Dynamic Vehicle Routing for Robotic Networks
Lecture #5: Extensions to vehicle networks and distributed algorithms

Francesco Bullo[1] Emilio Frazzoli[2] Marco Pavone[2]
Ketan Savla[2] Stephen L. Smith[2]

[1]CCDC
University of California, Santa Barbara
bullo@engineering.ucsb.edu

[2]LIDS and CSAIL
Massachusetts Institute of Technology
{frazzoli,pavone,ksavla,slsmith}@mit.edu

Workshop at the 2010 American Control Conference
Baltimore, Maryland, USA, June 29, 2010, 8:30am to 5:00pm

Lecture outline

1. Motivation and inspiration from biology
2. Intro to comm models, multi-agent networks and distributed algorithms
3. Partitioning with synchronous proximity-graphs communication
4. Partitioning with gossip (asynchronous pair-wise) communication
5. Partitioning with no explicit inter-vehicle communication
   - No explicit communication policy
   - Game-theoretic interpretation

Territory partitioning via *centralized space planning*

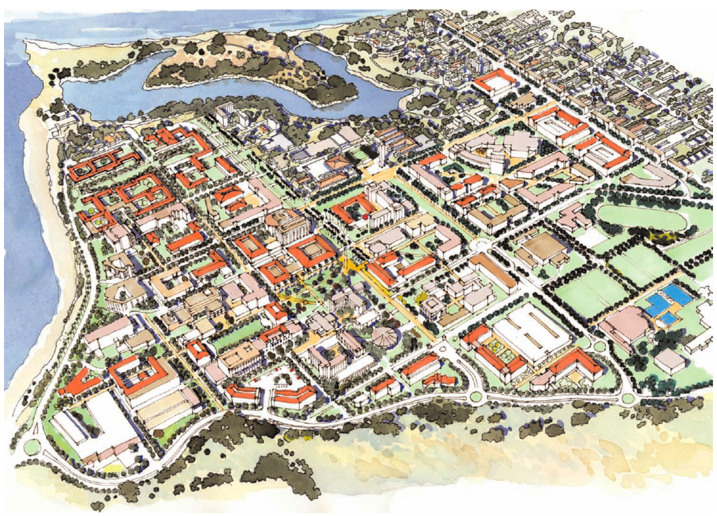

UCSB Campus Development Plan, 2008

Territory partitioning akin to *animal territory dynamics*

Tilapia mossambica, "Hexagonal Territories," Barlow et al, '74

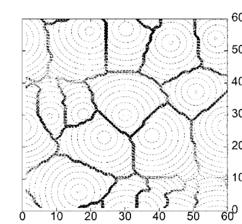

Red harvester ants, "Optimization, Conflict, and Nonoverlapping Foraging Ranges," Adler et al, '03

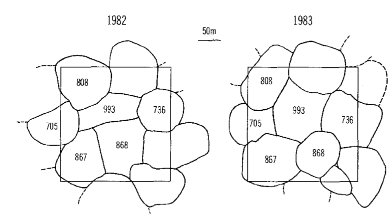

Sage sparrows, "Territory dynamics in a sage sparrows population," Petersen et al '87

Territory partitioning: behaviors and optimality

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?
3. where to place mailboxes in a city / cache servers on the internet?

ANALYSIS of cooperative distributed behaviors

how do animals share territory?
how do they decide foraging ranges?
how do they decide nest locations?

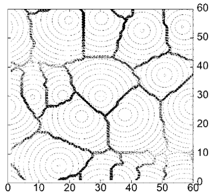

4. what if each robot goes to "center" of own dominance region?
5. what if each robot moves away from closest vehicle?

Lecture outline

1. Motivation and inspiration from biology
2. Intro to comm models, multi-agent networks and distributed algorithms
3. Partitioning with synchronous proximity-graphs communication
4. Partitioning with gossip (asynchronous pair-wise) communication
5. Partitioning with no explicit inter-vehicle communication
   - No explicit communication policy
   - Game-theoretic interpretation

Intro to communication models, multi-agent networks and distributed algorithms

References

1. I. Suzuki and M. Yamashita. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM Journal on Computing*, 28(4):1347–1363, 1999
2. N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1997
3. D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997
4. S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli. On synchronous robotic networks – Part I: Models, tasks and complexity. *IEEE Transactions on Automatic Control*, 52(12):2199–2213, 2007

Objective

1. meaningful + tractable model
2. information/control/communication tradeoffs

Preliminary: Processor network and distributed algorithm

Processor network: group of processors capable to exchange messages along edges and perform local computations

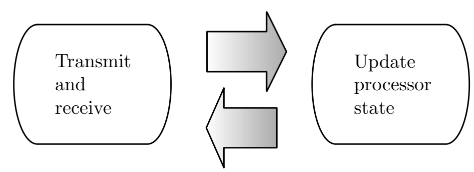

Distributed algorithm for a network of processors consists of

1. $W^{[i]}$, the **processor state set**
2. $\mathbb{A}$, the **communication alphabet**
3. $\mathrm{stf}^{[i]} : W^{[i]} \times \mathbb{A}^n \to W^{[i]}$, the **state-transition map**
4. $\mathrm{msg}^{[i]} : W^{[i]} \to \mathbb{A}$, the **message-generation map**

