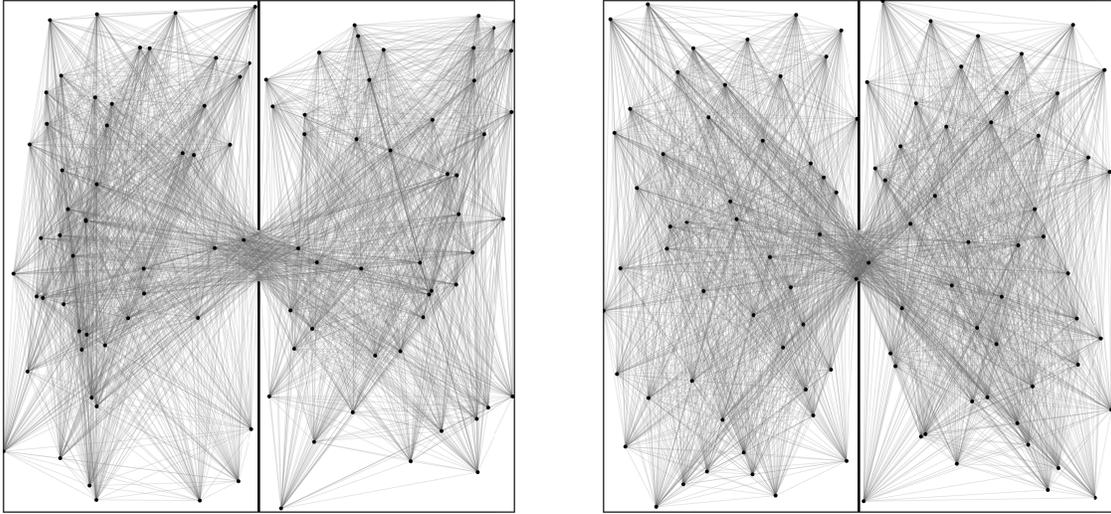
For each class $[p]$, take $G_{[p]} \cap T_{[p]}$ with all (directed) edges exiting p and take the union of digraphs over all classes.

Figure 11: Directed Delaunay network reduced construction

The obvious question to ask is *how significant an improvement does the modified algorithm represent over the original?* Moreover, is the number of equivalence classes a quantity of independent interest? To explore the first question, consider the following *two-room experiment*. Let Ω be the square $(-1, 1)^2$ bisected by a (zero thickness) partition parallel to two sides. Allow a gap at the center of the divider whose width is some fraction δ of the square's side length. That is, $\Omega = (-1, 1)^2 \setminus D_\delta$ where $D_\delta = \{0\} \times ([-1, -\delta] \cup [\delta, 1])$. Thus, Ω consists of two rooms with a doorway between them. Note that $D_\delta \subset \partial\Omega$. Now fix n and consider random sets P of n generators.

Two observations are immediate. If $\delta = 1$, i.e. the doorway fills the entire wall, the situation reduces to the convex case and the number of coteries is exactly one for all P . At the other extreme, if $\delta = 0$, the domain becomes disconnected and each component is convex so the number of coteries is exactly two for all P . For intermediate values of δ , on the other hand, we expect many coteries! It certainly is not clear whether the extremes correspond to *continuous limits* in any natural sense but it seems reasonable that they should. In the intermediate case we have great visibility between the left and right rooms on account of the wide doorway but also we have significant occlusion due to the large walls. As δ increases, occlusion diminishes and as δ decreases, visibility diminishes.

For $n = 100$, probing over P from the uniform distribution and from a quasi-random process (see Figure 12) yield the same results. Histograms over 10,000 configurations are shown for different values of δ in Figure 13 and Figure 14. For uniformly distributed generating sets, even for the least favorable values of δ , the mean number of coteries is about 15% less than than the number of generators n . We believe this result is generic for some appropriately nice domains and generator distributions. Conversely, one might assume the mean number of coteries (as a fraction of n , say) over appropriate distributions of generators as a “niceness” measure of the domain. The coterie statistics for the two-room experiment in particular have a notable feature. Letting δ approach its extreme values, we find not only that the mean decreases but also strong preferences emerge for the parity of the number of



(a) Generating set P from uniform distribution.

(b) Quasi-random generating set P .

Figure 12: Example initializations for the two-room experiment. Here $\delta = 0.1$.

coterics. As $\delta \rightarrow 1$, even numbers of coterics become much rarer than odd and the opposite holds as $\delta \rightarrow 0$. Moreover, in all cases, the even and odd histograms individually look like bell curves with the diminished one of the two having the higher mean. The quasi-random configurations perform slightly worse in the sense that the mean number of coterics is higher in the unfavorable δ range but this is to be expected as “clumping” should promote larger (and hence fewer) coterics. The same qualitative distributions also appear for the analogous experiment in a circular domain. This shows that the apparent difference between the rate of convergence to each of the extreme cases is not due to shape effects of the square. The profile of distributions over the entire range of δ does change substantially with n . In particular, for smaller n , the unfavorable range shrinks.

Although the algorithm for computing the DDN is costly, it is also highly parallelizable. With the fast preprocessing step of computing the coterics, the rest of the construction may be executed in $O(n \log n)$ -time on $O(k)$ processors where k is the number of coterics. The feasibility of this algorithm depends very strongly on the domain under consideration. Consider the *corner experiment* in contrast to the two-room experiment. Here the domain is

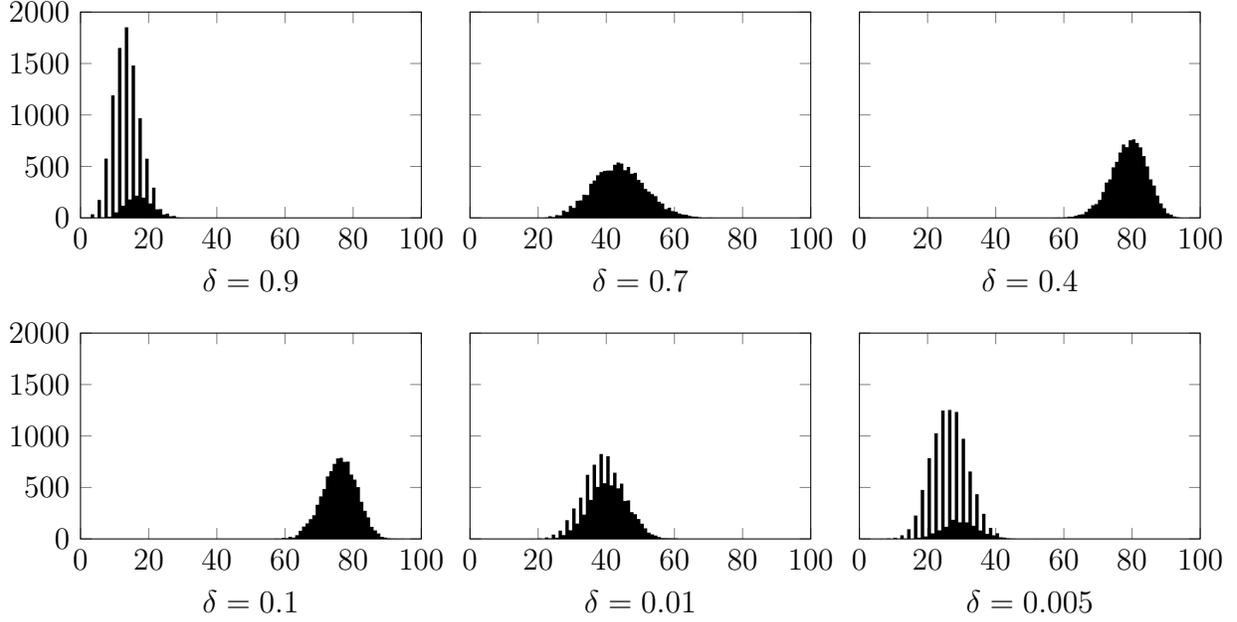


Figure 13: Number of coteries over 10,000 generating sets P drawn from the uniform distribution where $n = 100$. The unfavorable range roughly $0.1 \leq \delta \leq 0.4$ is characterized by a very high mean (≈ 80). As δ moves outside this range however, the mean number of coteries drops radically. As $\delta \rightarrow 1$, even numbers of coteries become exceedingly rare and as $\delta \rightarrow 0$, odd numbers become rare.

$\Omega = [0, 1] \setminus [0, \delta]$ where $\delta \in [0, 1]$ represents a kind of aspect ratio. Here, we consider consider density of 133.33 generators per unit area (corresponding to 100 agents when $\delta = \frac{1}{2}$) and vary δ . See Figure 15. The histograms are shown in Figure 16. Here, in all cases, considering only coteries rather than all generators represents a huge savings.

