Personal Perspectives on Robotic Coordination and Bioinspiration

Francesco Bullo

Center for Control,
Dynamical Systems & Computation
University of California at Santa Barbara
http://motion.me.ucsb.edu

The 3rd International Symposium on Swarm Behavior and Bio-Inspired Robotics
November 20-22, Okinawa, Japan

Acknowledgments

Jorge Cortés*	G. Notarstefano	Stephen Smith	Fabio Pasqualetti
Sonia Martínez*	Anurag Ganguli	Shaunak Bopardikar	Rush Patel
Emilio Frazzoli*	Ketan Savla	Karl Obermeyer	Pushkarini Agharkar
Marco Pavone*	Ruggero Carli*	Joey Durham	Jeff Peters
Paolo Frasca*	Sara Susca	Vaibhav Srivastava	Mishel George

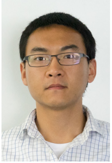

Xiaoming Duan, UC Santa Barbara

Rush Patel, Systems Technology

Pushkarini Agharkar, Google

Mishel George, Doordash

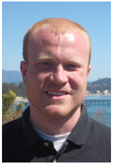

Jeff Peters, UTRC

Coordination in multi-agent systems

Animals and robots:

- each agent senses its immediate environment,
- communicates/interacts with others,
- processes information gathered, and
- takes local action in response

Classic examples of motion coordination

Geese flying in formation

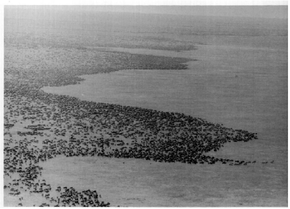

Wildebeest herd in the Serengeti

Fish swarm

Outline

1. **Territory partitioning**
2. **Routing through known locations**
3. **Searching evaders**

Territory partitioning is … art

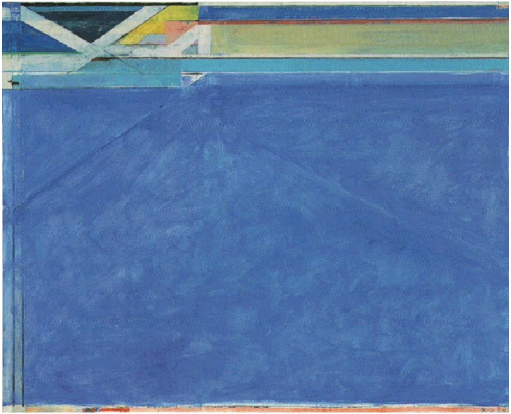

abstract expressionism
"Ocean Park No. 27" and "Ocean Park No. 129"
by Richard Diebenkorn (1922-1993), inspired by aerial landscapes