Robotic network

A robotic network is

1. set of robots moving in space $\mathcal{Q}$
2. interaction graph

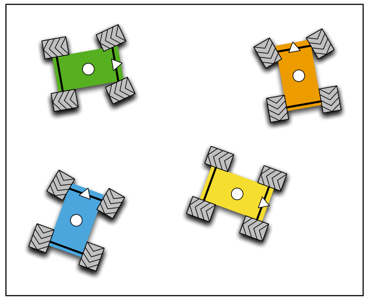

Disk, visibility and Delauney graphs

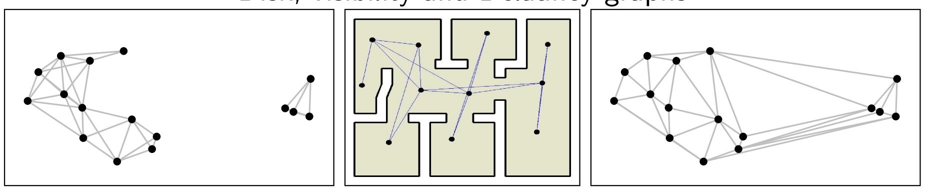

Communication models for robotic networks

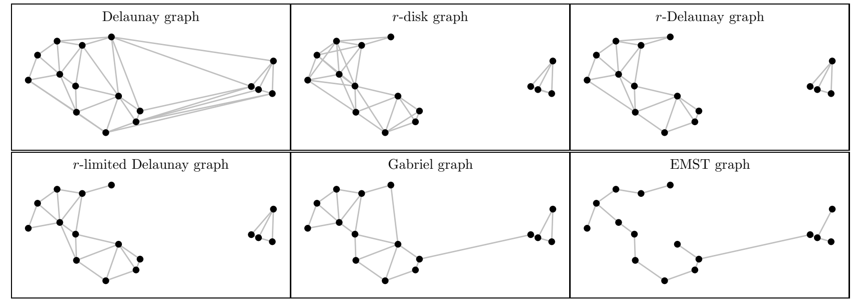

Relevant graphs

1. fixed, directed, balanced
2. switching
3. proximity/geometric or state-dependent
4. random, random geometric (packet losses)

Message model

1. message
2. packet/bits

Sensing model

1. absolute coords other robots
2. absolute coords environment boundary

Communication models for robotic networks

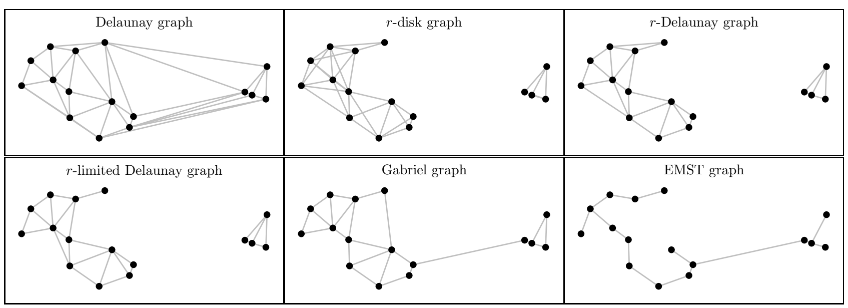

Relevant graphs

1. fixed, directed, balanced
2. switching
3. proximity/geometric or state-dependent
4. random, random geometric (packet losses)

Message model

1. message
2. packet/bits

Sensing model

1. absolute coords other robots
2. absolute coords environment boundary

Synchronous control and communication

1. communication schedule $\mathbb{T} = \{t_\ell\}_{\ell \in \mathbb{N}_0} \subset \mathbb{R}_{\geq 0}$
2. communication alphabet $\mathbb{A}$
3. set of values for processor vars W
4. message-generation function $\text{msg} : \mathbb{T} \times \mathcal{Q} \times W \to \mathbb{A}$
5. state-transition functions $\text{stf} : \mathbb{T} \times W \times \mathbb{A}^N \to W$
6. control function $\text{ctrl} : \mathbb{R}_{\geq 0} \times \mathcal{Q} \times W \times \mathbb{A}^N \to U$

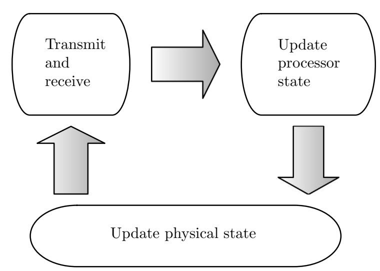

Lecture outline

1. Motivation and inspiration from biology
2. Intro to comm models, multi-agent networks and distributed algorithms
3. Partitioning with synchronous proximity-graphs communication
4. Partitioning with gossip (asynchronous pair-wise) communication
5. Partitioning with no explicit inter-vehicle communication
   - No explicit communication policy
   - Game-theoretic interpretation

M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo. Equitable partitioning policies for mobile robotic networks. *IEEE Transactions on Automatic Control*, 2010. (Submitted Dec 2008 and Aug 2009) to appear

Spatially-distributed policies for DVR

Key idea

Distributed multi-vehicle policy = single-vehicle policy + optimal partitioning + distributed algorithm for partitioning

Light load

Optimal pre-positioning
⇒ median Voronoi diagrams

Heavy load

Workload balance
⇒ equitable partitions

Spatially-distributed policies for DVR

Key idea

Distributed multi-vehicle policy = single-vehicle policy + optimal partitioning + distributed algorithm for partitioning

Light load

Optimal pre-positioning
⇒ median Voronoi diagrams

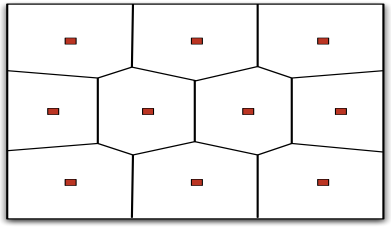

Heavy load

Workload balance
⇒ equitable partitions

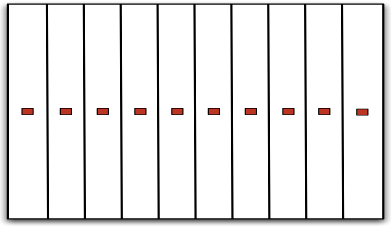

Median Voronoi diagrams (and beyond) with synchronous proximity-graphs communication

Voronoi+centering law

```
At each comm round:
 1: acquire neighbors' positions
 2: compute own dominance region
 3: move towards center of own
    dominance region
```

