Appendix: Emergence of collective behaviors from local Voronoi topological perception

Ivan Gonzalez, Jack Tisdell, Rustum Choksi, and Jean-Christophe Nave

Department of Mathematics and Statistics, McGill University

1 The definition of the VTP model compact set $\Omega \subset \mathbb{R}^2$ is at most 6, i.e.

1.1 Scope on the alignment term

To gain some insight on our definition of the alignment component (see §2.1.2), we perform a linearization of the term \mathbf{a}_i by: i) using the simplified transition function $g(s) = 1 - s/\pi$ for $s \in [0, \pi]$ and, ii) by making the first order approximation $\theta_{ij} = \arccos(\hat{\mathbf{u}}_i \cdot \hat{\mathbf{u}}_j) \approx \frac{\pi}{2}(1 - \hat{\mathbf{u}}_i \cdot \hat{\mathbf{u}}_j)$. For ease of notation let us denote the number of Voronoi neighbors of agent i by n_i and recall that ϕ_i is a rescaling; it follows that

$$\frac{\mathbf{a}_{i}}{\phi_{i}} \approx \frac{1}{n_{i}} \sum_{j \sim i} \left[1 - \frac{1}{\pi} \frac{\pi}{2} \left(1 - \hat{\mathbf{u}}_{i} \cdot \hat{\mathbf{u}}_{j} \right) \right] \hat{\mathbf{u}}_{j}$$

$$= \frac{1}{n_{i}} \sum_{j \sim i} \frac{1}{2} \left[\hat{\mathbf{u}}_{i} + \hat{\mathbf{u}}_{j} + \sin(\theta_{ij}) \left(\hat{\mathbf{u}}_{j} \right)^{\perp} \right]$$

$$= \underbrace{\frac{1}{2}}_{\text{inertial term}} + \underbrace{\frac{1}{2n_{i}} \sum_{j \sim i}}_{\text{traditional alignment}} \hat{\mathbf{u}}_{j} + \underbrace{\frac{1}{2n_{i}} \sum_{j \sim i}}_{\text{curling term}} \sin(\theta_{ij}) \left(\hat{\mathbf{u}}_{j} \right)^{\perp}$$

where we have used the identity $(\mathbf{A} \cdot \mathbf{C}) \mathbf{B} = (\mathbf{A} \times \mathbf{B}) \times \mathbf{C} + (\mathbf{B} \cdot \mathbf{C}) \mathbf{A}$ and written $(\hat{\mathbf{u}}_i \times \hat{\mathbf{u}}_j) \times \hat{\mathbf{u}}_j = \sin(\theta_{ij}) (\hat{\mathbf{u}}_j)^{\perp}$ with $(\hat{\mathbf{u}}_j)^{\perp}$ a unit vector orthogonal to $\hat{\mathbf{u}}_j$ determined by the right-hand rule. We thus see that our alignment component can be interpreted as the sum of three terms, the first one retaining the agent's orientation, the second one being a simple average of the neighbor's headings and the last term inducing a nonlinear behavior.

1.2 Scope on the director vectors $\{\mathbf{d}_i\}$

We show below that, on average, the magnitudes $\{\|\mathbf{d}_i\|\}_{i \in \Lambda}$ are bounded above by $1 + \frac{1}{1+\nu}$ and thus, the director vectors $\{\mathbf{d}_i\}_{i \in \Lambda}$ defined in (2) give a sensible collection of directions of motion to our model. We first recall an important Lemma about Voronoi tessellations

Lemma 1. Let n_i be the number of neighbors of agent i; then the average number of neighbors per Voronoi cell in the Voronoi tessellation of the either the entire plane \mathbb{R}^2 or of a

$$\frac{1}{n}\sum_{i\in\Lambda}n_i<6$$

Proof. i) for the plane \mathbb{R}^2 : let v and n be the total number of vertices and cells respectively in the Voronoi tessellation VT(X). Consider the planar graph G obtained from VT(X) by truncating the unbounded edges outside some sufficiently large disc, and joining their new ends (without crossings) at some new sufficiently distant vertex. (Assuming $n \ge 2$) G has exactly as many edges and cells as VT(X) (but only a single noncompact cell) and exactly one additional vertex. Let ebe the number of edges. Since every edge in G has exactly two vertices and at least three edges meet at every vertex we have $2e \ge 3v$. Next, from Euler characteristic on the plane we have (v + 1) - e + n = 2. Combining these we have that $n - 1 = e - v \ge e/3$. By noticing that the average number of neighbors per cell is 2e/n where we double count e as each edge is shared by two cells, we obtain

$$\frac{2e}{n} \le \frac{6(n-1)}{n} < 6.$$

which translates to the alternate expression in the proposition as two neighbors in the tessellation can only share one edge in .