Territory partitioning ... centralized space planning

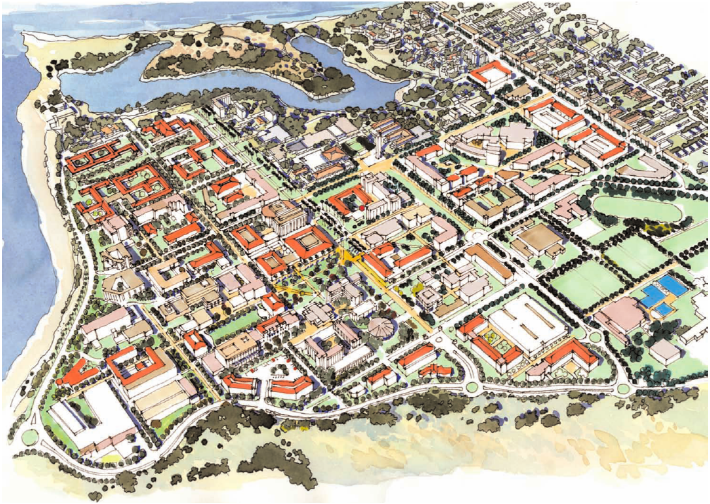

UC Santa Barbara Campus Development Plan, 2008

Territory partitioning ... undemocratic voting districts

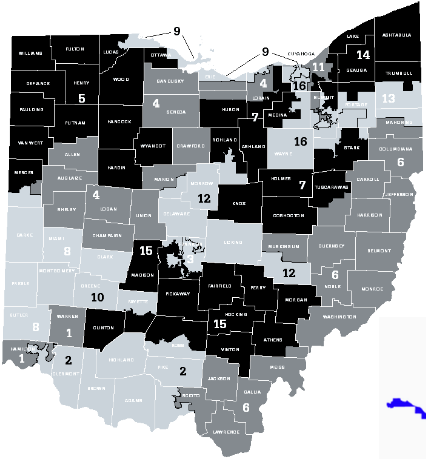

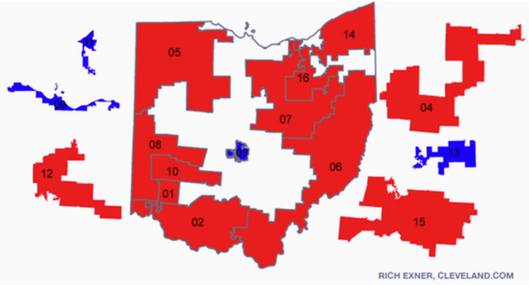

Gerrymandering the Ohio voting districts

Territory partitioning is … animal territory dynamics

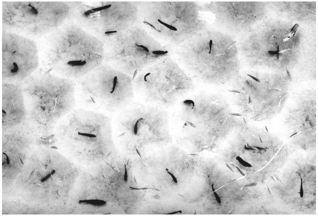

Tilapia mossambica, "Hexagonal Territories," Barlow '74

Territory partitioning is ... animal territory dynamics

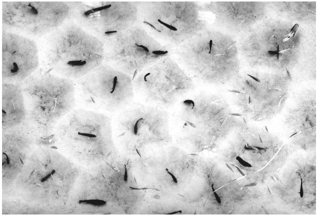

Tilapia mossambica, "Hexagonal Territories," Barlow '74

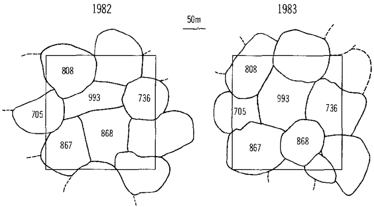

Sage sparrows, "Territory dynamics in a sage sparrows population," Petersen and Best '87

Territory partitioning is ... animal territory dynamics

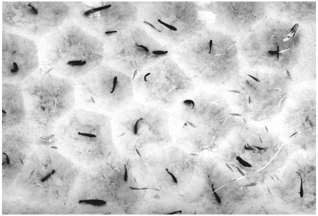

Tilapia mossambica, “Hexagonal Territories,” Barlow ’74

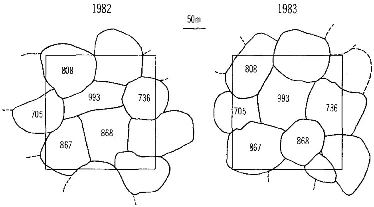

Sage sparrows, “Territory dynamics in a sage sparrows population,” Petersen and Best ’87

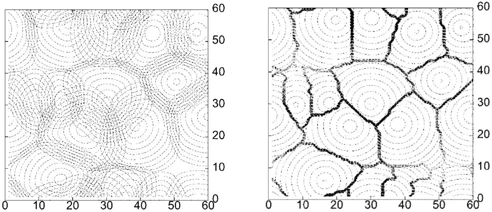

Red harvester ants, “Optimization, Conflict, and Nonoverlapping Foraging Ranges,” Adler and Gordon ’03

Territory partitioning: behaviors and optimality

ANALYSIS of cooperative distributed behaviors

1. how do animals share territory?
   how do they decide where to forage?
   how do they decide nest locations?

2. what if each robot goes to "center" of own dominance region?

DESIGN of performance metrics

3. how to cover a region with n minimum-radius overlapping disks?
4. how to design a minimum-distortion (fixed-rate) vector quantizer?