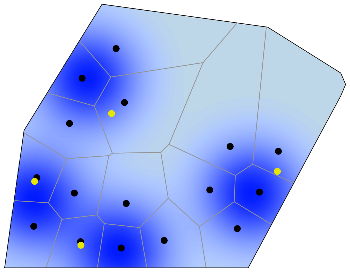

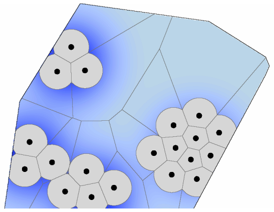

Area-center

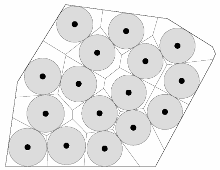

Incenter

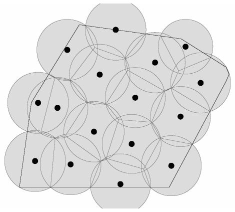

Circumcenter

S. Martínez, J. Cortés, and F. Bullo. Motion coordination with distributed information. *IEEE Control Systems Magazine*, 27(4):75–88, 2007

Experimental Partitioning

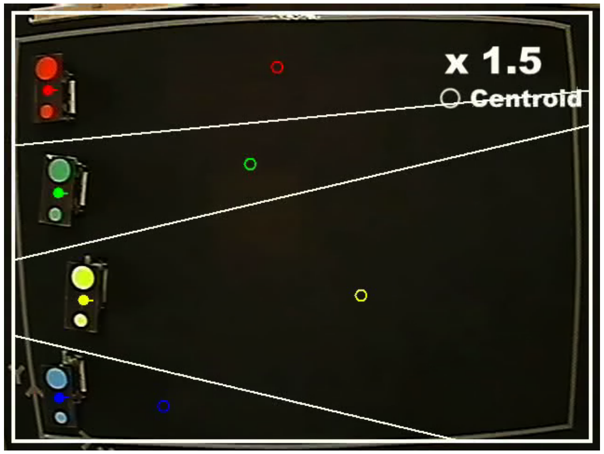

Takahide Goto, Takeshi Hatanaka, Masayuki Fujita
Tokyo Institute of Technology

Experimental Partitioning

Mac Schwager, Brian Julian, Daniela Rus
Distributed Robots Laboratory, MIT

Hardware-in-the-loop Partitioning and DVR for UAVs

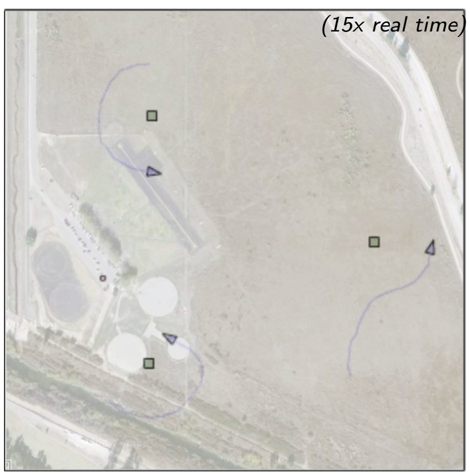

John J. Enright, Chung Hsieh, Emilio Frazzoli
ARES Group, MIT and UCLA

Equitable and median Voronoi diagrams with synchronous proximity-graphs communication

"Ambitious" goal:

Distributed algorithm to partition the workspace according to:

1. median Voronoi diagram (relevant in light-load)
2. equitable (relevant in heavy load)

Voronoi Diagrams

Voronoi partition $\{V_1, \dots, V_m\}$ generated by points $(p_1, \dots, p_m)$:

$$V_i = \{x \in Q \mid \|x - p_i\|^2 \leqslant \|x - p_j\|^2, \forall j \neq i\}$$

In general, an equitable Voronoi Diagram fails to exist...

Equitable and median Voronoi diagrams with synchronous proximity-graphs communication

"Ambitious" goal:

Distributed algorithm to partition the workspace according to:

1. median Voronoi diagram (relevant in light-load)
2. equitable (relevant in heavy load)

Voronoi Diagrams

Voronoi partition $\{V_1, \dots, V_m\}$ generated by points $(p_1, \dots, p_m)$:

$$V_i = \{x \in \mathcal{Q} \mid \|x - p_i\|^2 \leqslant \|x - p_j\|^2, \forall j \neq i\}$$

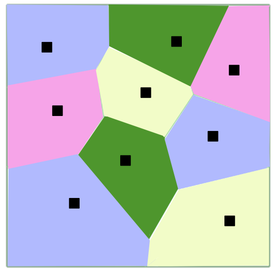

In general, an equitable Voronoi Diagram fails to exist...

Equitable and median Voronoi diagrams with synchronous proximity-graphs communication

"Ambitious" goal:

Distributed algorithm to partition the workspace according to:

1. median Voronoi diagram (relevant in light-load)
2. equitable (relevant in heavy load)

Voronoi Diagrams

Voronoi partition $\{V_1, \dots, V_m\}$ generated by points $(p_1, \dots, p_m)$:

$$V_i = \{x \in \mathcal{Q} \mid \|x - p_i\|^2 \leqslant \|x - p_j\|^2, \forall j \neq i\}$$

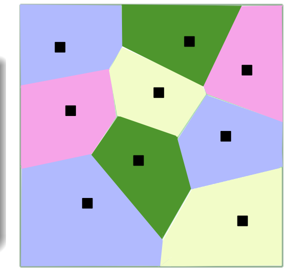

In general, an equitable Voronoi Diagram fails to exist...

FB, EF, MP, KS, SLS (UCSB, MIT) Dynamic Vehicle Routing (Lecture 5/8) 29jun10 @ Baltimore, ACC 18 / 44

Partitioning using Power Diagrams

Power distance

- $p = (p_1, \dots, p_m)$ collection of points in $\mathcal{Q} \subset \mathbb{R}^2$
- each p_i has assigned a weight $w_i \in \mathbb{R}$
- power distance function $d_P(x, p_i;\ w_i) = \|x - p_i\|^2 - w_i$

Power Diagrams

Power diagram $\{V_1, \dots, V_m\}$ generated by weighted points $\big((p_1, w_1), \dots, (p_m, w_m)\big)$:

- $V_i = \{x \in \mathcal{Q} \mid \|x - p_i\|^2 - w_i \leqslant \|x - p_j\|^2 - w_j, \forall j \neq i\}$