ii) for a compact convex set $\Omega \subset \mathbb{R}^2$ with boundary: let e be the number of edges shared by two Voronoi regions and e_b be boundary edges that live on $\partial\Omega$. Let v be the total number of Voronoi vertices. That is, vertices of the \mathbb{R}^2 diagram on the interior of Ω as well all points where the \mathbb{R}^2 diagram edges meet the $\partial\Omega$. No other points of $\partial\Omega$ (e.g. corners) are taken as vertices of this graph. Since every edge has exactly two vertices and at least three edges meet at every vertex we have $2(e + e_b) \geq 3v$. Next, from Euler characteristic on the plane we have $v - (e + e_b) + n = 2$. Combining these we have that $e + e_b \leq 3(n-2)$. Finally, the average number of neighbors per cell is 2e/n where we double count e as each edge is shared by two cells, we obtain

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

Note that one may replace the convexity assumption by the much weaker assumption that $\partial\Omega$ is the union of disjoint simple curves and Ω is connected with nonempty interior using essentially the same argument except that the average number of neighbors per cell as *at most* 2e/n since the number

^{*}ivan.gonzalez@mail.mcgill.ca

[†]jack.tisdell@mail.mcgill.ca

[‡]rustum.choksi@mcgill.ca

[§]jean-christophe.nave@mcgill.ca

of neighbors for a given cell is at most the number of its non boundary edges. Similar results are obtained for the 2-torus and the 2-sphere. $\hfill\square$

Proposition 2. For all positions $X \in \Omega^n$ and velocities $U \in (\mathbb{R}^2)^n$, the average of all displacement vectors $\{\mathbf{d}_i\}_{i \in \Lambda}$ with $\#\Lambda = n$ satisfies

$$\frac{1}{n}\sum_{i\in\Lambda} \|\mathbf{d}_i(X,U)\| \le 1 + \frac{1}{1+\nu}$$

Proof. Let us start with the definitions (3, 4, 5) from which we clearly have that $\|\hat{\mathbf{h}}_i\|$, $\|\hat{\mathbf{r}}_i\|$ and $\|\mathbf{a}_i\|/\phi_i$ are all at most 1 for each $i \in \Lambda$. Moreover, for generality of the scenario, let us distinguish between agents having an active target T_i and those who do not have one (target $T_i = \emptyset$ is empty), i.e. let us write $\Lambda = \Lambda_T \cup \Lambda_{\emptyset}$ where $\Lambda_T \cap \Lambda_{\emptyset} = \emptyset$. For conciseness let $\overline{\sigma}_i := 1 - \sigma_i$ and, as before, let n_i denote the number of Voronoi neighbors of agent *i*. It ensues that, by using the binary "switch"

$$b_i := \begin{cases} 1 & \text{if } i \in \Lambda_2 \\ 0 & \text{if } i \in \Lambda_{\emptyset} \end{cases}$$

the definition of \mathbf{d}_i (e.q. (2)) generalizes to obtain

$$\begin{split} \frac{1}{n} \sum_{i \in \Lambda} & \|\mathbf{d}_i\| = \frac{1}{n} \sum_{i \in \Lambda} \left\| \frac{\sigma_i \hat{\mathbf{r}}_i + \nu \, \mathbf{a}_i + b_i (1 - \sigma_i) \hat{\mathbf{h}}_i}{\sigma_i + \nu + b_i (1 - \sigma_i)} \right\| \\ & \leq \frac{1}{N} \sum_{i \in \Lambda} \frac{\sigma_i + b_i \overline{\sigma}_i}{\sigma_i + \nu + b_i \overline{\sigma}_i} + \frac{1}{n} \sum_{i \in \Lambda} \frac{\nu \, \phi_i}{\sigma_i + \nu + b_i \overline{\sigma}_i} \\ & \leq \frac{1}{n} \sum_{i \in \Lambda} \frac{\sigma_i + b_i \overline{\sigma}_i}{\sigma_i + \nu + b_i \overline{\sigma}_i} + \frac{1}{n} \sum_{i \in \Lambda} \phi_i \\ & = \frac{1}{n} \sum_{i \in \Lambda} \frac{\sigma_i + b_i \overline{\sigma}_i}{\sigma_i + \nu + b_i \overline{\sigma}_i} + \frac{1}{6n} \sum_{i \in \Lambda} n_i \\ & \leq \frac{1}{n} \sum_{i \in \Lambda} \frac{\sigma_i + b_i \overline{\sigma}_i}{\sigma_i + \nu + b_i \overline{\sigma}_i} + 1 \\ & = 1 + \frac{1}{n} \left(\sum_{i \in \Lambda_T} + \sum_{i \in \Lambda_{\emptyset}} \right) \frac{\sigma_i + b_i \overline{\sigma}_i}{\sigma_i + \nu + b_i \overline{\sigma}_i} \end{split}$$