7.3 The DDN dual diagram and VTP implementation

To define the Voronoi analog dual to the DDN, we need the notion of visibility subsets.

Definition 7.3.1. If $\Omega \subseteq \mathbb{R}^2$ is precompact and $p \in \Omega$, we define *the visible part of Ω from p* by $\Omega|_p = \{q \in \Omega : \overline{pq} \subset \Omega\}$. If Ω is a polygon, then $\Omega|_p$ is called the *visibility polygon* inside Ω from p .

Lemma 7.3.1. If $\Omega \subseteq \mathbb{R}^2$, then $\Omega|_p = \bigcup\{C \subseteq \Omega : p \in C \text{ and } C \text{ convex}\}$. for all $p \in \Omega$.

Proof. The proof is elementary. If $q \in \Omega|_p$, then by definition $\overline{pq} \subset \Omega$. But $p \in \overline{pq}$ and \overline{pq}

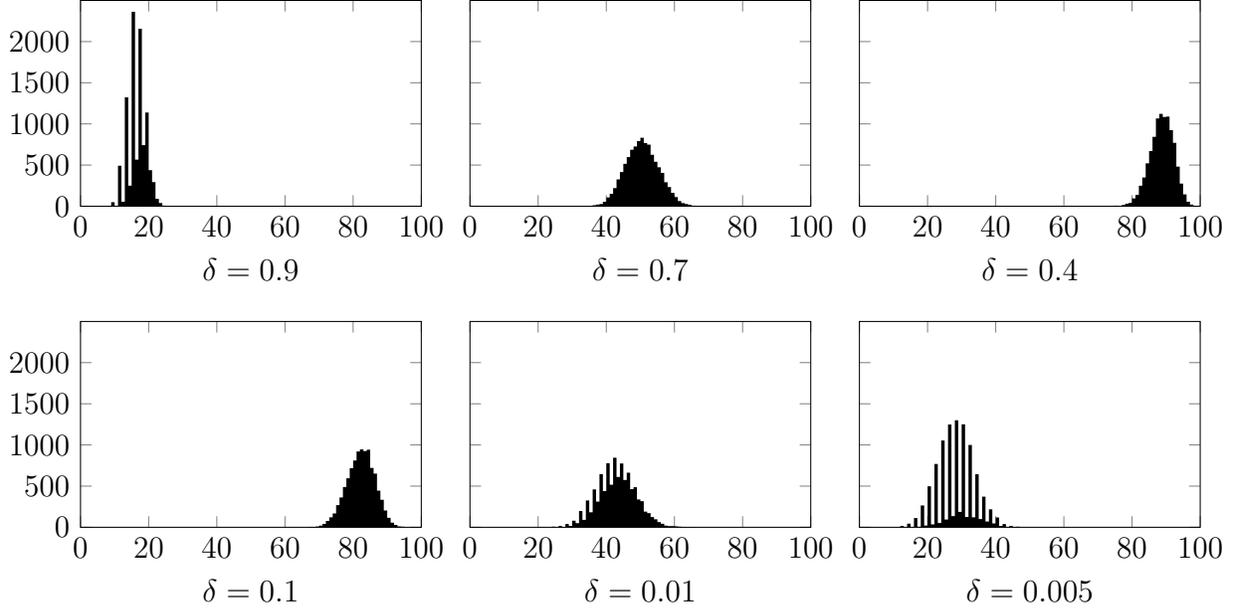
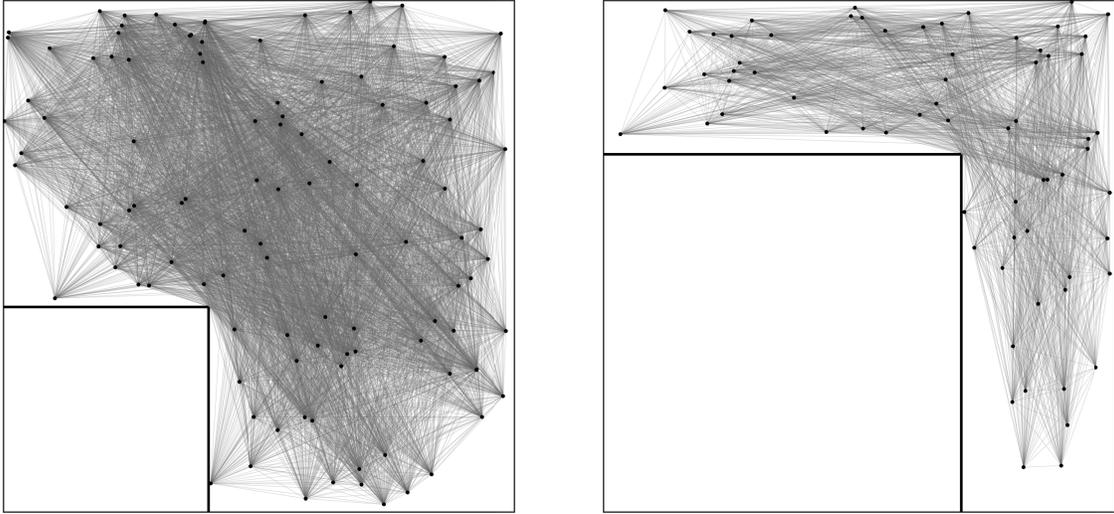


Figure 14: Number of coteries over 10,000 quasirandom generating sets P where $n = 100$. Notice the same qualitative features as the uniformly distributed samples.

is convex and therefore $C = \overline{pq}$ witnesses the union. Since $q \in \overline{pq}$, this shows $\Omega|_p \subseteq \bigcup\{C \subseteq \Omega : p \in C \text{ and } C \text{ convex}\}$. On the other hand, if q belongs to the union of convex sets, then q and p both belong to some convex set $C \subseteq \Omega$. Then by convexity, $\overline{pq} \subseteq C \subseteq \Omega$ and $q \in \Omega|_p$. \square

The directed Delaunay network gives us a particular analog of the Voronoi diagram which we call the directed Delaunay dual diagram and denote by $\text{DDN}^*(P)$. It is a collection of regions $\{V_p\}_{p \in P}$ defined as follows. We say that q is a directed Voronoi neighbor of p if $p \rightarrow q$ in $\text{DDN}(P)$. For every $p \in P$, we define the proto-cell \tilde{V}_p to be the (standard) Voronoi cell about p in \mathbb{R}^2 generated by p and its directed neighbors. Then $V_p = \tilde{V}_p \cap \Omega|_p$. Figure 17 shows the DDN^* dual diagram for the continuing example. The diagram lacks many of the desirable properties of the standard Voronoi diagram. In general, cells overlap and their union need not be all of Ω . However, it is very natural in the sense that V_p is precisely the standard Voronoi cell about p in $\Omega|_p$ generated by $P \cap \Omega|_p$, which is all p “knows about”. Clearly, certain properties of the $\text{DDN}^*(P)$ cells require computation of $\Omega|_p$ for arbitrary $p \in P$. The literature on computing visibility polygons is quite rich (c.f. [18]). However,



(a) Generating set P from uniform distribution with $\delta = 0.4$.

(b) Generating set P from uniform distribution with $\delta = 0.7$.

Figure 15: Example initializations for the corner experiment.

even in the case of simple polygons, to do this for every generator will be costly ($O(nm)$ where m is the total number of vertices of Ω). Instead, we can modify VTP to avoid computation of the $\text{DDN}^*(P)$ cells all together. Implementation of VTP with the speed scales ρ_i given by Equation (17) only requires computation of the proto-cells \tilde{V}_i . We then need to find the intersections of the ray $\{x_i + lu_i : l \geq 0\}$ with $\partial\tilde{V} \cup \partial\Omega$. The intersection closest to x_i will necessarily lie in $\Omega|_{x_i}$ since, by definition, there are no boundary points of Ω between this intersection and x_i .

Aside from this, VTP works just as in convex domains but everywhere we considered Voronoi neighbors before, we now consider directed Voronoi neighbors. The normalization factor $\#\mathcal{N}_i/6$ appearing in the alignment term still applies—even though we cannot say anything about the average number of directed Voronoi neighbors—because of the guiding principle: every agent behaves as if its Voronoi diagram is standard.