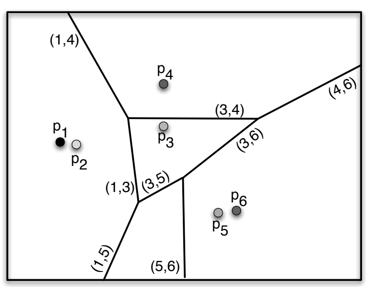

Existence theorem for Power diagrams

Existence theorem

Let $p = (p_1, \dots, p_m)$ be the positions of $m \geq 1$ distinct points in $\mathcal{Q}$. Then there exist weights $(w_1, \dots, w_m)$ such that the corresponding Power diagram is equitable with respect to φ

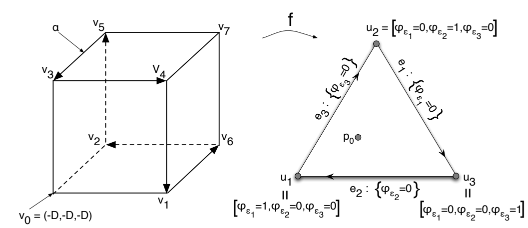

Existence theorem for Power diagrams

Existence theorem

Let $p = (p_1, \dots, p_m)$ be the positions of $m \geq 1$ distinct points in $\mathcal{Q}$. Then there exist weights $(w_1, \dots, w_m)$ such that the corresponding Power diagram is equitable with respect to φ

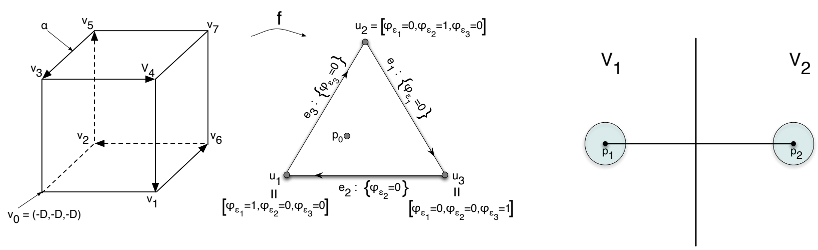

Existence theorem for Power diagrams

Existence theorem

Let $p = (p_1, \dots, p_m)$ be the positions of $m \geq 1$ distinct points in $\mathcal{Q}$. Then there exist weights $(w_1, \dots, w_m)$ such that the corresponding Power diagram is equitable with respect to φ

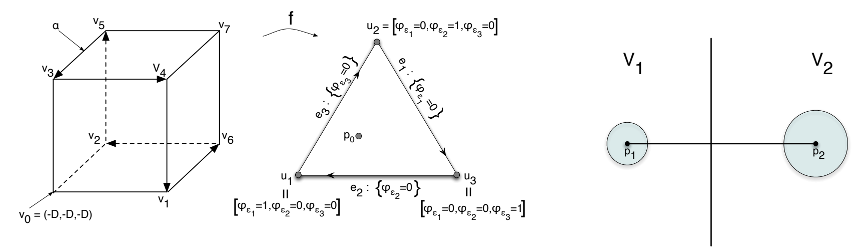

Existence theorem for Power diagrams

Existence theorem

Let $p = (p_1, \dots, p_m)$ be the positions of $m \geq 1$ distinct points in $\mathcal{Q}$. Then there exist weights $(w_1, \dots, w_m)$ such that the corresponding Power diagram is equitable with respect to φ

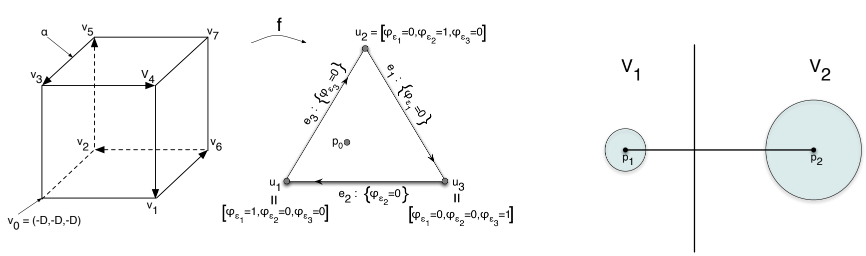

Existence theorem for Power diagrams

Existence theorem

Let $p = (p_1, \dots, p_m)$ be the positions of $m \geq 1$ distinct points in $\mathcal{Q}$. Then there exist weights $(w_1, \dots, w_m)$ such that the corresponding Power diagram is equitable with respect to φ

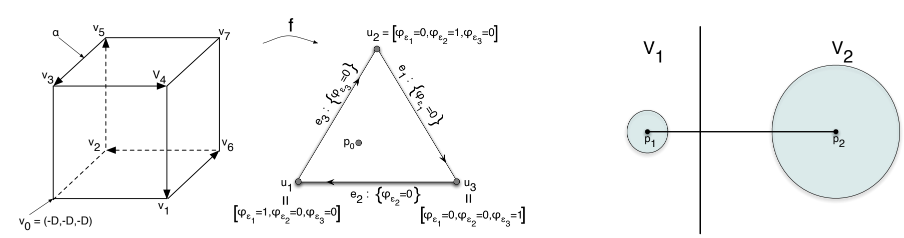

Gradient descent law for equitable partitioning

- w_i locally controlled by vehicle i
- locational optimization function

$$\mathcal{H}(w) \doteq \sum_{i=1}^{m} \Big(\int_{V_i(w)} \varphi(x)dx \Big)^{-1} = \sum_{i=1}^{m} |V_i(w)|_{\varphi}^{-1}$$

- spatially-distributed gradient: $\frac{\partial \mathcal{H}}{\partial w_i} = \sum_{j \in N_i} \alpha_{ij}^{\varphi} \Big(\frac{1}{|V_j|_{\varphi}^2} - \frac{1}{|V_i|_{\varphi}^2} \Big)$

Gradient law for equitable partitioning

At each comm round:

1: acquire neighbors' positions
2: compute own dominance region
3: $w_i \leftarrow w_i - \gamma \frac{\partial \mathcal{H}}{\partial w_i}$

Convergence result

Theorem (Convergence)