Where we have used the definition of the scaling $\phi_i := \frac{n_i}{6}$ and Lemma 1. Let us now examine the remaining sum by considering separately $i \in \Lambda_T$ and $i \in \Lambda_{\emptyset}$. Starting with the targeted agents for which $b_i = 1 \forall i \in \Lambda_T$ and recalling $\overline{\sigma}_i := 1 - \sigma_i$,

$$\frac{1}{n}\sum_{i\in\Lambda_T}\frac{\sigma_i+\overline{\sigma}_i}{\sigma_i+\nu+\overline{\sigma}_i}=\frac{1}{n}\sum_{i\in\Lambda_T}\frac{1}{1+\nu}=\frac{\#\Lambda_T}{n}\frac{1}{1+\nu}.$$

On the other hand, for the non targeted agents we have $b_i = 0$ for every $i \in \Lambda_{\emptyset}$ and

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

so $\frac{\sigma_i}{\sigma_i + \nu} \leq \frac{1}{1 + \nu}$ from which we recover

$$\frac{1}{n}\sum_{i\in\Lambda_{\emptyset}}\frac{\sigma_{i}+b_{i}\overline{\sigma}_{i}}{\sigma_{i}+\nu+b_{i}\overline{\sigma}_{i}}=\frac{1}{n}\sum_{i\in\Lambda_{\emptyset}}\frac{\sigma_{i}}{\sigma_{i}+\nu}\leq\frac{\#\Lambda_{\emptyset}}{n}\frac{1}{1+\nu}.$$

Combining these results and using $\#\Lambda_T + \#\Lambda_{\emptyset} = \#\Lambda = n$ we obtain the desired result

$$\frac{1}{n} \sum_{i \in \Lambda} \|\mathbf{d}_i\| \le 1 + \frac{\#\Lambda_T}{n} \frac{1}{1+\nu} + \frac{\#\Lambda_{\emptyset}}{n} \frac{1}{1+\nu} = 1 + \frac{1}{1+\nu}$$

1.3 Continuity of the planar Voronoi pressure

Lemma 3. Given distinct points $\mathbf{x}_1, \ldots, \mathbf{x}_n$ in \mathbb{R}^2 with corresponding Voronoi cells V_1, \ldots, V_n , the Voronoi pressure defined by $P = \sum_i \frac{1}{|V_i|}$ is continuous with respect to perturbations of the generators which preserve the identities (i.e., indices) of those on the convex hull.

Proof. Fix $\varepsilon > 0$. Let $\mathbf{x}'_1, \ldots, \mathbf{x}'_n$ be the perturbed sites and V'_1, \ldots, V'_n their Voronoi cells. Let P and P' be the respective pressures. Assume $\|\mathbf{x}_i - \mathbf{x}'_i\| \leq \delta$ for some positive δ to be determined for all i and that, as in the lemma statement, \mathbf{x}_i is in the convex hull of $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ if and only if \mathbf{x}'_i is in the convex hull of $\{\mathbf{x}'_1, \ldots, \mathbf{x}'_n\}$ (or equivalently, V_i is bounded if and only if V'_i is). Let $0 \le b < n$ be the number of bounded cells in either diagram. Without loss of generality, assume V_i is bounded just in case $i \leq b$. If b = 0, the result is trivial as P = 0 = P'. So assume $b \ge 1$. Let $K \subset \mathbb{R}^2$ be a closed convex set (say, a ball) which contains the 2ε fattening of $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \cup V_1 \cup \cdots \cup V_b$. By Reem's geometric stability result, for any $\varepsilon_1 > 0$, there is $0 < \delta_1 < \varepsilon$ such that if $\|\mathbf{x}_i - \mathbf{x}'_i\| < \delta_1$, then $h(V_i \cap K, V'_i \cap K) < \varepsilon_1$ for each $1 \le i \le n$ where h is the Hausdorff distance. (The condition that $\delta_1 < \varepsilon$ ensures that $\mathbf{x}'_i \in K$.) If V_i is bounded, then by definition of K, we know its closed ε -fattening $V_i + \overline{B}_{\varepsilon}$ lies in the interior of K where $\overline{B}_{\varepsilon}$ is the closed ball of radius ε . Taking $\varepsilon_1 \leq \varepsilon$, we find $V'_i \cap K \subseteq (V_i \cap K) + \overline{B}_{\varepsilon_1} \subseteq V_i + \overline{B}_{\varepsilon_1}$ so $V'_i \cap K$ is in the interior of K, hence V'_i itself is by convexity of V'_i . So in fact, $h(V_i, V'_i) = h(V_i \cap K, V'_i \cap K) < \varepsilon_1$ whenever V_i is bounded.