Multi-center functions

- n robots at $p = \{p_1, \ldots, p_n\}$
- environment is partitioned into $v = \{v_1, \ldots, v_n\}$
- customer arrives and waits for service:

$$H(p, v) = \int_{v_1} \|q - p_1\| dq \; + \; \cdots \; + \; \int_{v_n} \|q - p_n\| dq$$

Multi-center functions

- n robots at $p = \{p_1, \ldots, p_n\}$
- environment is partitioned into $v = \{v_1, \ldots, v_n\}$
- customer arrives and waits for service:

$$H(p, v) = \int_{v_1} \|q - p_1\| dq \; + \; \cdots \; + \; \int_{v_n} \|q - p_n\| dq$$

$$H(p, v) = \sum_{i=1}^{n} \int_{v_i} f(\|q - p_i\|)\phi(q) dq$$

- $\phi : \mathbb{R}^2 \to \mathbb{R}_{\geq 0}$ density
- $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ penalty function

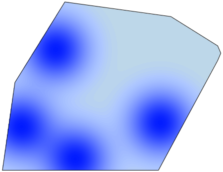

Optimal partitioning

The **Voronoi partition** $\{V_1, \ldots, V_n\}$ generated by points $(p_1, \ldots, p_n)$

$$
\begin{aligned}
V_i &= \{q \in Q \mid \ \|q - p_i\| \le \|q - p_j\|, \ \forall j \neq i\} \\
&= Q \bigcap_j (\text{half plane between } i \text{ and } j, \ \text{containing } i)
\end{aligned}
$$

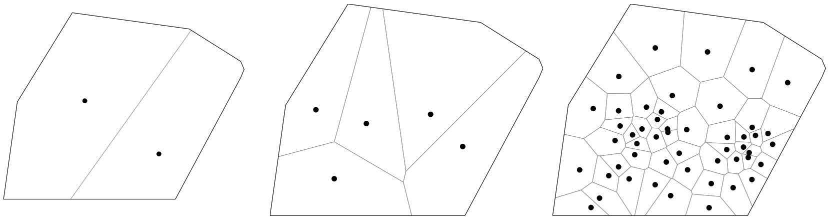

Optimal centering (for region v with density ϕ)

function of p	minimizer = center
$p \mapsto \int_v \|q-p\| \phi(q) dq$	**Fermat–Weber point** (or **median**)
$p \mapsto \int_v \|q-p\|^2 \phi(q) dq$	**centroid** (or **center of mass**)
$p \mapsto \text{area}(v \cap \text{disk}(p, r))$	***r*-area center**
$p \mapsto$ radius of largest disk centered at p enclosed inside v	**incenter**
$p \mapsto$ radius of smallest disk centered at p enclosing v	**circumcenter**

From online Encyclopedia of Triangle Centers

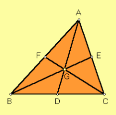

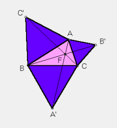

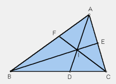

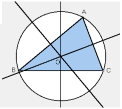

From optimality conditions to algorithms

$$H(p, v) = \int_{v_1} f(\|q - p_1\|)\phi(q)dq + \cdots + \int_{v_n} f(\|q - p_n\|)\phi(q)dq$$

1. at fixed positions, optimal partition is Voronoi
2. at fixed partition, optimal positions are "generalized centers"
3.

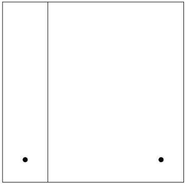

From optimality conditions to algorithms

$$H(p,v) = \int_{v_1} f(\|q-p_1\|)\phi(q)dq + \cdots + \int_{v_n} f(\|q-p_n\|)\phi(q)dq$$

1. at fixed positions, optimal partition is Voronoi
2. at fixed partition, optimal positions are "generalized centers"
3. alternate v-p optimization $\implies$ local opt = center Voronoi partition

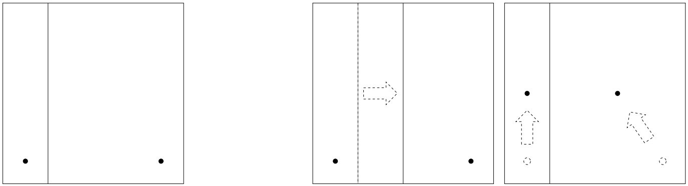

Voronoi+centering algorithm

Voronoi+centering law

At each comm round:

1: acquire neighbors' positions
2: compute own dominance region
3: move towards center of own dominance region