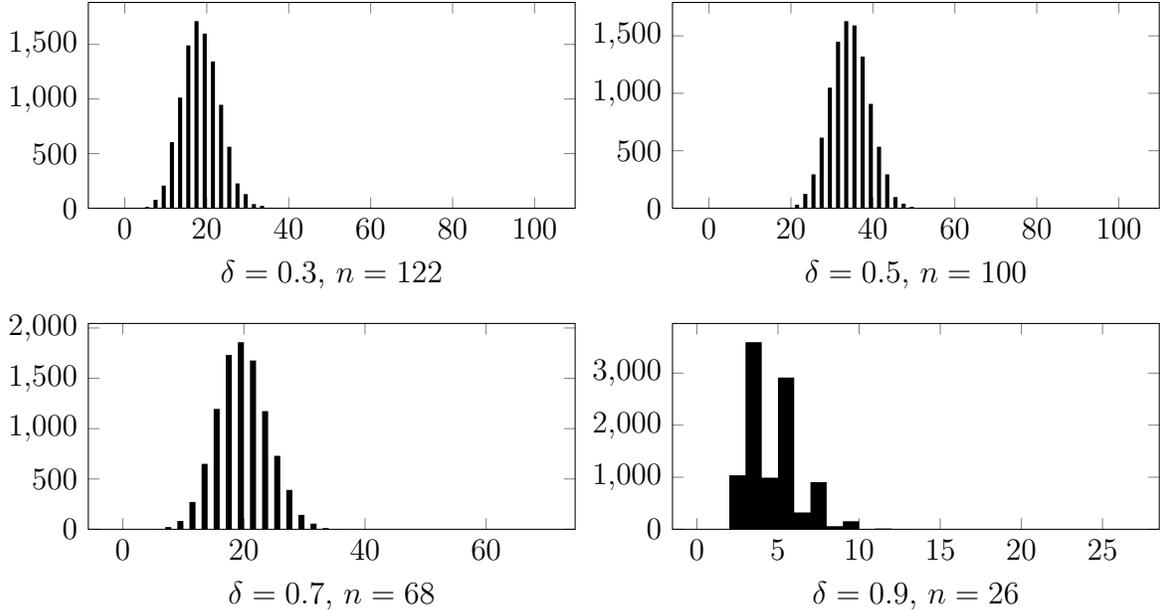


Figure 16: Number of coterie nodes in the right-angle corner domain over 10,000 random generating sets P where $n/\text{Area}(\Omega) = 133.33$ for all δ .

7.4 Summary

We have examined the challenges of extending the VTP framework to non-convex domains and presented a novel generalization for the Delaunay triangulation, the directed Delaunay network, as a solution. Although the DDN is motivated by the crowd modelling problem, its structure and corresponding dual structure may be of much broader applicability and interest, like most Voronoi/Delaunay generalizations. Moreover, the DDN construction suggests a handful of compelling avenues of research including the following. For the application to VTP simulation, we hope to develop a GPU (graphics processing unit) algorithm for computing DDNs leveraging the highly parallelizable nature of the construction. The notion of coterie nodes in a geometric graph which arises in the DDN construction is also potentially of broader interest in problems (like VTP) involving physically distributed nodes with possibly obstructed communications. We have here presented a brief study of the coterie properties of such networks with random and quasi-random nodes in the *two-room* and *corner* experiments.

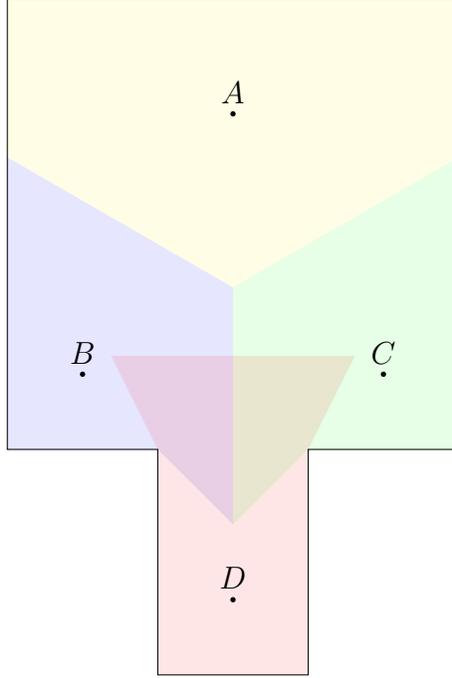


Figure 17: Directed Delaunay dual diagram $\text{DDN}^*(P)$.

A Lemmas

Lemma A.0.1. Let $V \subseteq \mathbb{R}^2$ be a measurable set and let $p \in \mathbb{R}^2$. Define $S(\phi)$ to be the half plane whose boundary contains p and whose inward normal makes an angle ϕ with some fixed reference line, say, the horizontal (inclusion of boundary points is irrelevant). Then, the average measure of $V \cap S(\phi)$ over $\phi \in [0, 2\pi)$ is half the measure of V .

Proof. Without loss of generality, take p to be the origin. Let $|\cdot|$ denote measure. Notice that $S(\phi)$ and $S(\phi + \pi)$ overlap on a set of measure zero for all ϕ and $S(\phi) \cup S(\phi + \pi)$ is the

entire plane up to a set of measure zero. Then the average measure of $V \cap S(\phi)$ is given by

$$\begin{aligned}
\frac{1}{2\pi} \int_0^{2\pi} |V \cap S(\phi)| d\phi &= \frac{1}{2\pi} \left(\int_0^\pi |V \cap S(\phi)| d\phi + \int_\pi^{2\pi} |V \cap S(\phi)| d\phi \right) \\
&= \frac{1}{2\pi} \int_0^\pi (|V \cap S(\phi)| + |V \cap S(\phi + \pi)|) d\phi \\
&= \frac{1}{2\pi} \int_0^\pi |V \cap (S(\phi) \cup S(\phi + \pi))| d\phi \\
&= \frac{1}{2\pi} \int_0^\pi |V| d\phi \\
&= \frac{1}{2} |V|,
\end{aligned}$$

as desired. □

Lemma A.0.2. If Ω is a compact, connected 2-manifold without boundary of Euler characteristic $\chi \geq 0$, then the average number of Voronoi edges per Voronoi cell in a Voronoi diagram is at most 6 (except possibly for degenerate configurations).

Proof. Let v , e , and n be the total numbers of vertices, edges, and cells respectively in a the Voronoi diagram. Assume that every edge has exactly two vertices (this is the nongeneracy assumption⁸) so that $2e \geq 3v$. By definition $v - e + n = \chi$. Combining these we have that $e \leq 3(n - \chi)$. Then, since every edge belongs to exactly two cells, the average number of edges per cell obeys

$$\frac{2e}{n} \leq \frac{6(n - \chi)}{n} = 6 - \frac{\chi}{n} < 6.$$

□

Lemma A.0.3. The average number of neighbors per Voronoi cell in a Voronoi diagram in a compact domain $\Omega \subset \mathbb{R}^2$ is at most six.

Proof. Let v and n be the total numbers of vertices and cells respectively in the Voronoi diagram. Let e be the number of edges where two Voronoi cells meet and e_b be the number

⁸the only degenerate configurations in the sphere are those with fewer than three generators. In the torus, the only degenerate configurations are those in which all generators lie on a single minimal geodesic in which case the edges will be parallel and closed, having no vertices.

of boundary edges, i.e., edges where a Voronoi cell meets the boundary of Ω . Since every edge has exactly two vertices and at least three edges meet at every vertex, $2(e + e_b) \geq 3v$. Recall that the Euler characteristic of the plane is 2 so $v - (e + e_b) + (n + 1) = 2$, where the $n + 1$ accounts for the one (unbounded) face $\mathbb{R}^2 \setminus \Omega$ to complete the planar graph. Combining these we have that $e + e_b \leq 3(n - 1)$. Then, the average number of neighbors per cell is exactly twice the number of non-boundary edges divided by the number of Voronoi cells and obeys

$$\frac{2e}{n} \leq \frac{2(e + e_b)}{n} \leq \frac{6(n - 1)}{n} = 6 - \frac{6}{n} < 6.$$

□

Lemma A.0.4. For all $X \in \Omega^n$ and $U \in \mathbb{R}^{2n}$,

$$\frac{1}{n} \sum_i \frac{|f_i(X, U)|}{\rho_i(X, U)} = \frac{1}{n} \sum_i \frac{|\bar{\sigma}_i h_i + \sigma_i u_i + \nu a_i|}{b_i \bar{\sigma}_i + \sigma_i + \nu} \leq \frac{2 + \nu}{1 + \nu}.$$

Moreover, if $b_i = 1$ for all i , this bound is tightened to 1.

