Distributed robotic networks: rendezvous, connectivity, and deployment

Sonia Martínez and Jorge Cortés

Mechanical and Aerospace Engineering
University of California, San Diego
{cortes,soniamd}@ucsd.edu

28th Benelux Meeting on Systems and Control
Spa, Belgium, March 17, 2009

Acknowledgments: Francesco Bullo
Anurag Ganguli

What we have seen in the previous lecture

Cooperative robotic network model

- proximity graphs
- control and communication law, task, execution
- time, space, and communication complexity
- analysis agree and pursue algorithm

Complexity analysis is **challenging** even in 1 dimension! Blend of math

- geometric structures
- distributed algorithms
- stability analysis
- linear iterations

What we will see in this lecture

Basic motion coordination tasks:

get together at a point, stay connected, deploy over a region

Design coordination algorithms that achieve these tasks and analyze their correctness and time complexity

Expand set of math tools: invariance principles for non-deterministic systems, geometric optimization, non-smooth stability analysis

Robustness against link failures, agents' arrivals and departures, delays, asynchronism

Image credits: jupiterimages and Animal Behavior

Outline

1. Rendezvous and connectivity maintenance
   - The rendezvous objective
   - Maintaining connectivity
   - Circumcenter algorithms
   - Correctness analysis via nondeterministic systems
2. Deployment
   - Expected-value deployment
   - Geometric-center laws
   - Disk-covering and sphere-packing deployment
   - Geometric-center laws
3. Conclusions

Rendezvous objective

Objective:
achieve multi-robot rendezvous; i.e. arrive at the same location of space, while maintaining connectivity

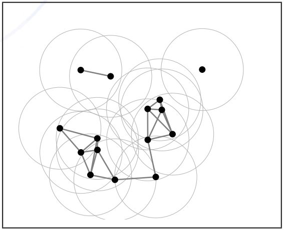

r-disk connectivity

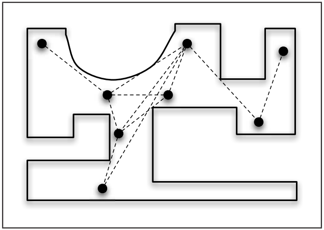

visibility connectivity

We have to be careful...

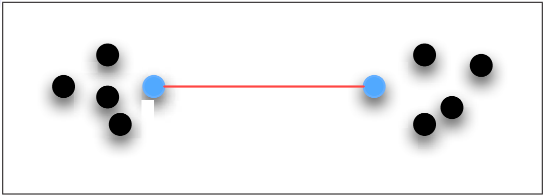

Blindly “getting closer” to neighboring agents might break overall connectivity

Network definition and rendezvous tasks

The objective is applicable for general robotic networks

$\mathcal{S}_{\text{disk}}$, $\mathcal{S}_{\text{LD}}$ and $\mathcal{S}_{\infty\text{-disk}}$,

and the relative-sensing networks $\mathcal{S}^{\text{rs}}_{\text{disk}}$ and $\mathcal{S}^{\text{rs}}_{\text{vis-disk}}$

We adopt the discrete-time motion model

$$p^{[i]}(\ell+1) = p^{[i]}(\ell) + u^{[i]}(\ell), \quad i \in \{1, \dots, n\}$$

Also for the relative-sensing networks

$$p^{[i]}_{\text{fixed}}(\ell+1) = p^{[i]}_{\text{fixed}}(\ell) + R^{[i]}_{\text{fixed}} u^{[i]}_i(\ell), \quad i \in \{1, \dots, n\}$$

The rendezvous task via aggregate objective functions

Coordination task formulated as function minimization

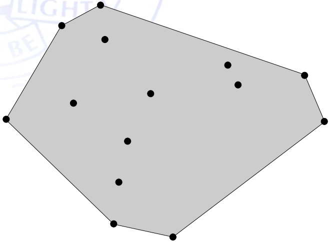

Diameter convex hull

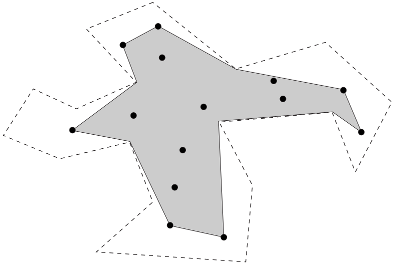

Perimeter relative convex hull

The rendezvous task formally

Let $\mathcal{S} = (\{1, \dots, n\}, \mathcal{R}, E_{\mathrm{cmm}})$ be a uniform robotic network
The **(exact) rendezvous task** $\mathcal{T}_{\mathrm{rendezvous}} \colon X^n \to \{\texttt{true}, \texttt{false}\}$ for $\mathcal{S}$ is

$$\mathcal{T}_{\mathrm{rendezvous}}(x^{[1]}, \dots, x^{[n]}) = \begin{cases} \texttt{true}, & \text{if } x^{[i]} = x^{[j]}, \text{ for all } (i,j) \in E_{\mathrm{cmm}}(x^{[1]}, \dots, x^{[n]}), \\ \texttt{false}, & \text{otherwise} \end{cases}$$

For $\epsilon \in \mathbb{R}_{>0}$, the **$\epsilon$-rendezvous task** $\mathcal{T}_{\epsilon\text{-rendezvous}} \colon (\mathbb{R}^d)^n \to \{\texttt{true}, \texttt{false}\}$ is

$$\mathcal{T}_{\epsilon\text{-rendezvous}}(P) = \texttt{true}$$
$$\iff \|p^{[i]} - \mathsf{avrg}\left(\left\{p^{[j]} \mid (i,j) \in E_{\mathrm{cmm}}(P)\right\}\right)\|_2 < \epsilon, \quad i \in \{1, \dots, n\}$$

Outline

1. Rendezvous and connectivity maintenance
   - The rendezvous objective
   - Maintaining connectivity
   - Circumcenter algorithms
   - Correctness analysis via nondeterministic systems

2. Deployment
   - Expected-value deployment
   - Geometric-center laws
   - Disk-covering and sphere-packing deployment
   - Geometric-center laws

3. Conclusions

Constraint sets for connectivity

Design constraint sets with key properties

- Constraints are flexible enough so that network does not get stuck
- Constraints change continuously with agents' position

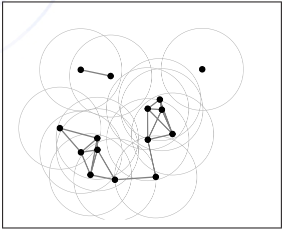

r-disk connectivity

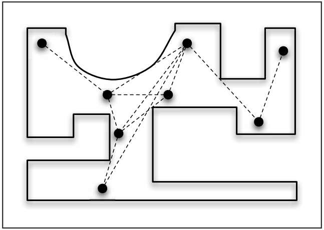

visibility connectivity

Enforcing range-limited links – pairwise

Pairwise connectivity maintenance problem:

Given two neighbors in $\mathcal{G}_{\mathrm{disk}}(r)$, find a rich set of control inputs for both agents with the property that, after moving, both agents are again within distance r

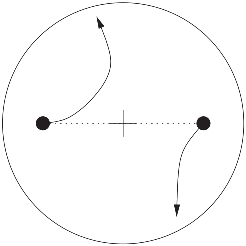

If $\|p^{[i]}(\ell) - p^{[j]}(\ell)\| \leq r$, and remain in connectivity set,
then $\|p^{[i]}(\ell+1) - p^{[j]}(\ell+1)\| \leq r$

Enforcing range-limited links – w/ all neighbors

Definition (Connectivity constraint set)

Consider a group of agents at positions $P = \{p^{[1]}, \dots, p^{[n]}\} \subset \mathbb{R}^d$. The *connectivity constraint set* of agent i with respect to P is

$$\mathcal{X}_{\text{disk}}(p^{[i]}, P) = \bigcap \left\{ \mathcal{X}_{\text{disk}}(p^{[i]}, q) \mid q \in P \setminus \{p^{[i]}\} \text{ s.t. } \|q - p^{[i]}\|_2 \leq r \right\}$$