Then by the Steiner formula for convex sets, there is a constant C > 0 depending only on K such that $|V'_i| \leq |V_i + \overline{B}_{\varepsilon_1}| \leq |V_i| + C\varepsilon_1$. So if $\varepsilon_1 \leq \varepsilon |V_i|/C$, then

$$P' = \sum_{i \le b} \frac{1}{|V'_i|} \ge \sum_{i \le b} \frac{1}{(1+\varepsilon)|V_i|} = \frac{1}{1+\varepsilon}P.$$

Symmetrically, $|V_i| \leq |V'_i + \overline{B}_{\varepsilon_1}| \leq |V'_i| + C\varepsilon_1$ so $|V'_i| \geq |V_i| - C\varepsilon_1 \geq (1 - \varepsilon)|V_i|$ and $P' \leq \frac{1}{1 - \varepsilon}P$.

2 The plane

2.1 Comment on choice of L = 1

Since the agents' step size is roughly constant (in a statistical sense) and independent of L, changing L effectively changes the time step. Recall that L essentially controls the minimum stable swarm density so if L is much less than 1, then the average step size is much larger than the average distance between agents, that is, agents tend to step over each other. We regard this situation as nonphysical, or at least contrary to the nature of the Voronoi model, as Voronoi cells will generally fail to maintain even their approximate shape over consecutive time steps. In other words, it is simply incoherent to adopt the VTP approach to modeling a swarm while working in space/time units that result in small L. On the other hand, if L is much larger than 1, then the average step size is tiny compared to the average distance between agents and so, macroscopically, the system evolves slowly in terms of number of iterations. We do not discount this regime outright as we did the small L case but we do not study it here for the simple reason that such small change per iteration prohibits computational exploration of long term behavior. It simply takes too long to simulate on a home computer to investigate in the way we intend. Thus we consider only L = 1 in the rest of this section as this value balances nicely the overall evolution speed of the system with the severity of geometric changes in the Voronoi diagram. (This is in sharp contrast to the next section, wherein we consider a bounded domain and different values of L, appropriately non-dimensionalized, play a central role.)

We remark that L is not the full story, but one piece of data about the full *repulsive falloff function*. Given a repulsive falloff function supported on [0, L], one could imagine a different function which decays from 1 to some $0 < \varepsilon \ll 1$ over [0, L] but which has a long (or even infinite) tail. Such a modification has minimal effect on the dynamics predicted by the model, and so the length of the support fails to capture the relevant information about the two falloff functions. Nonetheless, for a fixed functional shape with compact support on $[0, \infty)$, the length of the support L is a convenient parameter.

3 The bidirectional hallway

We present below the details originally omitted in the presentation of §4.

3.1 The entry process of agents at the sources

Having fixed the source length scale L_s , we describe below the stochastic process modeling the agents' entrance into the hallway Ω ; for simplicity we focus on the source $S_r \subset \partial \Omega$ which is the left-end of the corridor used by agents $i \in \Lambda_r$ that are moving to the right. The process is similar for the right-end source S_l .

At any fixed time t, given X(t), we insert new agents iteratively before computing the next state X(t+1) according to our governing equations (2). This iterative process (which we index by k in this discussion) occurs between each timestep and is not to be confused with time evolution (indexed by t). We follow these guidelines:

- i) a single agent is inserted on S_r at each iteration k.
- ii) the position on S_r of the k^{th} new agent (placed during the k^{th} iteration) depends on the overall position state X(t) of all agents currently in the hallway at time t and also on the position of the agents 1, 2, ..., k-1 previously placed in the earlier iterations.
- iii) the iterative placement is terminated whenever the preferred personal distance L_s can no longer be guaranteed on S_r for the agent k + 1.
- iv) all new agents are given a random unitary orientation $\{\hat{\mathbf{u}}_i\}$ whose angle with the horizontal is independently sampled from $(-\frac{\pi}{2}, \frac{\pi}{2})$.