Assume that the p_i's are distinct. Then, the w_i's converge asymptotically to a vector of weights that yields an equitable Power diagram

- guaranteed convergence for any set of *distinct* points
  $\Rightarrow$ **global convergence result**
- distributed over the dual graph of the induced Power diagram
  $\Rightarrow$ **communication, on average, with six neighbors**
- adjusting the weights sufficient to obtain an equitable diagram
  $\Rightarrow$ **move the p_i's to optimize secondary objectives**

Including the median Voronoi diagram property

Close to Voronoi:

- basic idea: keep the weights *close* to zero
- modify the gradient descent law as

$$\dot{w}_i = -\frac{\partial \mathcal{H}}{\partial w_i} - w_i, \qquad \frac{\partial \mathcal{H}}{\partial p_i} \cdot \dot{p}_i - \frac{\partial \mathcal{H}}{\partial w_i} w_i = 0$$

Motion toward the median:

- basic idea: add a term that enforces computation of the median
- gradient term for computation of the median:

$$\frac{\partial \mathcal{H}_{\mathrm{FW}}}{\partial p_i} = \int_{V_i} \frac{p_i - x}{\|p_i - x\|} \varphi(x) dx$$

- modify the gradient descent law as

$$\dot{w}_i = -\frac{\partial \mathcal{H}}{\partial w_i}, \qquad \dot{p}_i = \frac{\partial \mathcal{H}_{\mathrm{FW}}}{\partial p_i} \psi\Big(\frac{\partial \mathcal{H}}{\partial p_i}, \frac{\partial \mathcal{H}_{\mathrm{FW}}}{\partial p_i} \Big)$$

Simulation

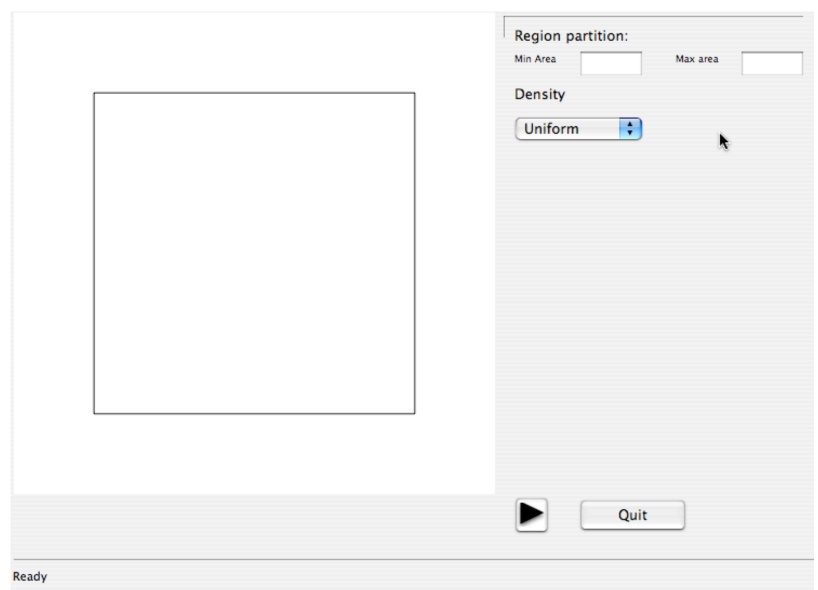

Lecture outline

Partitioning with gossip communication

Voronoi+centering law requires:

1. synchronous communication
2. communication along edges of dual graph

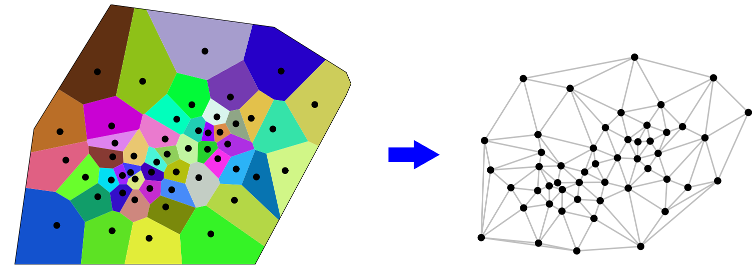

Minimalist coordination

- is synchrony necessary?
- is it sufficient to communicate peer-to-peer (gossip)?
- what are minimal requirements?

Gossip (asynchronous pair-wise) partitioning policy

1. Random communication between two regions
2. Compute two centers
3. Compute bisector of centers
4. Partition two regions by bisector

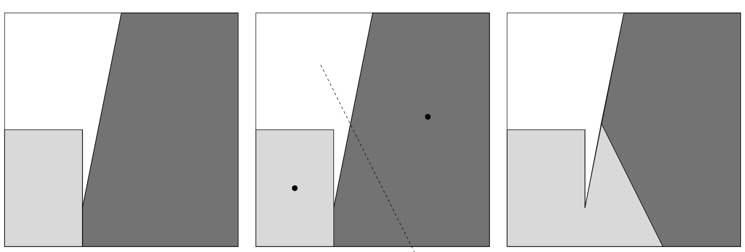

F. Bullo, R. Carli, and P. Frasca. Gossip coverage control for robotic networks: Dynamical systems on the the space of partitions. *SIAM Review*, January 2010. Submitted

Indoor example implementation

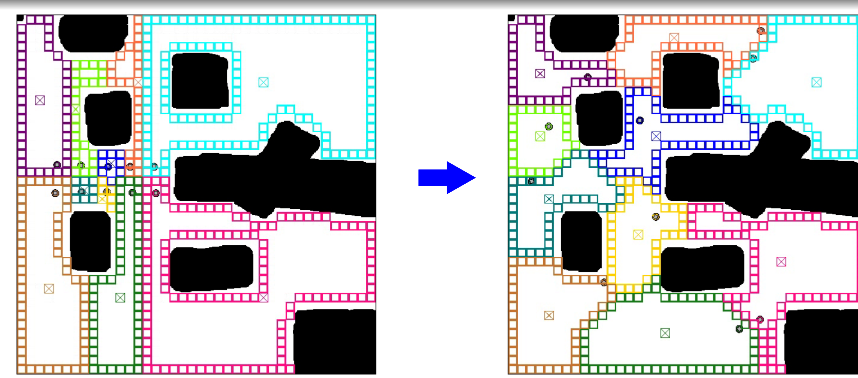

- Player/Stage platform
- realistic robot models in discretized environments
- integrated wireless network model & obstacle-avoidance planner