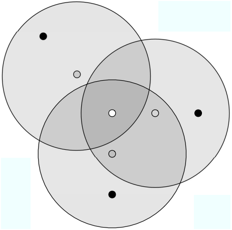

Same procedure over sparser graphs means fewer constraints: $\mathcal{G}_{\text{LD}}(r)$ has same connected components as $\mathcal{G}_{\text{disk}}(r)$ and is spatially distributed over $\mathcal{G}_{\text{disk}}(r)$

Enforcing range-limited line-of-sight links – pairwise

For $Q_\delta = \{q \in Q \mid \mathsf{dist}(q, \partial Q) \geq \delta\}$ δ-contraction of compact nonconvex $Q \subset \mathbb{R}^2$

Pairwise connectivity maintenance problem:

Given two neighbors in $\mathcal{G}_{\text{vis-disk},Q_\delta}$, find a rich set of control inputs for both agents with the property that, after moving, both agents are again within distance r and visible to each other in Q_δ

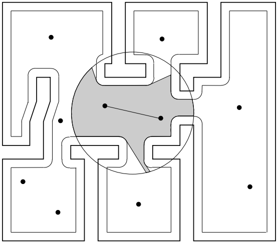

visibility region of agent i

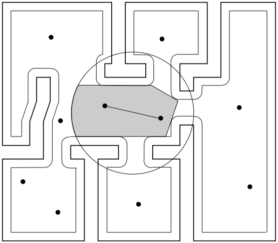

visibility pairwise constraint set

Enforcing range-limited line-of-sight links – w/ all neighbors

Definition (Line-of-sight connectivity constraint set)

Consider a group of agents at positions $P = \{p^{[1]}, \dots, p^{[n]}\}$ in a nonconvex allowable environment Q_δ. The **line-of-sight connectivity constraint sets** of agent i with respect to P is

$$\mathcal{X}_{\text{vis-disk}}(p^{[i]}, P; Q_\delta) = \bigcap \left\{ \mathcal{X}_{\text{vis-disk}}(p^{[i]}, q; Q_\delta) \mid q \in P \setminus \{p^{[i]}\} \right\}$$

Fewer constraints can be generated via sparser graphs with the same connected components and spatially distributed over

Outline

1. Rendezvous and connectivity maintenance
   - The rendezvous objective
   - Maintaining connectivity
   - Circumcenter algorithms
   - Correctness analysis via nondeterministic systems
2. Deployment
   - Expected-value deployment
   - Geometric-center laws
   - Disk-covering and sphere-packing deployment
   - Geometric-center laws
3. Conclusions

Circumcenter control and communication law

For $X = \mathbb{R}^d$, $X = \mathbb{S}^d$ or $X = \mathbb{R}^{d_1} \times \mathbb{S}^{d_2}$, $d = d_1 + d_2$, **circumcenter** $\mathsf{CC}(W)$ of a bounded set $W \subset X$ is center of closed ball of minimum radius that contains W

Circumradius $\mathsf{CR}(W)$ is radius of this ball

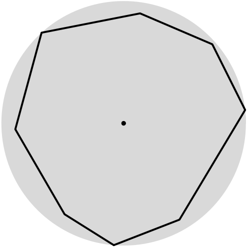

[Informal description:]

> At each communication round each agent performs the following tasks: (i) it transmits its position and receives its neighbors' positions; (ii) it computes the circumcenter of the point set comprised of its neighbors and of itself. Between communication rounds, each robot moves toward this circumcenter point while maintaining connectivity with its neighbors using appropriate connectivity constraint sets.

Illustration of the algorithm execution

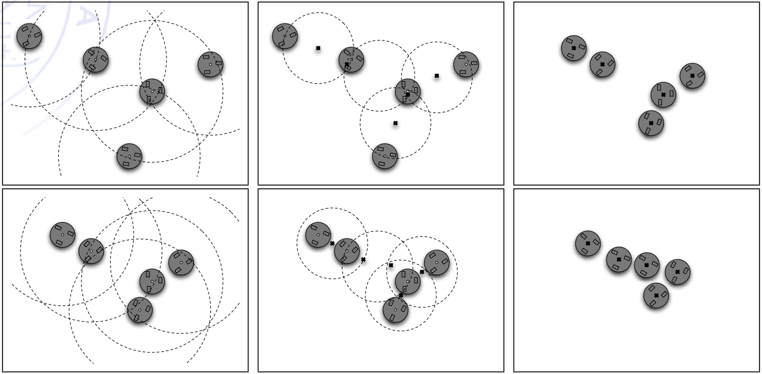

Circumcenter control and communication law

Formal algorithm description

`Robotic Network:` $\mathcal{S}_{\text{disk}}$ with a discrete-time motion model,
with absolute sensing of own position, and
with communication range r, in $\mathbb{R}^d$

`Distributed Algorithm:` circumcenter
`Alphabet:` $L = \mathbb{R}^d \cup \{\texttt{null}\}$

`function` $\text{msg}(p, i)$

1: **return** p

`function` $\text{ctrl}(p, y)$

1: $p_{\text{goal}} := \mathsf{CC}(\{p\} \cup \{p_{\text{rcvd}} \mid \text{for all non-null } p_{\text{rcvd}} \in y\})$
2: $\mathcal{X} := \mathcal{X}_{\text{disk}}(p, \{p_{\text{rcvd}} \mid \text{for all non-null } p_{\text{rcvd}} \in y\})$
3: **return** $\mathsf{fti}(p, p_{\text{goal}}, \mathcal{X}) - p$

Simulations

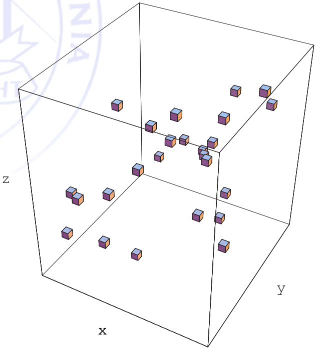

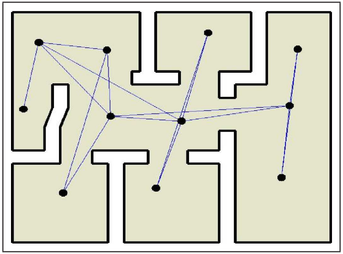

Outline

1. Rendezvous and connectivity maintenance
   - The rendezvous objective
   - Maintaining connectivity
   - Circumcenter algorithms
   - Correctness analysis via nondeterministic systems
2. Deployment
   - Expected-value deployment
   - Geometric-center laws
   - Disk-covering and sphere-packing deployment
   - Geometric-center laws
3. Conclusions

Some bad news...

Circumcenter algorithms are nonlinear discrete-time dynamical systems

$$x_{\ell+1} = f(x_\ell)$$

To analyze convergence, we need at least f continuous – to use classic Lyapunov/LaSalle results

But circumcenter algorithms are discontinuous because of changes in interaction topology

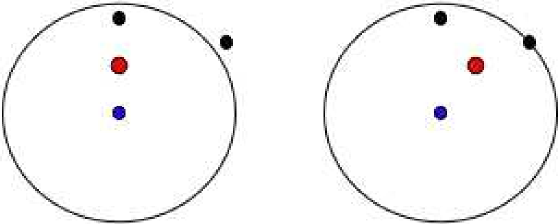

Alternative idea

Fixed undirected graph G, define **fixed-topology circumcenter algorithm**

$$f_G : (\mathbb{R}^d)^n \to (\mathbb{R}^d)^n, \quad f_{G,i}(p_1, \dots, p_n) = \mathsf{fti}(p, p_{\mathrm{goal}}, \mathcal{X}) - p$$

Now, there are no topological changes in f_G, hence f_G is continuous

Define set-valued map $T_{\mathcal{CC}} : (\mathbb{R}^d)^n \to \mathcal{P}((\mathbb{R}^d)^n)$

$$T_{\mathcal{CC}}(p_1, \dots, p_n) = \{f_G(p_1, \dots, p_n) \mid G \text{ connected}\}$$