More specifically, let some $\mathbf{y} \in S_r$ be the center of a disk of radius R that does not contain any agent $i \in \Lambda(t)$ currently in the hallway nor any of the k-1 agents recently introduced on S_r . If $R < L_s$ then S_r is too crowded near \mathbf{y} and there is a zero probability that the k^{th} agent enters through that location. On the other hand, if $R \geq L_s$ then there is a nonzero probability that the k^{th} new agent enters the hallway through **y** and this probability increases with R. Finally, as the number of newly placed agents increases, the source S_r will be saturated when no new agent can enter while having an empty radius $R \geq L_s$ around it. The pseudo-code from Algorithm 1 provides the specifics.

Algorithm 1 Stochastic entrance of agents on S_r

Input: 1) fixed L_s ; 2) current positions X(t) of agents at time t; 3) a nondecreasing transition function ξ : $\mathbb{R}^+ \to [0, 1]$ (e.g. $\xi(s) := 1 - \sigma(s)$, see §2.1.4)

set k = 0

set $Y = \emptyset$

while true do

for all $\mathbf{y} \in \mathcal{S}_r$ let $R_{\mathbf{y}} := \operatorname{dist}(\mathbf{y}; X \cup Y)$ and construct

$$f_k(\mathbf{y}) := \begin{cases} \xi \left(R_{\mathbf{y}} / L_s - 1 \right) & \text{if } R_{\mathbf{y}} \ge L_s \\ 0 & \text{if } R_{\mathbf{y}} < L_s \end{cases}$$

if $f_k(\mathbf{y}) \equiv 0$ then

break (meaning S_r is saturated)

end if

- (a) normalize $f_k(\mathbf{y})$ so that $\int_{\mathcal{S}_r} f_k(\mathbf{y}) d\mathbf{y} = 1$
- (b) generate the position \mathbf{y}_k of the k^{th} new agent on \mathcal{S}_r by a random sampling using the PDF $f_k(y)$

(c) updates: $Y \leftarrow Y \cup \{\mathbf{y}_k\}$ and $k \leftarrow k+1$

end while

Output: Y are the collected positions on S_r of the k newly inserted agents to be added to X(t).

<u>NOTE</u>: it is possible that k = 0 and $Y = \emptyset$, indicating that no new agent will enter at time t due to S_r being overcrowded from the beginning.

Note that this construction can trivially use two distinct values of L_s to each of Λ_r and Λ_l in order to have different inflows for each subgroup of agents; in our work however, we consider that the entering conditions and L_s are the same for both populations. At last, we point out for completeness, that slight variants of this entry process were tested but did not produce noticeable differences in the obtained regimes. Specifically, we tested variants of the PDF definition for $f_k(\mathbf{y})$ where a dependency in the orientation of agents near $\mathbf{y} \in S$ was added (on top of the dependency in position of nearby agents). Forcing the orientations $\{\hat{\mathbf{u}}_i\}$ of entering agents to be orthogonal to the sources S was also tested.

3.2 Queuing

We describe here our ad hoc construction of the queuing structures $\Xi_{r,l}(t)$ over which the queuing observable $Q(t) := \frac{1}{2} (Q_r(t) + Q_l(t))$ was calculated in this work. Once more, let us focus on the agents $i \in \Lambda_r(t)$ crossing the hallway from left to right; the graph Ξ_l for the subpopulation $\Lambda_l(t)$ is constructed analogously. We denote by DT(X) the Delaunay triangulation (Voronoi dual) generated by X. To systematically satisfy the four queuing postulates from $\S4.1$ at any given time t we follow these steps:

- 1. Let $\mathcal{D}_r := \mathrm{DT}(X)|_{\Lambda_r}$ be the restriction of $\mathrm{DT}(X)$ to the agents X_r , i.e. we remove from $\mathrm{DT}(X)$ all vertices from Λ_l and any Delaunay edges having an endpoint in X_l . Note that, while \mathcal{D}_r is an undirected graph, it is in general highly disconnected.
- 2. Let us momentarily construct $\overrightarrow{\mathcal{D}}_r$, a directed version of \mathcal{D}_r , where every edge $\overrightarrow{e}_k = (\overrightarrow{\mathbf{x}_{i_k}, \mathbf{x}_{j_k}}) \in \text{edge}(\overrightarrow{\mathcal{D}}_r)$ points "along" the homing $\hat{\mathbf{h}} = (1, 0)^{\top}$ common to all agents in Λ_r ; i.e. we orient the edges such that $\overrightarrow{e}_k \cdot \hat{\mathbf{h}} \geq$ $0 \forall k$. Let $\{\hat{e}_k\}$ denote the unitary vectors associated.
- 3. We now compute nonnegative edge weights $\{w_k\}$ over $\overrightarrow{\mathcal{D}}_r$ via the formula

$$w_k := \frac{1}{2} (\theta_{k,i_k} + \theta_{k,j_k})$$

= $\frac{1}{2} (\arccos(\hat{e}_k \cdot \hat{u}_{i_k}) + \arccos(\hat{e}_k \cdot \hat{u}_{j_k})$