Proof. From thier definitions, it is clear that $|h_i|$, $|u_i|$, and $|\tilde{a}_i|$ are all at most 1. So,

$$\begin{aligned} & \frac{1}{n} \sum_i \frac{|\bar{\sigma}_i h_i + \sigma_i u_i + \nu a_i|}{b_i \bar{\sigma}_i + \sigma_i + \nu} \\ &= \frac{1}{n} \sum_i \frac{|b_i \bar{\sigma}_i \hat{h}_i + \sigma_i u_i + \nu \varphi_i \tilde{a}_i|}{b_i \bar{\sigma}_i + \sigma_i + \nu} \\ &\leq \frac{1}{n} \sum_i \frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} + \frac{1}{n} \sum_i \frac{\nu \varphi_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} \\ &\leq \frac{1}{n} \sum_i \frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} + \frac{1}{n} \sum_i \varphi_i \\ &= \frac{1}{n} \sum_i \frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} + \frac{1}{6n} \sum_i \#\mathcal{N}_i \\ &\leq \frac{1}{n} \sum_i \frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} + 1 \end{aligned} \tag{21}$$

where the last inequality follows from Lemma A.0.2. For the remaining sum, consider the

$b_i = 0$ and $b_i = 1$ cases. If $b_i = 1$, then

$$\frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} = \frac{1}{1 + \nu}$$

and if $b_i = 0$, then

$$\frac{b_i \bar{\sigma}_i + \sigma_i + \nu}{b_i \bar{\sigma}_i + \sigma_i} = \frac{\sigma_i + \nu}{\sigma_i} = 1 + \frac{\nu}{\sigma_i} \geq 1 + \nu$$

so

$$\frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} \leq \frac{1}{1 + \nu}.$$

So continuing from (21), have

$$\frac{1}{n} \sum_i \frac{b_i \bar{\sigma}_i + \sigma_i}{b_i \bar{\sigma}_i + \sigma_i + \nu} + 1 \leq \frac{1}{1 + \nu} + 1$$

and we are done. In the special case $b_i = 1$ for all i , we have

$$\begin{aligned} & \frac{1}{n} \sum_i \frac{|\bar{\sigma}_i \tilde{h}_i + \sigma_i u_i + \nu \tilde{a}_i|}{b_i \bar{\sigma}_i + \sigma_i + \nu} \\ &= \frac{1}{n} \sum_i \frac{|\bar{\sigma}_i h_i + \sigma_i u_i + \nu \varphi_i a_i|}{1 + \nu} \\ &\leq \frac{1}{n(1 + \nu)} \sum_i (\bar{\sigma}_i + \sigma_i + \nu \varphi_i) \\ &= \frac{1}{n(1 + \nu)} \sum_i (1 + \nu \varphi_i) \\ &= \frac{1}{1 + \nu} \left(1 + \frac{\nu}{n} \sum_i \varphi_i \right) \\ &\leq 1 \end{aligned}$$

where again, the last inequality follows from Lemma A.0.2.

□

B Technicalities concerning vector translations in the torus

The torus $\mathbb{T}_l^2 = \mathbb{R}^2/l\mathbb{Z}^2$ is not a vector space but despite this, we employ seemingly vector-like operations on points in \mathbb{T}^2 , including a operator $+$, a difference operator $-$. There is however, no scalar multiplication as there is no “zero vector” in \mathbb{T}^2 . Thus, unlike in a vector space, the $-$ and $+$ operations function entirely differently. Indeed, $- : \mathbb{T}^2 \times \mathbb{T}^2 \rightarrow \mathbb{R}^2$ takes two points in the torus and gives a *vector* in \mathbb{R}^2 (more precisely in $[-l/\sqrt{2}, l/\sqrt{2}]^2$) while $+ : \mathbb{T}^2 \times \mathbb{R}^2 \rightarrow \mathbb{T}^2$ takes a point in \mathbb{T}^2 and a vector in \mathbb{R}^2 to give another point in \mathbb{T}^2 . It is the case that for all $x, y \in \mathbb{T}^2$, we have $x + (y - x) = y$ but $y - x$ is only one of countably many vectors with this property.

For real numbers s , we write $[s]_l$ for $[s]_l = l\lfloor s/l \rfloor$. (This may be familiar to programmers as $s\%l$.) For vectors $\xi \in \mathbb{R}^d$, we write $[\xi]_l = ([\xi_1]_l, \dots, [\xi_d]_l)$ for the component-wise operation. Points in \mathbb{T}_l^2 are by definition $l\mathbb{Z}^2$ cosets over the additive group of \mathbb{R}^2 . Thus, for $x \in \mathbb{T}_l^2$, we write $\hat{x} \in \mathbb{R}^2$ for the unique representative of the coset x in $[0, l)^2$.

Then, the map $+ : \mathbb{T}_l^2 \times \mathbb{R}^2 \rightarrow \mathbb{T}_l^2$ takes a point in \mathbb{T}_l^2 and a displacement vector in \mathbb{R}^2 to yield a new point in \mathbb{T}_l^2 via

$$x + \xi = (\hat{x} + \xi) + l\mathbb{Z}^2 = [\hat{x} + \xi]_l + l\mathbb{Z}^2 \in \mathbb{T}_l^2.$$

Now, the map $- : \mathbb{T}_l^2 \times \mathbb{T}_l^2 \rightarrow \mathbb{R}^2$, gives a partial solution to the inverse problem of finding a displacement vector $\xi = y - x$ between two points such that for all $x, y \in \mathbb{T}_l^2$, we have $x + (y - x) = y$, as one would expect. This property however only distinguishes a countably infinite family of vectors. In most cases (in a sense we will make precise momentarily), there is a unique shortest possibility. With this in mind, we write

$$\|x - y\| = \min_{a \in l\mathbb{Z}^2} \|\hat{x} - \hat{y} + a\|_{\mathbb{R}^2} = \min_{a \in \{-1, 0, 1\}^2} \|\hat{x} - \hat{y} + la\|_{\mathbb{R}^2}. \quad (*)$$

The last equality can easily be seen to follow from the fact that \hat{x} and \hat{y} belong to $[0, l]^2$. So $\hat{x} - \hat{y} \in (-l, l)^2$ and shifting by $\pm kl$ in either component for $k \geq 2$ can only put us further from the origin, increasing the magnitude. If $(*)$ has a unique minimizer a^* , then we define

$$x - y = \hat{x} - \hat{y} + la^* \in \mathbb{R}^2, \quad a^* \in \arg \min_{a \in \{-1, 0, 1\}^2} \|\hat{x} - \hat{y} + la\|_{\mathbb{R}^2}$$

Indeed, for all $x, y \in \mathbb{T}_l^2$,

$$\begin{aligned} y + (x - y) &= y + (\hat{x} - \hat{y} + la^*) \\ &= (\hat{y} + \hat{x} - \hat{y} + la^*) + l\mathbb{Z}^2 \\ &= (\hat{x} + la^*) + l\mathbb{Z}^2 \\ &= \hat{x} + l\mathbb{Z}^2 \\ &= x \end{aligned}$$

The question remains, for which $r \in \mathbb{R}^2$ does the minimization problem $\arg \min_{a \in \{-1, 0, 1\}^2} \|r + la\|_{\mathbb{R}^2}$ fail to have a unique solution? To get a feel for this, suppose $r = (0, l/2)$. Then we find $\arg \min_{a \in \{-1, 0, 1\}^2} \|r + la\|_{\mathbb{R}^2} = \{(0, 0), (0, -1)\}$ because $\|r + (0, 0)\|_{\mathbb{R}^2} = \|r - (0, l)\|_{\mathbb{R}^2} = l/2$. This reflects the fact that there are two shortest displacements to translate $(0, 0) + l\mathbb{Z}^2$ to $(0, l/2) + l\mathbb{Z}^2$, namely translation by $(0, l/2) = r + (0, 0)$ or by $(0, -l/2) = r - (0, l)$.

To answer this question, consider the following. For two points $\xi, \xi' \in \mathbb{R}^2$, let $J(\xi, \xi')$ be the set

$$J(\xi, \xi') = \{\zeta \in \mathbb{R}^2 : \|\zeta + \xi\| = \|\zeta + \xi'\|\}$$

such that for $\zeta \in J(\xi, \xi')$, the translates $\zeta + \xi$ and $\zeta + \xi'$ are equal in magnitude. We can

find an explicit form for $J(\xi, \xi')$ as follows.

$$\begin{aligned}\zeta \in J(\xi, \xi') &\iff \|\zeta + \xi\| = \|\zeta + \xi'\| \\ &\iff \|\zeta + \xi\|^2 - \|\zeta + \xi'\|^2 = 0 \\ &\iff 2\zeta \cdot (\xi - \xi') + \|\xi\|^2 - \|\xi'\|^2 = 0.\end{aligned}$$

Thus, $J(\xi, \xi')$ is a line for all $\xi \neq \xi'$.