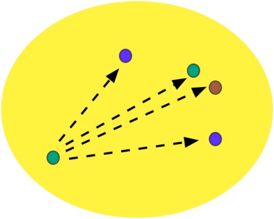

Non-deterministic dynamical systems

Given $T : X \to \mathcal{P}(X)$, a **trajectory** of T is sequence $\{x_m\}_{m\in\mathbb{N}_0} \subset X$ such that

$$x_{m+1} \in T(x_m), \quad m \in \mathbb{N}_0$$

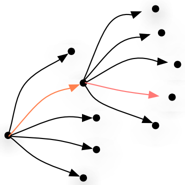

T is **closed** at x if $x_m \to x$, $y_m \to y$ with $y_m \in T(x_m)$ imply $y \in T(x)$
Every continuous map $T\colon \mathbb{R}^d \to \mathbb{R}^d$ is closed on $\mathbb{R}^d$

A set C is

- **weakly positively invariant** if, for any $p_0 \in C$, there exists $p \in T(p_0)$ such that $p \in C$
- **strongly positively invariant** if, for any $p_0 \in C$, all $p \in T(p_0)$ verifies $p \in C$

A point p_0 is a *fixed point of* T if $p_0 \in T(p_0)$

LaSalle Invariance Principle – set-valued maps

$V : X \to \mathbb{R}$ is non-increasing along T on $S \subset X$ if

$$V(x') \leq V(x) \text{ for all } x' \in T(x) \text{ and all } x \in S$$

Theorem (LaSalle Invariance Principle)

For S compact and strongly invariant with V continuous and non-increasing along closed T on S

Any trajectory starting in S converges to largest weakly invariant set contained in $\{x \in S \mid \exists x' \in T(x) \text{ with } V(x') = V(x)\}$

Correctness

$T_{\mathcal{CC}}$ is closed and diameter is non-increasing

Recall set-valued map $T_{\mathcal{CC}} : (\mathbb{R}^d)^n \to \mathcal{P}((\mathbb{R}^d)^n)$

$$T_{\mathcal{CC}}(p_1, \dots, p_n) = \{f_{\mathcal{G}}(p_1, \dots, p_n) \mid \mathcal{G} \text{ connected}\}$$

$T_{\mathcal{CC}}$ is **closed**: finite combination of individual continuous maps

Define

$$V_{\text{diam}}(P) = \mathsf{diam}(\mathsf{co}(P)) = \mathsf{max}\,\{\|p_i - p_j\| \mid i, j \in \{1, \dots, n\}\}$$

$$\mathsf{diag}((\mathbb{R}^d)^n) = \{(p, \dots, p) \in (\mathbb{R}^d)^n \mid p \in \mathbb{R}^d\}$$

Lemma

The function $V_{\text{diam}} = \mathsf{diam} \circ \mathsf{co} \colon (\mathbb{R}^d)^n \to \overline{\mathbb{R}}_+$ *verifies:*

1. V_{diam} *is continuous and invariant under permutations;*
2. $V_{\text{diam}}(P) = 0$ *if and only if* $P \in \mathsf{diag}((\mathbb{R}^d)^n)$*;*
3. V_{diam} *is non-increasing along* $T_{\mathcal{CC}}$

Correctness via LaSalle Invariance Principle

To recap

1. $T_{\mathcal{CC}}$ is closed
2. $V = \mathsf{diam}$ is non-increasing along $T_{\mathcal{CC}}$
3. Evolution starting from P_0 is contained in $\mathsf{co}(P_0)$ (compact and strongly invariant)

Application of **LaSalle Invariance Principle:** trajectories starting at P_0 converge to M, largest weakly positively invariant set contained in

$$\{P \in \mathsf{co}(P_0) \mid \exists P' \in T_{\mathcal{CC}}(P) \text{ such that } \mathsf{diam}(P') = \mathsf{diam}(P)\}$$

Have to **identify** M! In fact, $M = \mathsf{diag}((\mathbb{R}^d)^n) \cap \mathsf{co}(P_0)$

Convergence to a point can be concluded with a little bit of extra work

Correctness

Theorem (Correctness of the circumcenter laws)

For $d \in \mathbb{N}$, $r \in \mathbb{R}_{>0}$ and $\epsilon \in \mathbb{R}_{>0}$, the following statements hold:

1. *on $\mathcal{S}_{\text{disk}}$, the law $\mathcal{CC}_{\text{circumcenter}}$ (with control magnitude bounds and relaxed $\mathcal{G}$-connectivity constraints) achieves $\mathcal{T}_{\text{rendezvous}}$;*
2. *on $\mathcal{S}_{\text{LD}}$, the law $\mathcal{CC}_{\text{circumcenter}}$ achieves $\mathcal{T}_{\epsilon\text{-rendezvous}}$*

Furthermore,

1. *if any two agents belong to the same connected component at $\ell \in \mathbb{N}_0$, then they continue to belong to the same connected component subsequently; and*
2. *for each evolution, there exists $P^* = (p_1^*, \dots, p_n^*) \in (\mathbb{R}^d)^n$ such that:*
    1. *the evolution asymptotically approaches P^*, and*
    2. *for each $i, j \in \{1, \dots, n\}$, either $p_i^* = p_j^*$, or $\|p_i^* - p_j^*\|_2 > r$ (for the networks $\mathcal{S}_{\text{disk}}$ and $\mathcal{S}_{\text{LD}}$) or $\|p_i^* - p_j^*\|_\infty > r$ (for the network $\mathcal{S}_{\infty\text{-disk}}$).*

Similar result for visibility networks in non-convex environments

Theorem (Time complexity of circumcenter laws)

For $r \in \mathbb{R}_{>0}$ and $\epsilon \in]0,1[$, the following statements hold:

1. *on the network $\mathcal{S}_{\mathrm{disk}}$, evolving on the real line $\mathbb{R}$ (i.e., with $d = 1$),* $\mathsf{TC}(\mathcal{T}_{\mathrm{rendezvous}}, \mathcal{CC}_{\mathrm{circumcenter}}) \in \Theta(n)$*;*
2. *on the network $\mathcal{S}_{\mathrm{LD}}$, evolving on the real line $\mathbb{R}$ (i.e., with $d = 1$),* $\mathsf{TC}(\mathcal{T}_{(r\epsilon)\text{-rendezvous}}, \mathcal{CC}_{\mathrm{circumcenter}}) \in \Theta(n^2 \log(n\epsilon^{-1}))$*; and*

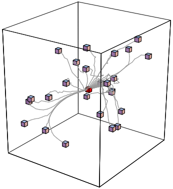

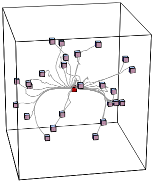

Similar results for visibility networks

Robustness of circumcenter algorithms

Push whole idea further!, e.g., for robustness against link failures

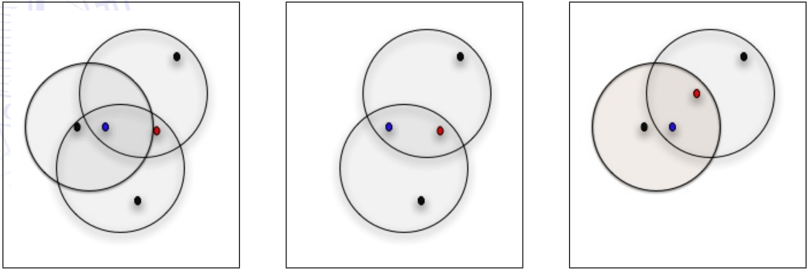

topology G_1 topology G_2 topology G_3

Look at **evolution under link failures** as outcome of nondeterministic evolution under multiple interaction topologies

$$P \longrightarrow \{\text{evolution under } G_1, \text{evolution under } G_2, \text{evolution under } G_3\}$$

Rendezvous

Corollary (Circumcenter algorithm over $\mathcal{G}_{\text{disk}}(r)$ on $\mathbb{R}^d$)

For $\{P_m\}_{m \in \mathbb{N}_0}$ synchronous execution with link failures such that union of any $\ell \in \mathbb{N}$ consecutive graphs in execution has globally reachable node

Then, there exists $(p^*, \dots, p^*) \in \mathsf{diag}((\mathbb{R}^d)^n)$ such that

$$P_m \to (p^*, \dots, p^*) \quad as \quad m \to +\infty$$

Proof uses

$T_{\mathcal{CC},\ell}(P) = \{f_{\mathcal{G}_\ell} \circ \cdots \circ f_{\mathcal{G}_1}(P) \mid$
$\cup_{s=1}^{\ell} \mathcal{G}_i$ has globally reachable node$\}$

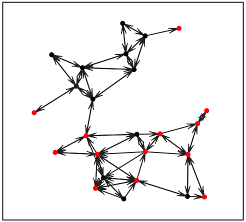

Outline

Objective: optimal task allocation and space partitioning
optimal placement and tuning of sensors

What notion of optimality? What algorithm design?