)

Under this definition, $\{w_k\}$ considers both **relative po**sition of the endpoint agents \mathbf{x}_{i_k} , \mathbf{x}_{j_k} as well as their **relative orientations** $\hat{\mathbf{u}}_{i_k}$, $\hat{\mathbf{u}}_{j_k}$. Specifically, $0 \leq w_k \leq \pi$ is minimal whenever the positions of both agents "align towards" their common target and their respective orientations coincide with the homing $\hat{\mathbf{h}}$. Conversely, w_k is maximal when agents "align away" from their target, i.e. whenever their position is aligned towards the target but their orientation is pointing in the opposite direction. In other words, $\{w_k\}$ uses basic geometrical features of X_r and \hat{U}_r to locally quantify alignment quality over

every neighboring pair of agents in Λ_r .

- 4. Now that all the pertinent local geometry has been encoded into $\{w_k\}$, we consider again the undirected \mathcal{D}_r now endowed with these weights. We obtain Ξ_r by running a minimum spanning forest algorithm on \mathcal{D}_r (e.g. Dijkstra's algorithm on every connected component of \mathcal{D}_r) so that the recovered Ξ_r is an acyclic graph. The intuition behind this is to span all agents in Λ_r while removing all edge-loops in \mathcal{D}_r and retaining "horizontal" edges with "better oriented" endpoint agents in the process.
- 5. At this stage, one could subjectively argue that each connected component of Ξ_r can be interpreted as a graph capturing a certain amount of oriented consensus and spatial cohesion. However, to us, the potentially large degree of vertices of Ξ_r (ramifications) do not constitute what can intuitively be called lanes. Thus, to showcase agents "lining up" behind one another, we refine Ξ_r by iteratively cutting (removing) edges of this forest in decreasing order of their associated weights; the removal process stops once $\deg_{\Xi_r}(\mathbf{x}_i) \leq 2$ for every $i \in \Lambda_r$. In other words, this refinement creates more non-ramified connected components to the detriment of breaking longer ramified chains; all while retaining edges with the smallest possible weights.
- 6. As an added property, we remove from Ξ_r any connected component having a single edge (two agents). This is because we consider that a lane needs to consist of at least three agents.

In summary, steps 1)-5) above guarantee that the four queuing postulates i)-iv) are met by our queuing structure Ξ_r and thus make it sensible to compute the queuing observable Q_r on it (see §4.1). For thoroughness, these steps are synthesized in Algorithm 2.

Finally, we remark that, should the reader's own interpretation of queuing allows for ramified queuing structures (i.e. postulate ii) needs not hold); one can introduce a tolerance value w_{tol} such that edge e_k is removed in step 5) above only if $w_k > w_{\text{tol}}$. Thus obtaining a graph $\Xi_r(X_r, \hat{U}_r; w_{\text{tol}})$ parametrized by the user-defined value w_{tol} . Another alternative is to solve the discrete optimization problem

$$\underset{w_{\text{tol}}}{\arg\min} Q_r(\Xi_r(X_r, \hat{U}_r; w_{\text{tol}}))$$

Which will generally yield a non-trivial graph since the objective function is not monotonically decreasing in w_{tol} . Such alternatives are not explored in this article so as to avoid introducing a supplementary degree of freedom and maintain simplicity; their use and more in-depth analysis is left for future work.

Algorithm 2 Our construction of the queuing structure Ξ_r

Input: 1) index set of agents Λ_r ; 2) positions $X_r := \{\mathbf{x}_i\}_{\in \Lambda_r}$; 3) unitary orientations $\hat{U}_r := \{\hat{\mathbf{u}}_i\}_{\in \Lambda_r}$; 4) Delaunay triangulation DT(X).