Clearly, if a^*, b^* are distinct solutions of $\arg \min_{a \in \{-1, 0, 1\}^2} \|\zeta + la\|$, then we have $\|\zeta + la^*\| = \min_{a \in \{-1, 0, 1\}^2} \|\zeta + la\| = \|\zeta + lb^*\|$ so therefore $\zeta \in J(la^*, lb^*)$. Thus, we have a necessary condition on ζ for the existence of distinct solutions of $\arg \min_{a \in \{-1, 0, 1\}^2} \|\zeta + la\|$, namely that $\zeta \in J(la, lb)$ for some $a, b \in \{-1, 0, 1\}^2$. That is

$$\zeta \in \bigcup_{\substack{a, b \in \{-1, 0, 1\}^2 \\ a \neq b}} J(la, lb).$$

The germane observation here, is that the set on the right hand side is a union of $\binom{9}{2} = 36$ lines in the plane (because $J(\xi, \xi')$ is symmetric in ξ and ξ').

So we have that $x - y$ for $x, y \in \mathbb{T}_l^2$ is well-defined except possibly if $\hat{x} - \hat{y} \in \bigcup_{a, b \in \{-1, 0, 1\}^2} J(la, lb)$.

Denote $S = [0, l]^2 \times [0, l]^2 \cong [0, l]^4 \subset \mathbb{R}^4$. Then

$$\begin{aligned}\mathcal{M} &= \left\{ (u, v) \in S : u - v \in \bigcup_{\substack{a, b \in \{-1, 0, 1\}^2 \\ a \neq b}} J(la, lb) \right\} \\ &= \bigcup_{\substack{a, b \in \{-1, 0, 1\}^2 \\ a \neq b}} \left\{ (u, v) \in S : u - v \in J(la, lb) \right\} \\ &= \bigcup_{\substack{a, b \in \{-1, 0, 1\}^2 \\ a \neq b}} \left\{ (u, v) \in S : 2(u - v) \cdot (a - b) + l(\|a\|^2 - \|b\|^2) = 0 \right\}.\end{aligned}$$

Writing $(u, v) = (u_1, u_2, v_1, v_2)$, the equations defining each of the 36 sets in the above union all take the form $C_1(u_1 - v_1) + C_2(u_2 - v_2) = C$ for real constants C, C_1, C_2 (depending on

a and b) where C_1 and C_2 are not both zero (because $a \neq b$), and thus define 36 smooth three-dimensional hypersurfaces in S .

C Alternate metrics

C.1 Ray-marching algorithm for computing approximate Voronoi diagrams in arbitrary metrics

In many discretized Voronoi algorithms, the domain is rasterized, every pixel is scanned, its metric distance to generators determined, and its Voronoi cell determined.

The algorithm presented here offers an alternate approach. The result is still an approximation of the true Voronoi diagram but we do not rasterize the domain. Instead, the procedure locates approximate points on the Voronoi edges within a user-specified tolerance (bounding both their distance from the true Voronoi edge and from each other). The essential idea is to emanate rays from each generator at constant speed⁹ in the metric with respect to which we define the Voronoi diagram. We stop pairs of rays from different generators if they come sufficiently near each other. Both points where rays have halted are taken to lie on the approximate Voronoi edge. The details involve setting the angular resolution of the emanating rays so that they are sufficiently dense when approaching the true Voronoi edges that the “collisions” are detected.

In more detail, we have a differentiable metric d , the generators x_1, \dots, x_k , and the tolerance $\varepsilon > 0$. We compute the angular tolerance $\delta > 0$ which depends on ε and $\text{diam}(\Omega)$ and a refined tolerance $\tilde{\varepsilon}$ depending on ε and d . We initialize a time parameter t at 0 and compute its increment Δt which will depend on ε and (partial derivatives of) d . For each generator x_i we initialize “rays” at angles of increment at most δ . Each ray consists of the identity of its generator, its (fixed) direction, its current (Euclidean) distance from

⁹Note that we can take the speed of the emanating rays to be 1 without loss of generality because changing this is equivalent to globally rescaling the metric which has no bearing on the Voronoi diagram. The crucial feature is that the speed is the same for every ray and every generator.

its generator (initially 0), and its status (propogating or terminated). We then iterate, propogating all (not yet terminated) rays in their direction with speed 1 *as measured in* d over time Δt . At each iteration, we terminate all rays which have crossed the domain boundary or are within d -distance $\tilde{\varepsilon}$ of a ray from a different generator.

We need to determine bounds on Δt and δ to ensure that (1) no ray propogates more than a distance $\tilde{\varepsilon}$ in d each iteration and (2) rays are dense enough that they pass within $\tilde{\varepsilon}$ of each other near Voronoi edges. Further $\tilde{\varepsilon}$ itself must be bounded so that the resulting points are within ε in the Euclidean metric of the Voronoi edges.

For each generator i and direction α , let n_α be the unit vector in the direction (indexed) α and define the function $u_{i\alpha} : [0, b] \rightarrow [0, \infty)$ where $b = \sup\{b' \geq 0 : x_i + b'n_\alpha \in \Omega\}$ by

$$u_{i\alpha}(s) = d(x_i, x_i + sn_\alpha).$$

That is, $u_{i\alpha}(s)$ gives the d -distance from x_i of the (i, α) ray after propogating a Euclidean distance s . The speed constraint on all rays amounts to varying the argument s with t so that

$$1 = \frac{d}{dt}u_{i\alpha}(s(t)) = \dot{s}u'_{i\alpha}(s)$$

is satisfied for each i and α . Rewrtting, we have the (independent) nonlinear ordinary differential equations

$$\dot{s}_{i\alpha} = \frac{1}{u'_{i\alpha}(s_{i\alpha})}, \quad s_{i\alpha}(0) = 0 \quad \forall i, \alpha$$

defining the growth paramters $s_{i\alpha}$. These equations give us a bound on Δt as the quantity

$$s(t + \Delta t) - s(t) = \int_t^{t+\Delta t} \dot{s}(\tau) d\tau \leq \Delta t \left(\inf_{j,\beta} \inf_{s \in \text{dom } u_{j\beta}} u'_{j\beta}(s) \right)^{-1} =: \Delta s$$

must satisfy

$$\sup_{\|x-y\| \leq \Delta s} d(x, y) \leq \tilde{\varepsilon}$$

to meet condition (1). If d is homogeneous and translation invariant, then it is metrically equivalent to the Euclidean distance d_2 and there are positive constants c, C such that $cd \leq d_2 \leq Cd$. Then the above becomes

$$\tilde{\varepsilon} \geq \sup_{d_2(x,y) \leq \Delta s} d(x,y) \geq \frac{1}{C} \sup_{d_2(x,y) \leq \Delta s} d_2(x,y) = \frac{1}{C} \Delta s = \frac{1}{C} \Delta t \left(\inf_{j,\beta} \inf_{s \in \text{dom } u_{j\beta}} u'_{j\beta}(s) \right)^{-1}$$

and

$$\Delta t \leq C \tilde{\varepsilon} \cdot \inf_{j,\beta} \inf_{s \in \text{dom } u_{j\beta}} u'_{j\beta}(s)$$

Writing $U_x : \Omega \rightarrow [0, \infty)$ by $U_x(y) = d(x, y)$, we may strengthen the bound on Δt to

$$\Delta t \leq \varepsilon \cdot \inf_{x,y \in \Omega} \|\nabla U_x(y)\| \leq \varepsilon \cdot \inf_{j,\beta} \inf_{s \in \text{dom } u_{j\beta}} u'_{j\beta}(s)$$

if the infimum on the left is more convenient than the one on the right. It may seem counterintuitive that we require $\inf u'_{j\beta} > 0$ over the whole domain so let's make sure we understand why this is so. If along some ray, $u_{j\beta}$ increases very slowly, then unit speed propagation of the ray with respect to d amounts to relatively fast propagation with respect to the Euclidean distance and thus, forces shorter time steps to keep the Euclidean step length small. So the more slowly $u_{j\beta}$ increases, the smaller Δt must be.

the bound on δ is more complicated. to see where it arises, we need that for every point y on a Voronoi edge, rays from each neighboring generator pass within $\tilde{\varepsilon}$ of y and of each other simultaneously and, secondly, that whenever two rays from distinct generators simultaneously pass within $\tilde{\varepsilon}$ of each other, they are each within $\tilde{\varepsilon}$ of the corresponding Voronoi edge.