- **top-down approach:** define aggregate function measuring "goodness" of deployment, then synthesize algorithm that optimizes function
- **bottom-up approach:** synthesize "reasonable" interaction law among agents, then analyze network behavior

Coverage optimization

DESIGN of performance metrics

1. how to cover a region with n minimum-radius overlapping disks?
2. how to design a minimum-distortion (fixed-rate) vector quantizer? (Lloyd '57)
3. where to place mailboxes in a city / cache servers on the internet?

ANALYSIS of cooperative distributed behaviors

4. how do animals share territory? what if every fish in a swarm goes toward center of own dominance region?

Barlow, Hexagonal territories, *Animal Behavior*, 1974

5. what if each vehicle goes to center of mass of own Voronoi cell?
6. what if each vehicle moves away from closest vehicle?

Expected-value multicenter function

Objective: Given sensors/nodes/robots/sites $(p_1, \dots, p_n)$ moving in environment Q achieve **optimal coverage**

- $\phi \colon \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ density
- $f \colon \mathbb{R}_{\geq 0} \to \mathbb{R}$ non-increasing and piecewise continuously differentiable, possibly with finite jump discontinuities

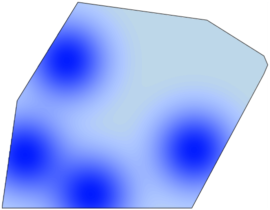

$$\text{maximize} \quad \mathcal{H}_{\text{exp}}(p_1, \dots, p_n) = E_\phi \left[\max_{i \in \{1,\dots,n\}} f(\|q - p_i\|) \right]$$

$\mathcal{H}_{\mathrm{exp}}$-optimality of the Voronoi partition

Alternative expression in terms of Voronoi partition,

$$\mathcal{H}_{\mathrm{exp}}(p_1, \dots, p_n) = \sum_{i=1}^{n} \int_{V_i(P)} f(\|q - p_i\|_2)\phi(q)dq$$

for $(p_1, \dots, p_n)$ distinct

Proposition

Let $P = \{p_1, \dots, p_n\} \in \mathbb{F}(S)$. For any performance function f and for any partition $\{W_1, \dots, W_n\} \subset \mathcal{P}(S)$ of S,

$$\mathcal{H}_{\mathrm{exp}}(p_1, \dots, p_n, V_1(P), \dots, V_n(P)) \geq \mathcal{H}_{\mathrm{exp}}(p_1, \dots, p_n, W_1, \dots, W_n),$$

and the inequality is strict if any set in $\{W_1, \dots, W_n\}$ differs from the corresponding set in $\{V_1(P), \dots, V_n(P)\}$ by a set of positive measure

Distortion problem
$f(x) = -x^2$

$$\mathcal{H}_{\rm dist}(p_1, \dots, p_n) = -\sum_{i=1}^{n} \int_{V_i(P)} \|q - p_i\|_2^2 \phi(q) dq = -\sum_{i=1}^{n} \mathsf{J}_\phi(V_i(P), p_i)$$

($\mathsf{J}_\phi(W, p)$ is moment of inertia). Note

$$\mathcal{H}_{\rm dist}(p_1, \dots, p_n, W_1, \dots, W_n) \\ = -\sum_{i=1}^{n} \mathsf{J}_\phi(W_i, \mathsf{CM}_\phi(W_i)) - \sum_{i=1}^{n} \mathsf{area}_\phi(W_i) \|p_i - \mathsf{CM}_\phi(W_i)\|_2^2$$

Proposition

Let $\{W_1, \dots, W_n\} \subset \mathcal{P}(S)$ be a partition of S. Then,

$$\mathcal{H}_{\rm dist}\big(\mathsf{CM}_\phi(W_1), \dots, \mathsf{CM}_\phi(W_n), W_1, \dots, W_n\big) \\ \geq \mathcal{H}_{\rm dist}(p_1, \dots, p_n, W_1, \dots, W_n),$$

and the inequality is strict if there exists $i \in \{1, \dots, n\}$ for which W_i has non-vanishing area and $p_i \neq \mathsf{CM}_\phi(W_i)$

Area problem
$f(x) = 1_{[0,a]}(x)$, $a \in \mathbb{R}_{>0}$

$$\begin{aligned}
\mathcal{H}_{\mathrm{area},a}(p_1,\dots,p_n) &= \sum_{i=1}^{n} \int_{V_i(P)} 1_{[0,a]}(\|q-p_i\|_2)\phi(q)dq \\
&= \sum_{i=1}^{n} \int_{V_i(P)\cap \overline{B}(p_i,a)} \phi(q)dq \\
&= \sum_{i=1}^{n} \mathrm{area}_\phi(V_i(P)\cap \overline{B}(p_i,a)) = \mathrm{area}_\phi(\cup_{i=1}^{n}\overline{B}(p_i,a)),
\end{aligned}$$

Area, measured according to ϕ, covered by the union of the n balls $\overline{B}(p_1,a),\dots,\overline{B}(p_n,a)$

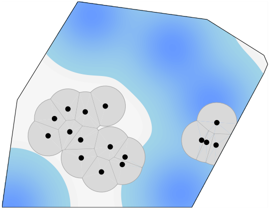

Mixed distortion-area problem

$f(x) = -x^2 \, 1_{[0,a]}(x) + b \cdot 1_{]a,+\infty[}(x)$, with $a \in \mathbb{R}_{>0}$ and $b \leq -a^2$

$$\mathcal{H}_{\text{dist-area},a,b}(p_1, \dots, p_n) = -\sum_{i=1}^{n} \mathsf{J}_\phi(V_{i,a}(P), p_i) + b \, \mathsf{area}_\phi(Q \setminus \cup_{i=1}^{n} \overline{B}(p_i, a)),$$

If $b = -a^2$, f is continuous, we write $\mathcal{H}_{\text{dist-area},a}$. Extension reads

$$\mathcal{H}_{\text{dist-area},a}(p_1, \dots, p_n, W_1, \dots, W_n) = -\sum_{i=1}^{n} \Big(\mathsf{J}_\phi(W_i \cap \overline{B}(p_i, a), p_i) + a^2 \, \mathsf{area}_\phi(W_i \cap (S \setminus \overline{B}(p_i, a))) \Big).$$

Proposition ($\mathcal{H}_{\text{dist-area},a}$-optimality of centroid locations)

Let $\{W_1, \dots, W_n\} \subset \mathcal{P}(S)$ be a partition of S. Then,

$$\mathcal{H}_{\text{dist-area},a}\big(\mathsf{CM}_\phi(W_1 \cap \overline{B}(p_1, a)), \dots, \mathsf{CM}_\phi(W_n \cap \overline{B}(p_n, a)), W_1, \dots, W_n\big) \geq \mathcal{H}_{\text{dist}}(p_1, \dots, p_n, W_1, \dots, W_n),$$

and the inequality is strict if there exists $i \in \{1, \dots, n\}$ for which W_i has non-vanishing area and $p_i \neq \mathsf{CM}_\phi(W_i \cap \overline{B}(p_i, a))$.