(1) restrict DT(X) to X_r to create \mathcal{D}_r

(2) create oriented version $\vec{\mathcal{D}}_r$ such that $\vec{e}_k \cdot \hat{\mathbf{h}} \ge 0$ for all $\vec{e}_k \in \text{edge}(\vec{\mathcal{D}}_r)$

(3) compute weights $\{w_k\}$ associated to $\{\overrightarrow{e}_k\}$

$$w_k := \frac{1}{2} (\arccos\left(\hat{e}_k \cdot \hat{\mathbf{u}}_{i_k}\right) + \arccos\left(\hat{e}_k \cdot \hat{\mathbf{u}}_{j_k}\right))$$

(4) endow the undirected graph \mathcal{D}_r with $\{w_k\}$ and obtain a minimum spanning forest Ξ_r , i.e. every connected component of Ξ_r is a minimum spanning tree. (5) cut edges $e_k \in \text{edge}(\Xi_r)$ in descending order of their weights w_k until $\text{deg}_{\Xi_r}(\mathbf{x}_i) \leq 2$ for all $i \in \Lambda_r$.

(6) optional: cut all single-edge connected components

Output: an undirected weighted forest Ξ_r satisfying all four queuing postulates and whose connected components we identify as lanes formed by the subpopulation Λ_r .

3.3 Phase diagrams for Model II on the bidirectional hallway

Figures 3.1 & 3.2 depict the $(L; L_s)$ phase portraits using Model II for $\nu = 1$ and $\nu = 2$ respectively; these are to be compared with Figures 5 and 6. For the weak alignment case, i.e. $\nu = 1$, we observe an uncanny resemblance between Models I and II; this includes the presence of a clear sharp bifurcation line separating the two extreme cases of full occlusion



Figure 3.1: Phase diagram (L, L_s) for **Model II** on the bidirectional hallway under **weak** alignment $\nu = 1$ (with resolution of 65 × 65 points). Much alike Figure 5 obtained for Model I, a bifurcation line shows a clear phase transition between the two extreme cases of full occlusion and sustained percolation. We remark that, as opposed to using the forward Voronoi area (i.e. Model I), using the forward available length Model II allows for slightly better percolation.

and sustained migration. For the case $\nu = 2$, the diagram of Model II is similar to the one of Model I up to a shift and scaling; qualitatively, all the exhibited behaviors remain very similar. In conclusion, when Ω is a bounded domain, there is no significant difference between agents using the forward area of their personal space (Model I) versus them using their personal forward length (Model II).

4 Compact Domains without Boundary

We have adapted VTP (Model I) to two compact domains Ω without boundary, the square torus $\mathbb{R}^2/l\mathbb{Z}^2$ of primary domain $[0, l)^2$ and the 2-dimensional sphere. Here it is natural to define a dimensionless parameter $\mu > 0$ capturing the relative length scales based upon n, L, and $|\Omega|$. Roughly speaking μ should represents the ratio of the repulsive length scale to the average inter-agent distance. We take

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

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

4.1 Phase diagram for homing-free systems

The μ - ν phase diagram for the scenario devoid of targets (untargeted case) is sketched in Figure 4.1 In the supplementary material we include simulations for these regimes as well as a brief description of their labels. We include as well simulations of all these regimes for the untargeted sphere (wherein a similar phase diagram is observed).

As in the case of the infinite plane $(\S3)$, the dynamics are simplified for lack of boundary interactions. We remark that in the absence of unbounded Voronoi cells, the dynamics of models I and II are not qualitatively different on compact domains and a lot of structure is preserved from one to the other. The governing equations are further simplified in case that no targets are assigned, we call such systems homing-free. Roughly speaking, the (μ, ν) phase space for homing-free systems on the torus and sphere are quite similar. Changes in μ and ν correlate tightly with changes in clustering energy and polarization respectively. Here clustering energy is as described in the previous section. On the torus, polarization is as defined in the previous section with respect to the (arbitrary) coordinates inherited from $[0, l)^2$. For the sphere, by "polarization" we mean the angular momentum of the ensemble with respect to the (unit) sphere's center viewing the position an orientation vectors as vectors in \mathbb{R}^3 in the natural way. In this section, we shall use the term "polarization" to talk about the sphere and torus at once. We find the following clearly distinct extremes.

- **Gaseous.** (Small μ , small ν) When ν is small, orientations are spatiotemporally uncorrelated achieving polarization near zero. With small μ , agents only interact repulsively at distances small compared to the domain scale and so positions as well are essentially random and modest clustering energy persists (although no consistent clusters propagate noticeably through the agent medium).
- **Solid.** (Large μ , small ν) As above, since ν is small, orientations remain uncorrelated and polarization near zero. But with large μ , mutual repulsion at lengths comparable to the domain scale force the system toward a uniformly spaced hexagonal crystalline structure (nec-



Figure 3.2: Phase diagram (L, L_s) for **Model II** on the bidirectional hallway under **strong** alignment $\nu = 2$ (with resolution of 65×65 points).