We use latin indices i, j, \dots for generators and greek indices α, β, \dots for ray directions (that is, greek letters index the angles at increments of at most δ). Generator-direction pairs (i, α) are then in correspondence with rays. For each generator-direction pair (i, α) let $r_{i\alpha}^t \in \Omega$ be the point in direction α from its generator x_i where $d(x_i, r_{i\alpha}^t) = t$. Also let

$\vec{r}_{i\alpha} = \{r_{i\alpha}^t : t \geq 0\}$ be the entire (i, α) ray. Suppose then that at iteration T , the (i, α) ray and the (j, β) ray are within $\tilde{\varepsilon}$ of each other in d , that is, $d(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t}) < \tilde{\varepsilon}$. Then

$$\begin{aligned} d(x_i, r_{j\beta}^{T\Delta t}) - d(x_j, r_{j\beta}^{T\Delta t}) &\leq d(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t}) + d(x_i, r_{i\alpha}^{T\Delta t}) - d(x_j, r_{j\beta}^{T\Delta t}) \\ &= d(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t}) + T - T \\ &\leq \tilde{\varepsilon} \end{aligned}$$

and

$$\begin{aligned} d(x_j, r_{j\beta}^{T\Delta t}) - d(x_i, r_{j\beta}^{T\Delta t}) &\leq d(x_j, r_{j\beta}^{T\Delta t}) - (d(x_i, r_{i\alpha}^{T\Delta t}) - d(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t})) \\ &= T - T + d(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t}) \\ &\leq \tilde{\varepsilon} \end{aligned}$$

so $|d(x_i, r_{j\beta}^{T\Delta t}) - d(x_j, r_{j\beta}^{T\Delta t})| \leq \tilde{\varepsilon}$. Then by continuity of the metric, $r_{j\beta}^{T\Delta t}$ is within $\tilde{\varepsilon}$ (w.r.t. d) of the curve $\Gamma_{ij} = \{z \in \Omega : d(x_i, z) - d(x_j, z) = 0\}$ which contains the ij Voronoi edge. Knowing then that $r_{j\beta}^{T\Delta t}$ is within $\tilde{\varepsilon}$ of Γ_{ij} , it is either within $\tilde{\varepsilon}$ of the ij Voronoi edge or it is in a Voronoi cell other than i or j 's. But in the latter case we have the following. If $r_{j\beta}^{T\Delta t}$ lies in a third Voronoi cell k , then it must have already (i.e., at some iteration $T' \leq T$) crossed the jk Voronoi edge, where it would have been terminated upon grazing some (k, γ) ray at a point $r_{k\gamma}^{T'\Delta t}$. So it must be that $T' = T$ (since by assumption the ray is still propogating up to time T) and $r_{i\alpha}^{T\Delta t}$, $r_{j\beta}^{T\Delta t}$, and $r_{k\gamma}^{T\Delta t}$ are all within $\tilde{\varepsilon}$ of each other *simultaneously*, putting them within $\tilde{\varepsilon}$ of the ijk Voronoi vertex. Exactly the same holds for $r_{i\alpha}^{T\Delta t}$.

Now we need to check that every point along a Voronoi edge is seen by some pair of rays simultaneously. Here it will be crucial to distinguish not only between d - and Euclidean distances but also between derived figures such as circles and balls.

For all $r > 0$, define the quantities

$$m_r := \inf\{d_2(x, y) : x, y \in \Omega, d(x, y) \geq r\},$$

$$M_r := \sup\{d_2(x, y) : x, y \in \Omega, d(x, y) \leq r\}.$$

Note that $m_r \leq M_r$ for all r . This is because otherwise, if $M_r < m_r$, then fix $M_r < a < m_r$. By definition of M_r and m_r , pairs of points in Ω achieve d_2 separations in both intervals $(M_r, a]$ and $[a, m_r)$. By convexity then, there are points $x, y \in \Omega$ where $d_2(x, y) = a$. But then, since $d_2(x, y) = a < m_r$, necessarily $d(x, y) < r$ and, similarly, since $d_2(x, y) = a > M_r$, we have $d(x, y) > r$ and we have a contradiction.

Now, fix $y \in \Omega$ on the ij Voronoi edge, say. Consider the continuous pulse fronts $\Phi_i^t = \{z \in \Omega : d(x_i, z) = t\}$. Let $t^* = d(x_i, y) = d(x_j, y)$. Let T be the greatest integer such that $T\Delta t \leq t^*$ and write $y_i \in \Phi_i^{T\Delta t} \cap \overline{x_i y}$ (this intersection contains exactly one point) and y_j similarly. We can take Δt small enough that $\Delta s \leq m_\varepsilon/4$. Then $d_2(y_i, y) < \Delta s \leq m_\varepsilon/4$. The ray (i, α) nearest to y_i is at an angular separation of $\theta \leq \delta/2$. Since $d(x_i, r_{i\alpha}^{T\Delta t}) = d(x_i, y_i) = T\Delta t$, we have by Lemma C.1.1 that $d_2(r_{i\alpha}^{T\Delta t}, y_i) \leq \text{diam}(\Omega)^2 K(\frac{\delta}{2}, \kappa)$ so long as $\delta/2 < \arcsin \frac{1}{\kappa}$ where κ is the condition number of d . By choosing δ small enough, we can make this as small as we like, in particular, so that $d_2(r_{i\alpha}^{T\Delta t}, y_i) \leq m_\varepsilon/4$. Doing the same for j , we have by the triangle inequality

$$\begin{aligned} d_2(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t}) &\leq d_2(r_{i\alpha}^{T\Delta t}, y_i) + d_2(y_i, y) + d_2(y, y_j) + d_2(y_j, r_{j\beta}^{T\Delta t}) \\ &< \frac{m_\varepsilon}{4} + \frac{m_\varepsilon}{4} + \frac{m_\varepsilon}{4} + \frac{m_\varepsilon}{4} \\ &= m_\varepsilon \end{aligned}$$

for some α and β . By definition then, $d(r_{i\alpha}^{T\Delta t}, r_{j\beta}^{T\Delta t}) < \tilde{\varepsilon}$, as desired. Finally we must take $\tilde{\varepsilon}$ so that $M_{\tilde{\varepsilon}} \leq \varepsilon$.

Considering the runtime, because rays propagate at unit speed with respect to d , each

ray travels a d -distance of Δt each step and the total number of iterations is at worst $\text{diam}_d(\Omega)/\Delta t$ where $\text{diam}_d(\Omega) = \sup_{x,y \in \Omega} d(x,y)$. The number of rays emanating from each generator is $2\pi/\delta$ and at each iteration t , we must compute for every ray (i, α) the integral $\int_t^{t+\Delta t} \dot{s}_{i\alpha}(\tau) d\tau$. Say this problem takes computation time M . Finally, we need to find the n -th nearest neighbors in d among all $k \cdot 2\pi/\delta$ propagating rays to test the termination condition. (If we can dominate d by some scaling of the Euclidean metric, we can speed up this step but at the cost of a tighter tolerance $\tilde{\varepsilon}$ which will multiplicatively increase the number of comparisons we must perform. The best way to do this depends entirely on the ease of evaluating d .) The main draw of this algorithmic approach is that one only needs global information about the metric during the preprocessing stages when computing the condition number in order to determine the angular resolution δ . During the execution of the ray-marching phase, only local information—the partial derivative of the metric and presence of other rays sufficiently nearby—is required.

One might point out that the standard rasterization approach to computing approximate Voronoi diagrams locates the approximate interiors of Voronoi cells as well as edges while the algorithm presented above only locates edges. To locate interiors, we need only the following simple modification. At each iteration, after checking the termination conditions on each ray, if the ray is not yet terminated, assign its generator to its current location. At the end, every point where a ray was checked will have been assigned a generator or become an edge point. Interestingly, unlike the rasterization approach, the resolution of the resulting image will not be constant but rather, will be highest at the generators and maintain constant *angular* resolution about each generator within its Voronoi cell. The resolution of the image is as good or better than one dot per area ε^2 everywhere.