Smoothness properties of $\mathcal{H}_{\rm exp}$

$\mathsf{Dscn}(f)$ (finite) discontinuities of f
f_- and f_+, limiting values from the left and from the right

Theorem

Expected-value multicenter function $\mathcal{H}_{\rm exp} : S^n \to \mathbb{R}$ *is*

1. *globally Lipschitz on* S^n*; and*
2. *continuously differentiable on* $S^n \setminus \mathcal{S}_{\rm coinc}$*, where*

$$\frac{\partial \mathcal{H}_{\rm exp}}{\partial p_i}(P) = \int_{V_i(P)} \frac{\partial}{\partial p_i} f(\|q - p_i\|_2)\phi(q)dq + \sum_{a \in \mathsf{Dscn}(f)} \big(f_-(a) - f_+(a)\big) \int_{V_i(P) \cap \partial \overline{B}(p_i,a)} \mathsf{n}_{\mathrm{out}, \overline{B}(p_i,a)}(q)\phi(q)dq$$

$$= \textit{integral over } V_i + \textit{integral along arcs in } V_i$$

Therefore, the gradient of $\mathcal{H}_{\rm exp}$ *is spatially distributed over* $\mathcal{G}_{\rm D}$

Particular gradients

Distortion problem: continuous performance,

$$\frac{\partial \mathcal{H}_{\text{dist}}}{\partial p_i}(P) = 2\,\mathsf{area}_\phi(V_i(P))(\mathsf{CM}_\phi(V_i(P)) - p_i)$$

Area problem: performance has single discontinuity,

$$\frac{\partial \mathcal{H}_{\text{area},a}}{\partial p_i}(P) = \int_{V_i(P)\cap\partial\overline{B}(p_i,a)} \mathsf{n}_{\text{out},\overline{B}(p_i,a)}(q)\phi(q)dq$$

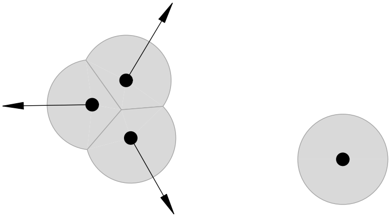

Mixed distortion-area: continuous performance ($b = -a^2$),

$$\frac{\partial \mathcal{H}_{\text{dist-area},a}}{\partial p_i}(P) = 2\,\mathsf{area}_\phi(V_{i,a}(P))(\mathsf{CM}_\phi(V_{i,a}(P)) - p_i)$$

Tuning the optimization problem

Gradients of $\mathcal{H}_{\text{area},a}$, $\mathcal{H}_{\text{dist-area},a,b}$ are distributed over $\mathcal{G}_{\text{LD}}(r)2a$

Robotic agents with range-limited interactions can compute gradients of $\mathcal{H}_{\text{area},a}$ and $\mathcal{H}_{\text{dist-area},a,b}$ as long as $r \geq 2a$

Proposition (Constant-factor approximation of $\mathcal{H}_{\text{dist}}$)

Let $S \subset \mathbb{R}^d$ be bounded and measurable. Consider the mixed distortion-area problem with $a \in]0, \operatorname{diam} S]$ and $b = -\operatorname{diam}(S)^2$. Then, for all $P \in S^n$,

$$\mathcal{H}_{\text{dist-area},a,b}(P) \leq \mathcal{H}_{\text{dist}}(P) \leq \beta^2\, \mathcal{H}_{\text{dist-area},a,b}(P) < 0,$$

where $\beta = \frac{a}{\operatorname{diam}(S)} \in [0, 1]$

Similarly, constant-factor approximations of $\mathcal{H}_{\text{exp}}$

Outline

Geometric-center laws

Uniform networks $\mathcal{S}_{\mathrm{D}}$ and $\mathcal{S}_{\mathrm{LD}}$ of locally-connected first-order agents in a polytope $Q \subset \mathbb{R}^d$ with the Delaunay and r-limited Delaunay graphs as communication graphs

All laws share similar structure

At each communication round each agent performs the following tasks:

- *it transmits its position and receives its neighbors' positions;*
- *it computes a notion of geometric center of its own cell determined according to some notion of partition of the environment*

Between communication rounds, each robot moves toward this center

VRN-CNTRD ALGORITHM

Optimizes distortion $\mathcal{H}_{\mathrm{dist}}$

`Robotic Network:` $\mathcal{S}_{\mathrm{D}}$ in Q, with absolute sensing of own position

`Distributed Algorithm:` VRN-CNTRD

`Alphabet:` $L = \mathbb{R}^d \cup \{\texttt{null}\}$

`function` $\mathrm{msg}(p, i)$

1: **return** p

`function` $\mathrm{ctrl}(p, y)$

1: $V := Q \cap \left(\bigcap \{H_{p,p_{\mathrm{rcvd}}} \mid \text{for all non-null } p_{\mathrm{rcvd}} \in y\} \right)$

2: **return** $\mathsf{CM}_{\phi}(V) - p$

Simulation

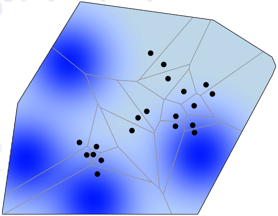

initial configuration

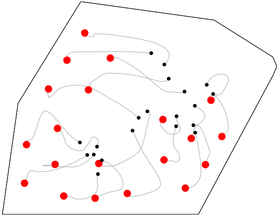

gradient descent

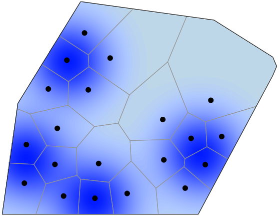

final configuration

For $\epsilon \in \mathbb{R}_{>0}$, the ϵ-*distortion deployment task*

$$\mathcal{T}_{\epsilon\text{-distor-dply}}(P) = \begin{cases} \texttt{true}, & \text{if } \left\| p^{[i]} - \mathsf{CM}_\phi(V^{[i]}(P)) \right\|_2 \leq \epsilon, \ i \in \{1, \dots, n\}, \\ \texttt{false}, & \text{otherwise}, \end{cases}$$

Voronoi-centroid law on planar vehicles

`Robotic Network:` $\mathcal{S}_{\text{vehicles}}$ in Q with absolute sensing of own position

`Distributed Algorithm:` VRN-CNTRD-DYNMCS

`Alphabet:` $L = \mathbb{R}^2 \cup \{\texttt{null}\}$

`function` $\text{msg}((p, \theta), i)$

1: **return** p

`function` $\text{ctrl}((p, \theta), (p_{\text{smpld}}, \theta_{\text{smpld}}), y)$

1: $V := Q \cap \left(\bigcap \left\{ H_{p_{\text{smpld}}, p_{\text{rcvd}}} \mid \text{for all non-null } p_{\text{rcvd}} \in y \right\} \right)$

2: $v := -k_{\text{prop}} (\cos\theta, \sin\theta) \cdot (p - \mathsf{CM}_\phi(V))$

3: $\omega := 2k_{\text{prop}} \arctan \dfrac{(-\sin\theta, \cos\theta) \cdot (p - \mathsf{CM}_\phi(V))}{(\cos\theta, \sin\theta) \cdot (p - \mathsf{CM}_\phi(V))}$

4: **return** (v, ω)

Algorithm illustration

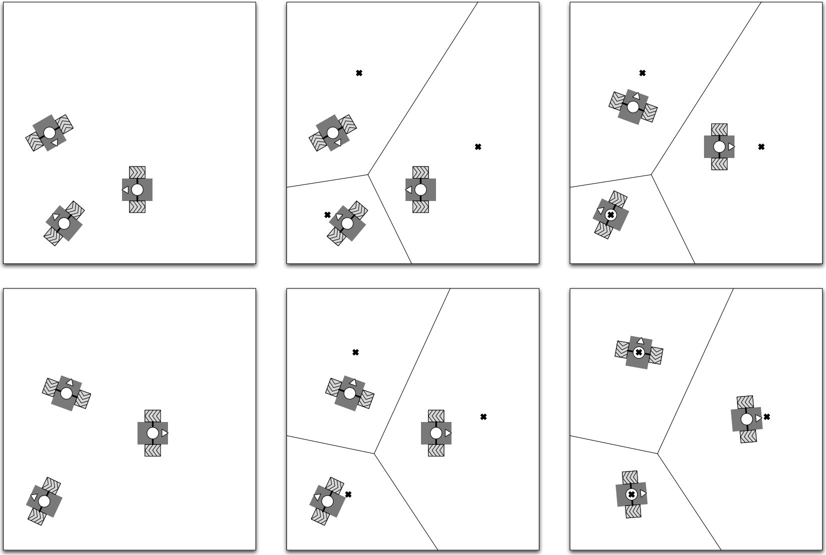

Simulation

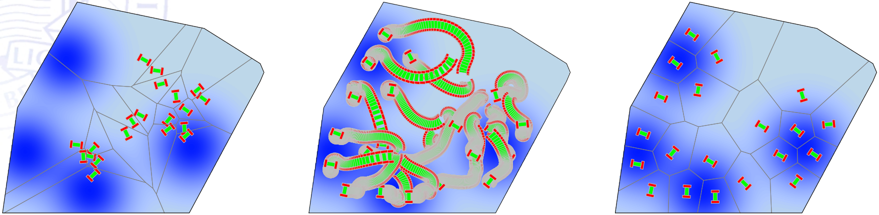

initial configuration | gradient descent | final configuration

`Robotic Network:` $\mathcal{S}_{\mathrm{LD}}$ in Q with absolute sensing of own position and with communication range r