Similar to Figure 6 obtained for Model I, the transition between the two extreme regions of full occlusion and sustained percolation is significantly blurry. This lack of sharp transition is again attributed to the sheer effects and ensuing nonlinearities produced by larger ν .



Figure 4.1: Phase diagram sketch with Model I for the torus with no targets. The dimensionless parameter μ is given by $\mu = L(n/|\Omega|)^{1/2} = \sqrt{nL/l}$. The dashed lines are merely conceptual delineations, not sharp bifurcation loci. Their rough shape is based on coarse probing of phase space with simulations. While the precise features of the phase diagram of course also depend on the particular choices of transition functions σ and g, the qualitative structure remains unchanged. In the digital version of this document, click the labels in the diagram to view corresponding simulations for n = 300. (For more, explore the site.)

essarily with defects due to the topological constraints) and extremely low clustering energy (near 1) is maintained. Furthermore, since the individual orientations are essentially random, the overall drift heading is near zero.

- **Flow.** (Small μ , large ν) With large ν , the system quickly attains high polarization with near consensus in orientation. Again, for small μ , only modest clustering energy persists with no clear structure in the agents' positions.
- **Drift.** (Large μ , large ν) Like the solid regime, large μ quickly drives the system to a near crystal structure but now also, neighbors tend to align their orientations and overall orientation consensus is attained (albeit slightly noisier than the flow regime due to significant nearest-neighbor interactions), thus achieving high polarization and an overall drift (with much less individual variation than the flow regime.

The transitional regimes are harder to characterize. For all values of μ , increasing through intermediate ν causes a gradual increase in polarization with no apparent bifurcations. Fixing ν and increasing through intermediate μ also does not show sharp bifurcations but at intermediate values of μ and moderate or large ν , one observes one or several (depending on the number of agents) persistent clusters which propagate backward (against the drift heading) through the agent medium, akin to what has been observes in traffic flow. Indeed, for sufficiently few agents (a few hundred), at intermediate μ and ν , one finds a single large patch of more densely packed agents with less correlated orientation propagating backward against a comparably large patch of sparser (hence faster) agents with near consensus in orientation.

4.2 Point targets

We further include simulations on the torus and the sphere for one, two and three point targets (wherein all agents always seek their nearest target point). The emergent dynamics are complicated, even for such simple target configurations. Nonetheless, we remark qualitatively.

In all cases, the behavior for very strong alignment is essentially similar to the homing free case (Section §4.1), achieving near orientation consensus with only slight deviation close to the target points. Compare this with the single-point target behavior in the (non-compact) planar case of Section §3, where the target has manifest influence for extremely strong alignment. This can be understood as a size effect of the compact domains. Whereas, in the strong alignment singlepoint target planar case, ensembles of agents may adopt wide orbits about the target point because (i) agents may always in principle spread out enough that homing dominates the homing-repulsion effect, giving each agent's steps a radial component with respect to the target and (ii) there exist wide enough orbits in the plane that this radial component feeble against the magnitude $\approx \nu$ alignment component—is sufficient. Clearly these on not both possible on compact domains, any minimum average distance between agents imposes a minimum path length for a size n ensemble, and sufficiently long paths in compact domains will interfere with themselves (to within the ensemble density).

For sufficiently weak alignment with a single point target, the behavior is much like in the planar case. Although this is what one would expect, it is not entirely trivial. Namely, alignment, as we have formulated it, exerts direct influence on neighboring agents and on these domains, a pinwheel or ring system has a rather different Delaunay topology than the analogous planar systems, with agents near opposite edges of the ensemble will neighbor each other.

For two and three point targets, we encourage readers to explore the simulations and we comment only on the following curiosity. In the torus, we (mostly arbitrarily) considered two target points placed bisecting a minimal (closed) geodesic. In most of the parameter space (again for sufficiently weak alignment), the system robustly organizes into a "cog" behavior, with counter-rotating ensembles surrounding the two target points. The analogous setup in the plane yields like behavior for appropriate parameter values. However, in a certain parameter range, the torus exhibited a behavior, which we call the "anti-cog", whereby the system would settle into two co-rotating ensembles about the target points, with some continual agent exchange between them. This anti-cog behavior could not be reliably reproduced on the plane. On compact domains, emergence of the anti-cog behavior was sensitive to initial conditions in the sense that, in this parameter range, for uniformly sampled position/orientation initial conditions, the system would either settle into the cog or anti-cog behavior. More specifically, we found two sub regimes in the (μ, ν) phase plane, one in which the cog behavior was by far dominant amongst random initializations and one in which the anti-cog emerged for a substantial fraction of initializations. Whether the emergence of the anti-cog behavior is a size effect or a topological one or both is unclear.