J. W. Durham, R. Carli, P. Frasca, and F. Bullo. Discrete partitioning and coverage control with gossip communication. In *ASME Dynamic Systems and Control Conference*, Hollywood, CA, October 2009

Peer-to-peer convergence analysis (proof sketch 1/3)

Lyapunov function for peer-to-peer territory partitioning

$$\mathcal{H}(v) = \sum_{i=1}^{n} \int_{v_i} f(\| \operatorname{center}(v_i) - q \|) \phi(q) dq$$

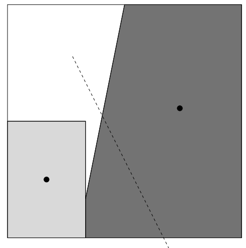

1. state space is not finite-dimensional
   - non-convex disconnected polygons
   - arbitrary number of vertices
2. peer-to-peer map is not deterministic, ill-defined and discontinuous
   - two regions could have same centers

Peer-to-peer convergence analysis (proof sketch 1/3)

Lyapunov function for peer-to-peer territory partitioning

$$\mathcal{H}(v) = \sum_{i=1}^{n} \int_{v_i} f(\| \operatorname{center}(v_i) - q \|) \phi(q) dq$$

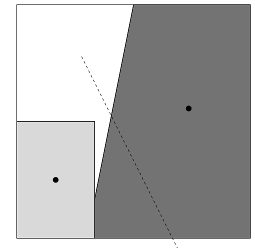

1. state space is not finite-dimensional
   - non-convex disconnected polygons
   - arbitrary number of vertices
2. peer-to-peer map is not deterministic, ill-defined and discontinuous
   - two regions could have same centers

The space of partitions (proof sketch 2/3)

Definition (Space of finitely-convex partitions)

Fix ℓ, the set v is collections of n subsets of Q, $\{v_1, \dots, v_n\}$, such that

1. $v_1 \cup \dots \cup v_n = Q$,
2. $\operatorname{interior}(v_i) \cap \operatorname{interior}(v_j) = \emptyset$ if $i \neq j$, and
3. each v_i is union of ℓ convex sets

Given sets A and B, symmetric distance is:

$$d_\Delta(A, B) = \operatorname{area}\Big((A \cup B) \setminus (A \cap B)\Big)$$

Theorem (topological properties of the space of finitely-convex partitions)

Partition space with $(u, v) \mapsto \sum_{i=1}^{n} d_\Delta(u_i, v_i)$ is metric and compact

Convergence with persistent switches (proof sketch 3/3)

- X is metric space
- finite collection of maps $T_i : X \to X$ for $i \in I$
- consider sequences $\{x_\ell\}_{\ell \geq 0} \subset X$ with

$$x_{\ell+1} = T_{i(\ell)}(x_\ell)$$

Assume:

1. $W \subset X$ compact and positively invariant for each T_i
2. $U : W \to \mathbb{R}$ decreasing along each T_i
3. U and T_i are continuous on W
4. there exists probability $p \in]0, 1[$ such that, for all indices $i \in I$ and times ℓ, we have $\operatorname{Prob}\big[x_{\ell+1} = T_i(x_\ell) \,|\, \text{past}\big] \geq p$

If $x_0 \in W$, then almost surely

$$x_\ell \to \text{(intersection of sets of fixed points of all } T_i) \cap U^{-1}(c)$$

Lecture outline

1. Motivation and inspiration from biology
2. Intro to comm models, multi-agent networks and distributed algorithms
3. Partitioning with synchronous proximity-graphs communication
4. Partitioning with gossip (asynchronous pair-wise) communication
5. Partitioning with no explicit inter-vehicle communication
   - No explicit communication policy
   - Game-theoretic interpretation

A. Arsie, K. Savla, and E. Frazzoli. Efficient routing algorithms for multiple vehicles with no explicit communications. *IEEE Transactions on Automatic Control*, 54(10):2302–2317, 2009

Motivation

Gradient policy

- Cost function: $\mathcal{H}(p) = \sum_{j=1}^{n} \int_{V_j(p)} \|q - p_j\| \varphi(q) dq$
- $\dot{p}_i = -\frac{\partial \mathcal{H}}{\partial p_i}(p) = -\int_{V_i(p)} \frac{\partial}{\partial p_i} \|q - p_i\| \varphi(q) dq$
- $p(t)$ converges to a critical point of $\mathcal{H}(p)$

- Similar result using the gossip partitioning policy

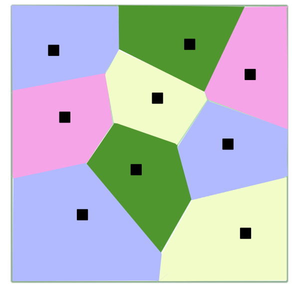

Salient Features

- Explicit agent-to-agent communication
- Needs knowledge of φ

Motivation

Gradient policy

- Cost function: $\mathcal{H}(p) = \sum_{j=1}^{n} \int_{V_j(p)} \|q - p_j\| \varphi(q) dq$
- $\dot{p}_i = -\frac{\partial \mathcal{H}}{\partial p_i}(p) = -\int_{V_i(p)} \frac{\partial}{\partial p_i} \|q - p_i\| \varphi(q) dq$
- $p(t)$ converges to a critical point of $\mathcal{H}(p)$