Lemma C.1.1. Suppose $\Omega \subset \mathbb{R}^2$ is compact, convex, and $0 \in \Omega$. Further assume Ω is equipped with a metric $d : \Omega \rightarrow \Omega$ that is continuously differentiable in each argument (except possibly where $d = 0$) and satisfies the well-conditioning condition described below. Then if $\theta > 0$ is sufficiently small, there is a constant K depending only on d , θ , and $\text{diam}_2(\Omega)$

such that if $y, y' \in \Omega$ are d -equidistant from the origin and have angular separation at most θ (with respect to the origin), then $\|y - y'\| \leq K$.

Proof. Write $\varphi_x(y) = d(x, x + y)$ for the univariate function obtained by fixing the first argument of d (and translating the second so that $\varphi_x(0) = 0$) and define

$$G = \sup_{x \in \Omega} \sup_{y \in (\Omega - x) \setminus \{0\}} \|\nabla \varphi_x(y)\| \quad \text{and} \quad g = \inf_{x \in \Omega} \inf_{y \in (\Omega - x) \setminus \{0\}} \partial_r \varphi_x(y)$$

where ∂_r denote the partial derivative in the radial direction. Note that G is finite because $\nabla \varphi_x$ is continuous and Ω is compact. That $g > 0$ is the *well-conditioning* requirement mentioned above. Then, if d is well-conditioned, write

$$\kappa = \frac{G}{g}.$$

In the following, we take $x = 0$ so that $\varphi_x(y) = \varphi_0(y) = d(0, y)$ for all y . We have defined φ_x more generally so that the bound we prove is translation invariant. That is, because G and g do not depend on x , we may take the origin to be any point of Ω and the derived bound will hold. Let y and y' in Ω with angular separation $\theta > 0$ and $\varphi_0(y) = \varphi_0(y')$. Assume $\|y'\| \geq \|y\|$. Define

$$h := \|y'\| - \cos \theta \|y\| = \|y' - \text{proj}_{y'}(y)\|.$$

We would like to find an explicit upper bound for h in terms of $\|y\|$, θ , and κ . We will then weaken this bound so as not to depend on $\|y\|$ then derive a bound on $\|y - y'\|$ from it. First, let's explain intuitively why the bound on h ought to exist given the assumptions. Using differentiability and the fact that φ_0 has no stationary points (because $g > 0$), the Mean Value Theorem guarantees a point w^* between y and y' where $\nabla \varphi_0(w^*)$ is perpendicular to $y' - y$. So by increasing h , the orientation of $y' - y$ becomes closer to radial and therefore $\nabla \varphi_0(w^*)$ closer to tangential. If θ is small enough, then eventually the gradient at w^* is so close to being tangential that its radial component cannot remain larger than g without its

own magnitude exceeding G . With this in mind, let us treat the proof in detail.

Let $R = \|y\|$. Let $w(\phi)$ for $\phi \in [0, \theta]$ parametrize the segment $\overline{yy'}$ so that the angle between $w(\phi)$ and y is ϕ . By the Mean Value Theorem, there is some $\phi^* \in [0, \theta]$ such that $w^* := w(\phi^*)$ satisfies $\nabla\varphi_0(w^*) \cdot (y - y') = \varphi_0(y) - \varphi_0(y') = 0$. By assumption, $y - y' \neq 0$ and because $\|\nabla\varphi_0\| \geq g$ everywhere, $\nabla\varphi_0(w^*) \neq 0$. Therefore, $\nabla\varphi_0(w^*)$ is perpendicular to $y - y'$. Let n be the unit normal to $\overleftrightarrow{yy'}$ which points away from the origin at w^* and let $r(\phi) = w(\phi)/\|w(\phi)\|$ be the unit vector in the radial direction at $w(\phi)$ and let $\alpha(\phi)$ be the angle between n and $r(\phi)$. See the diagram below. We have just shown that at w^* , the angle $\alpha^* = \alpha(\phi^*)$ is given by $\cos(\alpha^*) = \frac{\partial_r \varphi_0(w^*)}{\|\nabla\varphi_0(w^*)\|} \geq \frac{g}{G} = \kappa^{-1}$. We will show that if h is too large, this inequality will fail for all $\alpha(\phi)$.

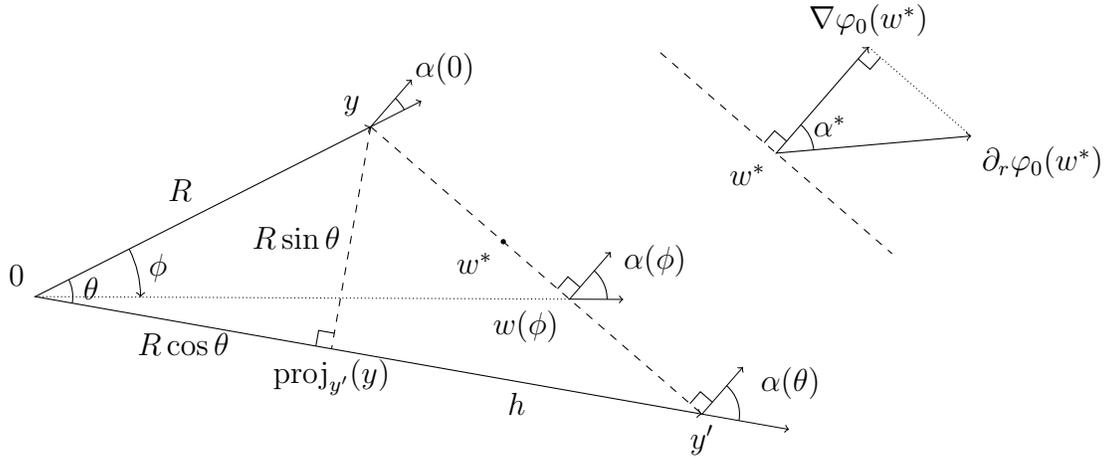


Figure 18: Diagram for the proof of Lemma C.1.1. By the Mean Value Theorem, for all h , there exist some w^* on $\overline{yy'}$ for which we have the situation depicted at the top right. The Mean Value Theorem and the finite condition number of d therefore impose a bound on α^* , and thus on h and $\|y - y'\|$.

Consider the right triangle Σ formed by 0 , y , and $\text{proj}_{y'}(y)$. Again, the angle at 0 is θ and the sides incident to 0 are of length $\|y\| = R$ and $\|\text{proj}_{y'}(y)\| = R \cos \theta + h$. Demand $\theta < \pi/2$. If $0 \leq h \leq R \tan \theta \sin \theta$, then it is easy to see that Σ is not obtuse and therefore the altitude from the vertex at 0 intersects the opposite side. At this point of intersection, $\alpha = 0$. So in order to violate the bound $\cos(\alpha) \geq \kappa^{-1}$ on the whole segment, we need only consider $h > R \tan \theta \sin \theta$, for which Σ has an obtuse angle at y . For h in this range, the angle at y

is obtuse and α increases monotonically as we move from y to y' , precisely, $\alpha(\phi) = \alpha(0) + \phi$ for all $\phi \in [0, \theta]$. To meet our condition then, we need that only the lower bound for $\alpha(0)$. So it suffices to ensure $\cos(\alpha_0) < \kappa^{-1}$ where $\alpha_0 = \alpha(0)$. Notice that $\alpha_0 = \gamma - \frac{\pi}{2}$ where γ is the internal angle of our triangle at y . Hence $\cos(\alpha_0) = \sin \gamma = \frac{\|y'\|}{\|y-y'\|} \sin \theta$ by the law of sines. But also, using the right triangle formed by y , $\text{proj}_{y'}(y)$, and y' , we see that $\|y - y'\|^2 = h^2 + R^2 \sin^2 \theta$. Combining these, we obtain the quadratic

$$h^2(\kappa^2 \sin^2 \theta - 1) + h2\kappa^2 R \cos \theta \sin^2 \theta + R^2 \sin^2 \theta(\kappa^2 \cos^2 \theta - 1) < 0$$

in h . Provided θ is small enough that the second-degree coefficient $\kappa^2 \sin^2 \theta - 1$ is negative, then we will satisfy this inequality for all $h > h^*$ where h^* is the larger root of the quadratic given by

$$h^* = R \sin \theta \frac{\kappa^2 \cos \theta \sin \theta + \sqrt{\kappa^2 - 1}}{1 - \kappa^2 \sin^2 \theta}$$