`Distributed Algorithm:` LMTD-VRN-NRML

`Alphabet:` $L = \mathbb{R}^d \cup \{\texttt{null}\}$

`function` $\mathrm{msg}(p, i)$

1: **return** p

`function` $\mathrm{ctrl}(p, y)$

1: $V := Q \cap \left(\bigcap \{H_{p,p_{\mathrm{rcvd}}} \mid \text{for all non-null } p_{\mathrm{rcvd}} \in y\} \right)$
2: $v := \int_{V \cap \partial \overline{B}(p,\frac{r}{2})} \mathsf{n}_{\mathrm{out},\overline{B}(p,\frac{r}{2})}(q)\phi(q)dq$
3: $\lambda_* := \max\left\{ \lambda \mid \delta \mapsto \int_{V \cap \overline{B}(p+\delta v,\frac{r}{2})} \phi(q)dq \text{ is strictly increasing on } [0,\lambda] \right\}$
4: **return** $\lambda_* v$

Simulation

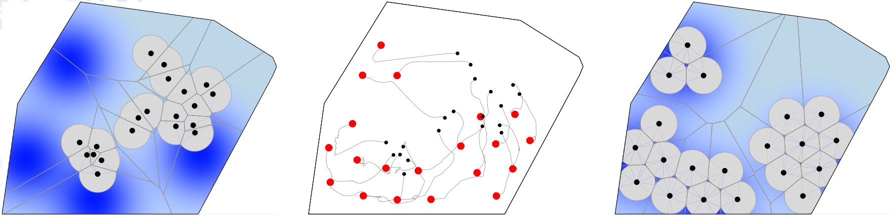

initial configuration | gradient descent | final configuration

For $r, \epsilon \in \mathbb{R}_{>0}$,

$$\mathcal{T}_{\epsilon\text{-}r\text{-area-dply}}(P)$$
$$= \begin{cases} \texttt{true}, & \text{if } \left\| \int_{V^{[i]}(P) \cap \partial \overline{B}(p^{[i]}, \frac{r}{2})} \mathsf{n}_{\mathrm{out}, \overline{B}(p^{[i]}, \frac{r}{2})}(q)\phi(q)dq \right\|_2 \le \epsilon, \; i \in \{1, \dots, n\}, \\ \texttt{false}, & \text{otherwise.} \end{cases}$$

LMTD-VRN-CNTRD algorithm

Optimizes $\mathcal{H}_{\text{dist-area},\frac{r}{2}}$

`Robotic Network:` $\mathcal{S}_{\text{LD}}$ in Q with absolute sensing of own position, and with communication range r

`Distributed Algorithm:` LMTD-VRN-CNTRD
`Alphabet:` $L = \mathbb{R}^d \cup \{\texttt{null}\}$

`function` $\text{msg}(p, i)$

1: **return** p

`function` $\text{ctrl}(p, y)$

1: $V := Q \cap \overline{B}(p, \frac{r}{2}) \cap \left(\bigcap \{H_{p,p_{\text{rcvd}}} \mid \text{for all non-null } p_{\text{rcvd}} \in y\} \right)$
2: **return** $\mathsf{CM}_\phi(V) - p$

Simulation

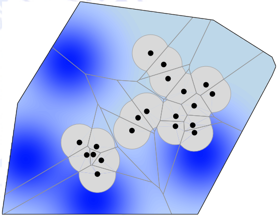

initial configuration

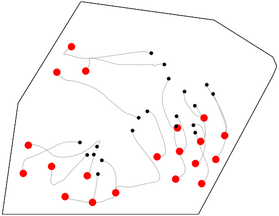

gradient descent

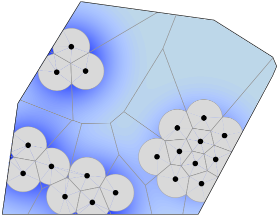

final configuration

For $r, \epsilon \in \mathbb{R}_{>0}$,

$$\mathcal{T}_{\epsilon\text{-}r\text{-distor-area-dply}}(P) = \begin{cases} \texttt{true}, & \text{if } \left\| p^{[i]} - \mathsf{CM}_\phi(V^{[i]}_{\frac{r}{2}}(P)) \right\|_2 \leq \epsilon, \ i \in \{1, \dots, n\}, \\ \texttt{false}, & \text{otherwise.} \end{cases}$$

Optimizing $\mathcal{H}_{\text{dist}}$ via constant-factor approximation

Limited range

run #1: 16 agents, density ϕ is sum of 4 Gaussians, time invariant, 1st order dynamics

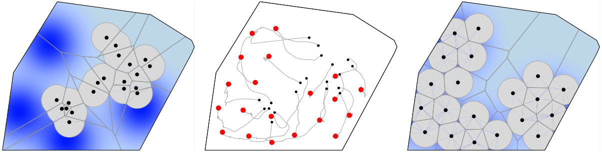

initial configuration | gradient descent of $\mathcal{H}_{\frac{r}{2}}$ | final configuration

Unlimited range

run #2: 16 agents, density ϕ is sum of 4 Gaussians, time invariant, 1st order dynamics

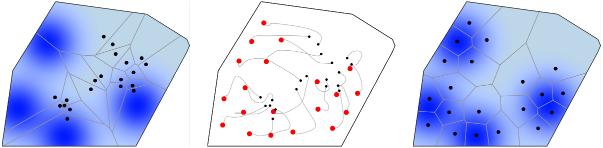

initial configuration | gradient descent of $\mathcal{H}_{\text{exp}}$ | final configuration

Correctness of the geometric-center algorithms

Theorem

For $d \in \mathbb{N}$, $r \in \mathbb{R}_{>0}$ and $\epsilon \in \mathbb{R}_{>0}$, the following statements hold.

1. *on the network $\mathcal{S}_{\mathrm{D}}$, the law $\mathcal{CC}_{\text{VRN-CNTRD}}$ and on the network $\mathcal{S}_{\text{vehicles}}$, the law $\mathcal{CC}_{\text{VRN-CNTRD-DYNMCS}}$ both achieve the ϵ-distortion deployment task $\mathcal{T}_{\epsilon\text{-distor-dply}}$. Moreover, any execution of $\mathcal{CC}_{\text{VRN-CNTRD}}$ and $\mathcal{CC}_{\text{VRN-CNTRD-DYNMCS}}$ monotonically optimizes the multicenter function $\mathcal{H}_{\text{dist}}$;*
2. *on the network $\mathcal{S}_{\mathrm{LD}}$, the law $\mathcal{CC}_{\text{LMTD-VRN-NRML}}$ achieves the ϵ-r-area deployment task $\mathcal{T}_{\epsilon\text{-}r\text{-area-dply}}$. Moreover, any execution of $\mathcal{CC}_{\text{LMTD-VRN-NRML}}$ monotonically optimizes the multicenter function $\mathcal{H}_{\text{area},\frac{r}{2}}$; and*
3. *on the network $\mathcal{S}_{\mathrm{LD}}$, the law $\mathcal{CC}_{\text{LMTD-VRN-CNTRD}}$ achieves the ϵ-r-distortion-area deployment task $\mathcal{T}_{\epsilon\text{-}r\text{-distor-area-dply}}$. Moreover, any execution of $\mathcal{CC}_{\text{LMTD-VRN-CNTRD}}$ monotonically optimizes the multicenter function $\mathcal{H}_{\text{dist-area},\frac{r}{2}}$.*