- Similar result using the gossip partitioning policy

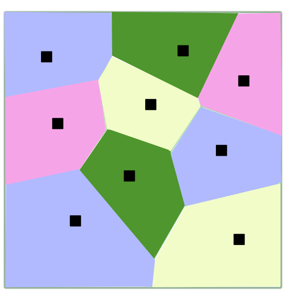

Salient Features

- Explicit agent-to-agent communication
- Needs knowledge of φ

Partitioning with no explicit inter-vehicle communication

Inspiration: Distributed MacQueen algorithm

- Pick any m generator points $(p_1, \dots, p_m) \in \mathcal{Q}^m$
- Iteratively sample points q_j according to probability density function φ
- At each iteration j:
  - Assign the sampled point to the nearest generator $i^*(q_j) \in \{1, \dots, m\}$
  - update the position of generator i^* as

$$p_{i^*} = \frac{(\#\text{pts assigned in past})\, p_{i^*} + q_j}{\#\text{pts assigned in past} + 1}$$

Algorithms

No sensor policy

For all time t, each vehicle moves towards:

- the nearest outstanding task; else,
- the (nearest) point minimizing the average distance to tasks *serviced in the past*

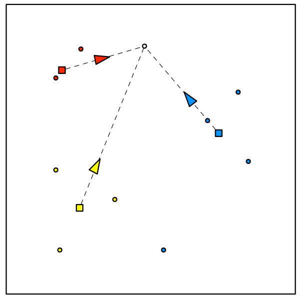

Sensor-based policy

For all time t, each vehicle moves towards:

- the nearest among outstanding tasks that is closest to it than other vehicles; else,
- the (nearest) point minimizing the average distance to tasks *serviced in the past*

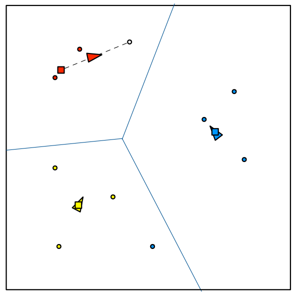

Algorithms

No sensor policy

For all time t, each vehicle moves towards:

- the nearest outstanding task; else,
- the (nearest) point minimizing the average distance to tasks *serviced in the past*

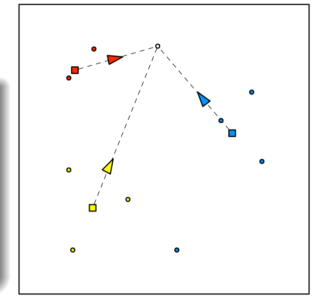

Sensor-based policy

For all time t, each vehicle moves towards:

- the nearest among outstanding tasks that is closest to it than other vehicles; else,
- the (nearest) point minimizing the average distance to tasks *serviced in the past*

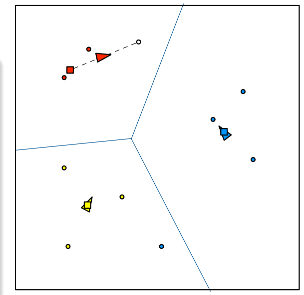

Illustration

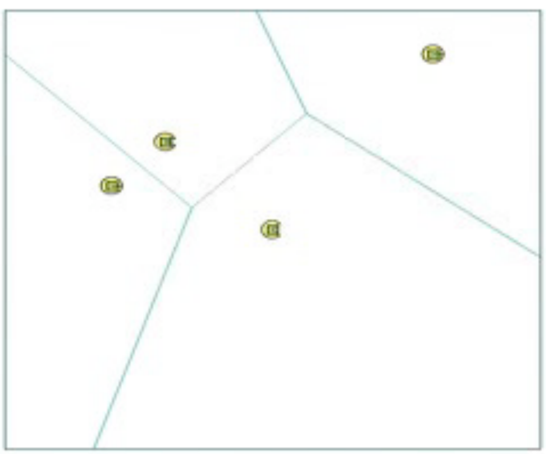

Differences with the MacQueen algorithm

- At each iteration, the no-communication algorithm computes the "Fermat-Weber (FW) point" with respect to the set of tasks serviced by a vehicle; MacQueen algorithm computes the mean

$$\mathrm{FW}_i = \mathrm{argmin}_{p_i \in \mathcal{Q}} \sum_{q \in \text{past tasks}_i} \|q - p_i\|$$

$$\mathrm{Mean}_i = \frac{1}{|\text{past tasks}_i|} \sum_{q \in \text{past tasks}_i} q$$

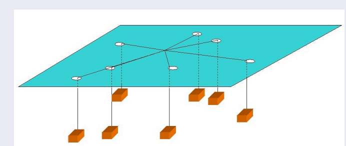

- No simple recursion like the MacQueen algorithm $\rightarrow$ need to store locations of all the tasks serviced in the past
- Sequence of FW points exhibit more complex behavior than the sequence of means.

Analysis of the algorithm

- $p_i(t)$: loitering location of agent i at time t
- Sufficient to study convergence of $(p_1(t), \dots, p_m(t))$

Convergence result

$p(t)$ converges to a critical point of $\mathcal{H}(p)$ with probability one.

Key steps in the proof

- Convergence of the sequence of Fermat-Weber points:
  - $C_i(t) := \{y \in \mathcal{Q} \mid \|\sum_{q \in \text{past tasks}_i} \text{vers}(y - q)\| \leq 1\}$
  - By the properties of the Fermat-Weber point, $p_i(t_j) \in C_i(t_j)$
  - Prove that $p_i(t_{j+1}) \in C_i(t_j)$
  - Prove that $\lim_{j \to \infty} \text{diam}(C_i(t_j)) = 0$ with prob. 1; this implies $p_i(t_j) \to p_i^*$ with prob 1
- p_i^* is the median of its own Voronoi cell

Analysis of the algorithm

- $p_i(t)$: loitering location of agent i at time t
- Sufficient to study convergence of $(p_1(t), \dots, p_m(t))$

Convergence result

$p(t)$ converges to a critical point of $\mathcal{H}(p)$ with probability one.