To recap, provided $\theta < \arcsin \frac{1}{\kappa}$, if $h > h^*$, then

$$\begin{aligned} & h^2(\kappa^2 \sin^2 \theta - 1) + h2\kappa^2 R \cos \theta \sin^2 \theta + R^2 \sin^2 \theta(\kappa^2 \cos^2 \theta - 1) < 0 \\ \implies & \kappa^2(h + R \cos \theta)^2 \sin^2 \theta < h^2 + R^2 \sin^2 \theta \\ \implies & \kappa^2 \|y'\|^2 \sin^2 \theta < \|y - y'\|^2 \\ \implies & \frac{\|y'\|}{\|y - y'\|} \sin \theta < \frac{1}{\kappa} \\ \implies & \sin \gamma < \frac{1}{\kappa} \\ \implies & \cos \alpha_0 < \frac{1}{\kappa} \\ \implies & \cos(\alpha(\phi)) < \frac{1}{\kappa} \quad \text{for all } \phi \in [0, \theta]. \end{aligned}$$

This contradicts the mean value theorem which ensures $\cos \alpha^* \geq \frac{1}{\kappa}$. So h cannot exceed the critical value $h^* = R(\sqrt{\kappa^2 - 1} \theta + \kappa^2 \theta^2 + O(\theta^3))$. Recalling that $R = \|y\|$ by definition, we can weaken this bound by replacing R by $\text{diam}(\Omega)$ for a global result. That is, regardless of

$\|y\|$, we have the upper bound

$$h < h^{\text{gl}} = \text{diam}(\Omega) \sin \theta \frac{\kappa^2 \cos \theta \sin \theta + \sqrt{\kappa^2 - 1}}{1 - \kappa^2 \sin^2 \theta}.$$

Finally, this gives us a bound on $\|y - y'\|$,

$$\begin{aligned} \|y - y'\|^2 &= h^2 + R^2 \sin^2 \theta \\ &\leq (h^{\text{gl}})^2 + \text{diam}(\Omega)^2 \sin^2 \theta \\ &= \text{diam}(\Omega)^2 \sin^2 \theta \left(\left(\frac{\kappa^2 \cos \theta \sin \theta + \sqrt{\kappa^2 - 1}}{1 - \kappa^2 \sin^2 \theta} \right)^2 + 1 \right) \\ &= \text{diam}(\Omega)^2 (\kappa^2 \theta + 2\kappa^2 \sqrt{\kappa^2 - 1} \theta^2 + O(\theta^3)). \end{aligned}$$

To complete the proof, we note that the right hand side is an increasing function of θ so the same bound holds as well for pairs y, y' with lesser angular separation.

□

References

- [1] Katsuhiko Nishinari, Andreas Schadschneider, Debashish Chowdhury. *Stochastic Transport in Complex Systems*. Elsevier, 2011.
- [2] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237, 2008.
- [3] Sara Bernardi and Marco Scianna. An agent-based approach for modelling collective dynamics in animal groups distinguishing individual speed and orientation. *Philosophy*

- ical Transactions of the Royal Society B: Biological Sciences*, 375(1807):20190383, jul 2020.
- [4] Andrew J Bernoff and Chad M Topaz. A primer of swarm equilibria. *SIAM Journal on Applied Dynamical Systems*, 10(1):212–250, 2011.
- [5] Dapeng Bi, J. H. Lopez, J. M. Schwarz, and M. Lisa Manning. A density-independent rigidity transition in biological tissues. *Nature Physics*, 11(12):1074–1079, sep 2015.
- [6] Dapeng Bi, Xingbo Yang, M. Cristina Marchetti, and M. Lisa Manning. Motility-driven glass and jamming transitions in biological tissues. 6(2), apr 2016.
- [7] Victor J. Blue and Jeffrey L. Adler. Cellular automata microsimulation of bidirectional pedestrian flows. 1678(1):135–141, jan 1999.
- [8] Carsten Burstedde, Ansgar Kirchner, Kai Klauck, Andreas Schadschneider, and Johannes Zittartz. Cellular automaton approach to pedestrian dynamics - applications.
- [9] Iain D Couzin, Jens Krause, Nigel R Franks, and Simon A Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [10] Iain D Couzin, Jens Krause, Richard James, Graeme D Ruxton, and Nigel R Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.
- [11] F. Cucker and S. Smale. Emergent behavior in flocks. *IEEE Transactions on Automatic Control*, 52(5):852–862, May 2007.
- [12] Johann Delcourt, Nikolai W. F. Bode, and Mathieu Denoël. Collective vortex behaviors: Diversity, proximate, and ultimate causes of circular animal group movements. 91(1):1–24, mar 2016.

- [13] Maria R D’Orsogna, Yao-Li Chuang, Andrea L Bertozzi, and Lincoln S Chayes. Self-propelled particles with soft-core interactions: patterns, stability, and collapse. *Physical review letters*, 96(10):104302, 2006.
- [14] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM review*, 41(4):637–676, 1999.
- [15] Qiang Du, Max Gunzburger, and Lili Ju. Advances in studies and applications of centroidal voronoi tessellations. *Numerical Mathematics: Theory, Methods and Applications*, 3(2):119–142, 2010.
- [16] Minoru Fukui and Yoshihiro Ishibashi. Self-organized phase transitions in cellular automaton models for pedestrians. 68(8):2861–2863, aug 1999.
- [17] Allen Gersho. Asymptotically optimal block quantization. *IEEE Transactions on information theory*, 25(4):373–380, 1979.
- [18] Subir Kumar Ghosh. *Visibility Algorithms in the Plane*. Cambridge University Press, 2007.
- [19] Francesco Ginelli and Hugues Chaté. Relevance of metric-free interactions in flocking phenomena. *Physical Review Letters*, 105(16), oct 2010.
- [20] Pablo Gómez-Moureló. From individual-based models to partial differential equations. 188(1):93–111, oct 2005.
- [21] Ivan Gonzalez, Rustum Choksi, and Jean-Christophe Nave. A simple geometric method for navigating the energy landscape of centroidal voronoi tessellations. *SIAM Journal on Scientific Computing*, 43(2):A1527–A1554, 2021.
- [22] Guillaume Grégoire, Hugues Chaté, and Yuhai Tu. Moving and staying together without a leader. *Physica D: Nonlinear Phenomena*, 181(3):157–170, 2003.

- [23] Peter Manfred Gruber. A short analytic proof of fejes tóth’s theorem on sums of moments. *aequationes mathematicae*, 58(3):291–295, 1999.
- [24] Daniel Grünbaum. Translating stochastic density-dependent individual behavior with sensory constraints to an eulerian model of animal swarming. 33(2):139–161, dec 1994.
- [25] R.L. Hughes. The flow of large crowds of pedestrians. 53(4-6):367–370, oct 2000.
- [26] Roger L. Hughes. A continuum theory for the flow of pedestrians. 36(6):507–535, jul 2002.
- [27] Roger L. Hughes. THE FLOW OF HUMAN CROWDS. 35(1):169–182, jan 2003.
- [28] A. Jadbabaie, Jie Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on automatic control*, 48(6):988–1001, jun 2003.
- [29] Herbert Levine, Wouter-Jan Rappel, and Inon Cohen. Self-organization in systems of self-propelled particles. *Physical Review E*, 63(1), dec 2000.
- [30] M. Lindhe, P. Ogren, and K.H. Johansson. Flocking with obstacle avoidance: A new distributed coordination algorithm based on voronoi partitions. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE.
- [31] Daniela Morale, Vincenzo Capasso, and Karl Oelschläger. An interacting particle system modelling aggregation behavior: from individuals to populations. 50(1):49–66, jul 2004.
- [32] Sebastien Motsch and Eitan Tadmor. A new model for self-organized dynamics and its flocking behavior. *Journal of Statistical Physics*, 144(5):923–947, aug 2011.
- [33] Masakuni Muramatsu, Tunemasa Irie, and Takashi Nagatani. Jamming transition in pedestrian counter flow. 267(3-4):487–498, may 1999.

- [34] Donald Newman. The hexagon theorem. *IEEE Transactions on information theory*, 28(2):137–139, 1982.
- [35] Atsuyuki Okabe and David George Kendall. *Spatial tessallations concepts and applications of Voronoi diagrams*. John Wiley & Sons, Chichester; New York, 2000.
- [36] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21(4):25–34, aug 1987.
- [37] Chad M Topaz, Andrew J Bernoff, Sheldon Logan, and Wyatt Toolson. A model for rolling swarms of locusts. *The European Physical Journal Special Topics*, 157(1):93–109, 2008.
- [38] L.. Fejes Tóth. *Lagerungen in der Ebene auf der Kugel und im Raum*. Springer-Verlag, 1953.
- [39] Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Phys Rev Lett*, 75:1226–1229, Aug 1995.