Time complexity of $\mathcal{CC}_{\text{LMTD-VRN-CNTRD}}$

Assume $\mathsf{diam}(Q)$ is independent of n, r and ϵ

Theorem (Time complexity of Lmtd-Vrn-cntrd law)

Assume the robots evolve in a closed interval $Q \subset \mathbb{R}$, that is, $d = 1$, and assume that the density is uniform, that is, $\phi \equiv 1$. For $r \in \mathbb{R}_{>0}$ and $\epsilon \in \mathbb{R}_{>0}$, on the network $\mathcal{S}_{\text{LD}}$

$$\mathsf{TC}(\mathcal{T}_{\epsilon\text{-}r\text{-distor-area-dply}}, \mathcal{CC}_{\text{LMTD-VRN-CNTRD}}) \in O(n^3 \log(n\epsilon^{-1}))$$

Outline

Deployment: basic behaviors

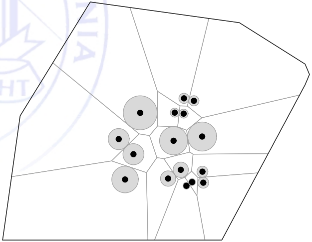

"move away from closest"

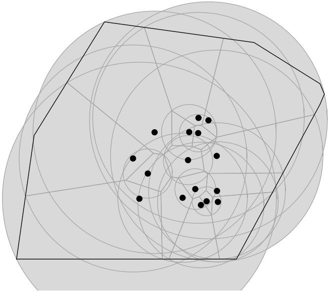

"move towards furthest"

Equilibria? Asymptotic behavior?
Optimizing network-wide function?

Deployment: 1-center optimization problems

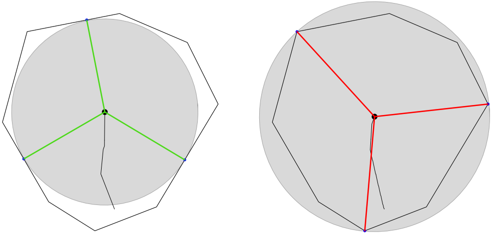

$\mathsf{sm}_Q(p) = \min\{\|p-q\| \mid q \in \partial Q\}$ Lipschitz $0 \in \partial\, \mathsf{sm}_Q(p) \Leftrightarrow p \in \mathsf{IC}(Q)$

$\mathsf{lg}_Q(p) = \max\{\|p-q\| \mid q \in \partial Q\}$ Lipschitz $0 \in \partial\, \mathsf{lg}_Q(p) \Leftrightarrow p = \mathcal{CC}(Q)$

Locally Lipschitz function V are differentiable a.e.

Generalized gradient of V is

$$\partial V(x) = \text{convex closure}\{\lim_{i\to\infty} \nabla V(x_i) \mid x_i \to x\,,\ x_i \notin \Omega_V \cup S\}$$

Deployment: 1-center optimization problems

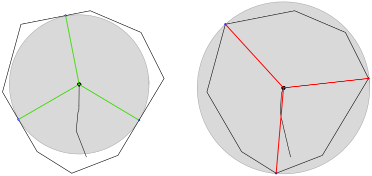

$+$ gradient flow of sm_Q $\quad \dot{p}_i = +\,\mathsf{Ln}[\partial\,\mathsf{sm}_Q](p)$ "move away from closest"

$-$ gradient flow of lg_Q $\quad \dot{p}_i = -\,\mathsf{Ln}[\partial\,\mathsf{lg}_Q](p)$ "move toward furthest"

For X essentially locally bounded, **Filippov solution** of $\dot{x} = X(x)$ is absolutely continuous function $t \in [t_0, t_1] \mapsto x(t)$ verifying

$$\dot{x} \in K[X](x) = \mathsf{co}\{\lim_{i\to\infty} X(x_i) \mid x_i \to x\,,\ x_i \notin S\}$$

For V locally Lipschitz, gradient flow is $\dot{x} = \mathsf{Ln}[\partial V](x)$

Ln = least norm operator

Nonsmooth LaSalle Invariance Principle

Evolution of V along Filippov solution $t \mapsto V(x(t))$ is differentiable a.e.

$$\frac{\mathrm{d}}{\mathrm{d}t} V(x(t)) \in \underbrace{\widetilde{\mathcal{L}}_X V(x(t)) = \{a \in \mathbb{R} \mid \exists v \in K[X](x) \text{ s.t. } \zeta \cdot v = a\,,\ \forall \zeta \in \partial V(x)\}}_{\text{set-valued Lie derivative}}$$

LaSalle Invariance Principle

For S compact and strongly invariant with $\max \widetilde{\mathcal{L}}_X V(x) \leq 0$

Any Filippov solution starting in S converges to largest weakly invariant set contained in $\overline{\left\{x \in S \mid 0 \in \widetilde{\mathcal{L}}_X V(x)\right\}}$

E.g., **nonsmooth gradient flow** $\dot{x} = -\operatorname{Ln}[\partial V](x)$ converges to critical set "move away from closest": $\dot{p}_i = +\,\mathsf{Ln}(\partial\,\mathsf{sm}_{V_i(P)})(p_i)$ — at fixed $V_i(P)$

"move towards furthest": $\dot{p}_i = -\,\mathsf{Ln}(\partial\,\mathsf{lg}_{V_i(P)})(p_i)$ — at fixed $V_i(P)$

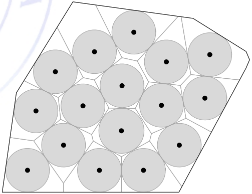

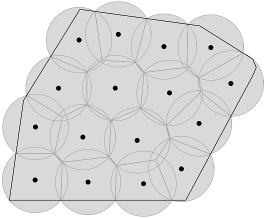

Aggregate objective functions!

$$\mathcal{H}_{\mathrm{sp}}(P) = \min_i \mathsf{sm}_{V_i(P)}(p_i) = \min_{i \neq j} \left[\tfrac{1}{2}\|p_i - p_j\|, \mathsf{dist}(p_i, \partial Q)\right]$$

$$\mathcal{H}_{\mathrm{dc}}(P) = \max_i \mathsf{lg}_{V_i(P)}(p_i) = \max_{q \in Q} \left[\min_i \|q - p_i\|\right]$$

Deployment: multi-center optimization

Critical points of $\mathcal{H}_{\mathbf{sp}}$ and $\mathcal{H}_{\mathbf{dc}}$ (locally Lipschitz)

- If $0 \in \operatorname{int} \partial \mathcal{H}_{\mathrm{sp}}(P)$, then P is strict local maximum, all agents have same cost, and P is **incenter Voronoi configuration**
- If $0 \in \operatorname{int} \partial \mathcal{H}_{\mathrm{dc}}(P)$, then P is strict local minimum, all agents have same cost, and P is **circumcenter Voronoi configuration**

Aggregate functions **monotonically optimized** along evolution

$$\boxed{\min \widetilde{\mathcal{L}}_{\mathsf{Ln}(\partial\, \mathsf{sm}_{\mathcal{V}(P)})} \mathcal{H}_{\mathbf{sp}}(P) \geq 0} \qquad \boxed{\max \widetilde{\mathcal{L}}_{-\,\mathsf{Ln}(\partial\, \mathsf{lg}_{\mathcal{V}(P)})} \mathcal{H}_{\mathbf{dc}}(P) \leq 0}$$

Asymptotic convergence to center Voronoi configurations via nonsmooth LaSalle

Outline

1. Rendezvous and connectivity maintenance
   - The rendezvous objective
   - Maintaining connectivity
   - Circumcenter algorithms
   - Correctness analysis via nondeterministic systems
2. Deployment
   - Expected-value deployment
   - Geometric-center laws
   - Disk-covering and sphere-packing deployment
   - Geometric-center laws
3. Conclusions