Key steps in the proof

- Convergence of the sequence of Fermat-Weber points:
  - $C_i(t) := \{y \in \mathcal{Q} \mid \|\sum_{q \in \text{past tasks}_i} \text{vers}(y - q)\| \leq 1\}$
  - By the properties of the Fermat-Weber point, $p_i(t_j) \in C_i(t_j)$
  - Prove that $p_i(t_{j+1}) \in C_i(t_j)$
  - Prove that $\lim_{j \to \infty} \text{diam}(C_i(t_j)) = 0$ with prob. 1; this implies $p_i(t_j) \to p_i^*$ with prob 1
- p_i^* is the median of its own Voronoi cell

Lecture outline

Coverage as a geometric game

Strategies

- $p = (p_1, \dots, p_m) \in \mathcal{Q}^m$
- When a new task is generated, every vehicle move towards its location

Utility Function

- Upon its generation, each task offers continuous reward at rate unity
- A task expires as soon as two vehicles are present at its location or after $\text{diam}(\mathcal{Q})$ time, whichever occurs first.
- Utility function: expected time spent alone at the next task location

$$\mathcal{U}_i(p_i, p_{-i}) = \mathbb{E}_\varphi[R_i(p, q)] = \mathbb{E}_\varphi\left[\max\left\{0, \min_{j \neq i} \|p_j - q\| - \|p_i - q\|\right\}\right]$$

Coverage as a geometric game

Strategies

- $p = (p_1, \ldots, p_m) \in \mathcal{Q}^m$
- When a new task is generated, every vehicle move towards its location

Utility Function

- Upon its generation, each task offers continuous reward at rate unity
- A task expires as soon as two vehicles are present at its location or after diam($\mathcal{Q}$) time, whichever occurs first.
- Utility function: expected time spent alone at the next task location

$$\mathcal{U}_i(p_i, p_{-i}) = \mathbb{E}_\varphi[R_i(p, q)] = \mathbb{E}_\varphi\left[\max\left\{0, \min_{j \neq i} \|p_j - q\| - \|p_i - q\|\right\}\right]$$

Properties of the Game

- Potential function: $\psi(p) = -\sum_{i=1}^m \int_{V_i(p)} \|p_i - q\|\varphi(q)dq$
- The coverage spatial game is a potential game ($\mathcal{U}_i(p) = \psi(p) - \psi(p_{-i})$)
- $\mathcal{U}$ is a Wonderful Life utility function

Characterization of Equilibria

critical point of $\mathcal{H}$ $\Longleftrightarrow$ pure Nash equilibrium

Properties of the Game

- Potential function: $\psi(p) = -\sum_{i=1}^m \int_{V_i(p)} \|p_i - q\|\varphi(q)dq$
- The coverage spatial game is a potential game ($\mathcal{U}_i(p) = \psi(p) - \psi(p_{-i})$)
- $\mathcal{U}$ is a Wonderful Life utility function

Characterization of Equilibria

critical point of $\mathcal{H}$ $\Longleftrightarrow$ pure Nash equilibrium

No communication policy as a learning algorithm

Complete Information

$\dot{p}_i = \frac{\partial}{\partial p_i}\mathcal{U}_i(p) = -\int_{V_i(p)} \frac{p_i - q}{\|p_i - q\|}\varphi(q)dq \implies$ gradient descent policy

Limited information

- No knowledge of φ
- No inter-agent communication

Approximations

- Empirical Utility Maximization: $p_i(t) = \text{argmax}_{x \in \mathcal{Q}} \sum_{q \sim \varphi} R_i(x, p_{-i}, q)$
- $\hat{R}_i(x, p_{-i}, q) = \text{diam}(\mathcal{Q}) - \|x - q\|$ if vehicle i reaches task located at q first, else $\hat{R}_i(x, p_{-i}, q) = 0$.

No communication policy as a learning algorithm

Complete Information

$\dot{p}_i = \frac{\partial}{\partial p_i}\mathcal{U}_i(p) = -\int_{V_i(p)} \frac{p_i - q}{\|p_i - q\|}\varphi(q)dq \implies$ gradient descent policy

Limited information

- No knowledge of φ
- No inter-agent communication

Approximations

- Empirical Utility Maximization: $p_i(t) = \text{argmax}_{x \in \mathcal{Q}} \sum_{q \sim \varphi} R_i(x, p_{-i}, q)$
- $\hat{R}_i(x, p_{-i}, q) = \text{diam}(\mathcal{Q}) - \|x - q\|$ if vehicle i reaches task located at q first, else $\hat{R}_i(x, p_{-i}, q) = 0$.

No communication policy as a learning algorithm

Complete Information

$\dot{p}_i = \frac{\partial}{\partial p_i}\mathcal{U}_i(p) = -\int_{V_i(p)} \frac{p_i - q}{\|p_i - q\|}\varphi(q)dq \implies$ gradient descent policy

Limited information

- No knowledge of φ
- No inter-agent communication

Approximations

- Empirical Utility Maximization: $p_i(t) = \text{argmax}_{x \in \mathcal{Q}} \sum_{q \sim \varphi} R_i(x, p_{-i}, q)$
- $\hat{R}_i(x, p_{-i}, q) = \text{diam}(\mathcal{Q}) - \|x - q\|$ if vehicle i reaches task located at q first, else $\hat{R}_i(x, p_{-i}, q) = 0$.

Lecture outline

1. Motivation and inspiration from biology
2. Intro to comm models, multi-agent networks and distributed algorithms
3. Partitioning with synchronous proximity-graphs communication
4. Partitioning with gossip (asynchronous pair-wise) communication
5. Partitioning with no explicit inter-vehicle communication
   - No explicit communication policy
   - Game-theoretic interpretation

Workshop Structure and Schedule

8:00-8:30am	*Coffee Break*	
8:30-9:00am	Lecture #1:	Intro to dynamic vehicle routing
9:05-9:50am	Lecture #2:	Prelims: graphs, TSPs and queues
9:55-10:40am	Lecture #3:	The single-vehicle DVR problem
10:40-11:00am	*Break*	
11:00-11:45pm	Lecture #4:	The multi-vehicle DVR problem
11:45-1:10pm	*Lunch Break*	
1:10-2:10pm	Lecture #5:	Extensions to vehicle networks
2:15-3:00pm	Lecture #6:	Extensions to different demand models
3:00-3:20pm	*Coffee Break*	
3:20-4:20pm	Lecture #7:	Extensions to different vehicle models
4:25-4:40pm	Lecture #8:	Extensions to different task models
4:45-5:00pm		Final open-floor discussion