Voronoi-circumcenter algorithm

`Robotic Network:` $\mathcal{S}_{\mathrm{D}}$ in Q with absolute sensing of own position

`Distributed Algorithm:` VRN-CRCMCNTR
`Alphabet:` $L = \mathbb{R}^d \cup \{\mathtt{null}\}$

`function` $\mathrm{msg}(p, i)$

1: **return** p

`function` $\mathrm{ctrl}(p, y)$

1: $V := Q \cap \left(\bigcap \{H_{p,p_{\mathrm{rcvd}}} \mid \text{for all non-null } p_{\mathrm{rcvd}} \in y\} \right)$
2: **return** $\mathsf{CC}(V) - p$

Voronoi-incenter algorithm

`Robotic Network:` $\mathcal{S}_{\mathrm{D}}$ in Q with absolute sensing of own position

`Distributed Algorithm:` VRN-NCNTR
`Alphabet:` $L = \mathbb{R}^d \cup \{\texttt{null}\}$

`function` $\mathrm{msg}(p, i)$

1: **return** p

`function` $\mathrm{ctrl}(p, y)$

1: $V := Q \cap \left(\bigcap \{ H_{p, p_{\mathrm{rcvd}}} \mid \text{for all non-null } p_{\mathrm{rcvd}} \in y \} \right)$
2: **return** $x \in \mathsf{IC}(V) - p$

Correctness of the geometric-center algorithms

For $\epsilon \in \mathbb{R}_{>0}$, the ϵ-*disk-covering deployment task*

$$\mathcal{T}_{\epsilon\text{-dc-dply}}(P) = \begin{cases} \texttt{true}, & \text{if } \|p^{[i]} - \mathsf{CC}(V^{[i]}(P))\|_2 \leq \epsilon, \; i \in \{1, \dots, n\}, \\ \texttt{false}, & \text{otherwise}, \end{cases}$$

For $\epsilon \in \mathbb{R}_{>0}$, the ϵ-*sphere-packing deployment task*

$$\mathcal{T}_{\epsilon\text{-sp-dply}}(P) = \begin{cases} \texttt{true}, & \text{if } \mathsf{dist}_2(p^{[i]}, \mathsf{IC}(V^{[i]}(P))) \leq \epsilon, \; i \in \{1, \dots, n\}, \\ \texttt{false}, & \text{otherwise}, \end{cases}$$

Theorem

For $d \in \mathbb{N}$, $r \in \mathbb{R}_{>0}$ and $\epsilon \in \mathbb{R}_{>0}$, the following statements hold.

1. *on the network $\mathcal{S}_{\mathrm{D}}$, any execution of the law $\mathcal{CC}_{\mathrm{VRN\text{-}CRCMCNTR}}$ monotonically optimizes the multicenter function $\mathcal{H}_{\mathrm{dc}}$;*
2. *on the network $\mathcal{S}_{\mathrm{D}}$, any execution of the law $\mathcal{CC}_{\mathrm{VRN\text{-}NCNTR}}$ monotonically optimizes the multicenter function $\mathcal{H}_{\mathrm{sp}}$.*

Summary and conclusions

Examined three basic motion coordination tasks

1. **rendezvous:** circumcenter algorithms
2. **connectivity maintenance:** flexible constraint sets in convex/nonconvex scenarios
3. **deployment:** gradient algorithms based on geometric centers

Correctness and **(1-d) complexity analysis** of geometric-center control and communication laws via

1. Discrete- and continuous-time nondeterministic dynamical systems
2. Invariance principles, stability analysis
3. Geometric structures and geometric optimization

Motion coordination is emerging discipline

Literature is full of exciting problems, solutions, and tools we have not covered

> *Formation control, consensus, cohesiveness, flocking, collective synchronization, boundary estimation, cooperative control over constant graphs, quantization, asynchronism, delays, distributed estimation, spatial estimation, data fusion, target tracking, networks with minimal capabilities, target assignment, vehicle dynamics and energy-constrained motion, vehicle routing, dynamic servicing problems, load balancing, robotic implementations,...*

Too long a list to fit it here!

Book coming out in June 2009

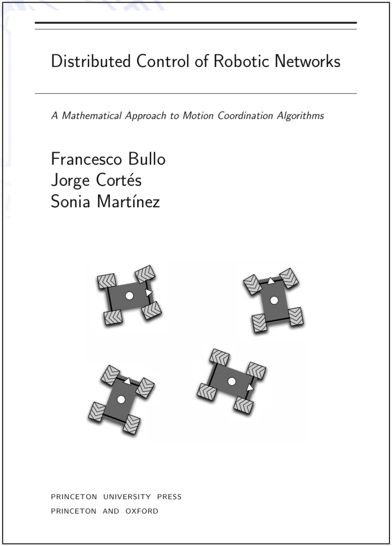
Distributed Control of Robotic Networks

A Mathematical Approach to Motion Coordination Algorithms

Francesco Bullo
Jorge Cortés
Sonia Martínez

PRINCETON UNIVERSITY PRESS
PRINCETON AND OXFORD

Freely available online (forever) at www.coordinationbook.info

- Self-contained exposition of graph-theoretic concepts, distributed algorithms, and complexity measures
- Detailed treatment of averaging and consensus algorithms interpreted as linear iterations
- Introduction of geometric notions such as partitions, proximity graphs, and multicenter functions
- Detailed treatment of motion coordination algorithms for deployment, rendezvous, connectivity maintenance, and boundary estimation

Voronoi partitions

Let $(p_1, \dots, p_n) \in Q^n$ denote the positions of n points

The **Voronoi partition** $\mathcal{V}(P) = \{V_1, \dots, V_n\}$ generated by $(p_1, \dots, p_n)$

$$\begin{aligned} V_i &= \{q \in Q | \ \|q - p_i\| \leq \|q - p_j\|, \forall j \neq i\} \\ &= Q \cap_j \mathcal{HP}(p_i, p_j) \qquad \text{where } \mathcal{HP}(p_i, p_j) \text{ is half plane } (p_i, p_j) \end{aligned}$$

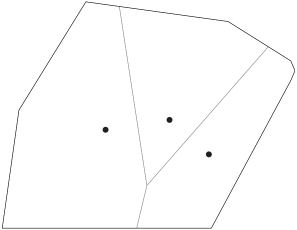
3 generators

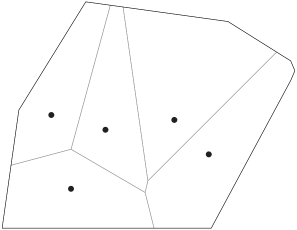
5 generators

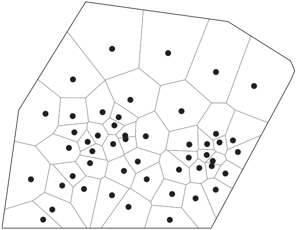
50 generators

Return

Distributed Voronoi computation

Assume: agent with sensing/communication radius R_i

Objective: smallest R_i which provides sufficient information for V_i

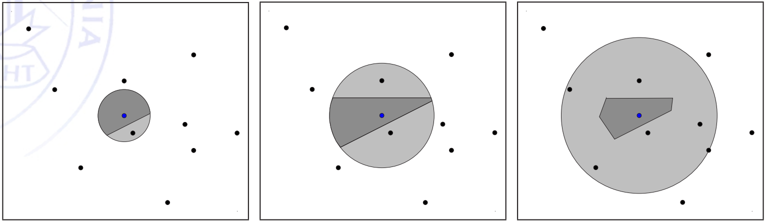

For all i, agent i performs:

1: initialize R_i and compute $\widehat{V}_i = \cap_{j:\|p_i - p_j\| \leq R_i} \mathcal{HP}(p_i, p_j)$
2: **while** $R_i < 2 \max_{q \in \widehat{V}_i} \|p_i - q\|$ **do**
3: $\quad R_i := 2R_i$
4: $\quad$ detect vehicles p_j within radius R_i, recompute $\widehat{V}_i$

Return