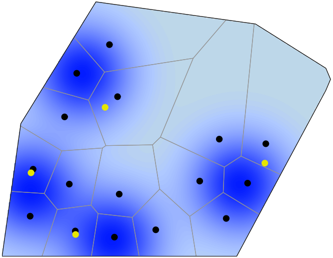

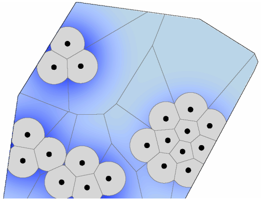

Area-center

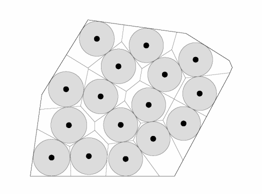

Incenter

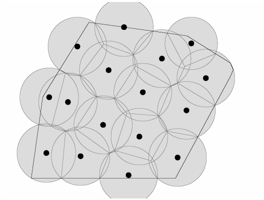

Circumcenter

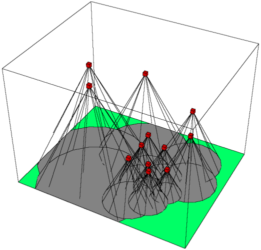

3D coverage

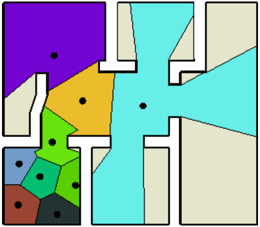

nonconvex converage

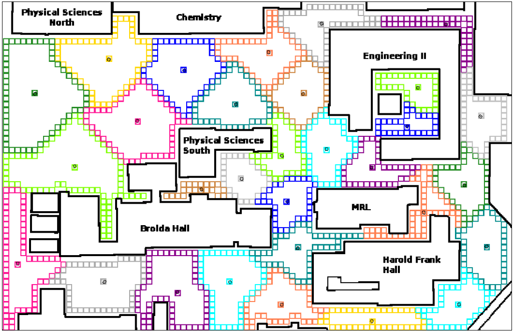

discrete peer-to-peer

territory partitioning:

- well developed in engineering
- existing connection with the study of animal behavior

  ... even if cost functions may differ

Potential for future research and collaborations

1. do animals achieve optimal territory partitioning?
   recent game-theoretic work proposes "elaborate" coordination algorithms to achieve constant-factor optimality
2. how to incorporate exploration in robotic territory partitioning?
3. how about animal behavior in nonconvex environment

Outline

1. **Territory partitioning**
2. **Routing through known locations**
3. **Searching evaders**

Routing through known locations

- customers appear sequentially randomly space/time
- robotic network *knows* locations and provides service
- Goal: distributed algorithm that minimizes wait time

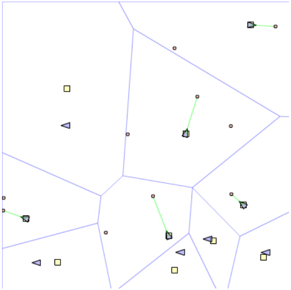

Algo #1: Receding-horizon shortest-path policy

Receding-horizon Shortest-Path (RH-SP)

For $\eta \in (0, 1]$, single agent performs:

1: while no customers, move to center

2: while customers waiting

1. compute shortest path through current customers
2. service η-fraction of path

Algo #1: Receding-horizon shortest-path policy

Receding-horizon Shortest-Path (RH-SP)

For $\eta \in (0, 1]$, single agent performs:

```
1: while no customers, move to center
2: while customers waiting
```

1. compute shortest path through current customers
2. service η-fraction of path

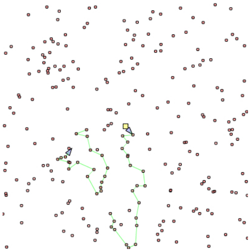

- shortest path is NP-hard, but effective heuristics available
- delay is optimal in light traffic
- delay is constant-factor optimal in high traffic

Algo #2: Load balancing via territory partitioning

RH-SP + Partitioning

For $\eta \in (0, 1]$, agent i performs:

1: compute own cell v_i in optimal partition

2: apply RH-SP policy on v_i

Asymptotically constant-factor optimal in light and high traffic

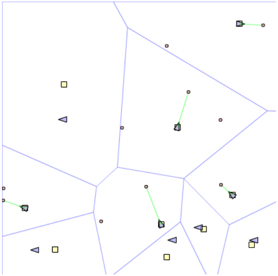

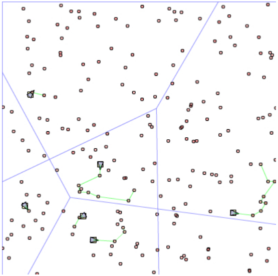

I am unaware of comparable animal behavior

Potential for future research and collaborations

1. can animals can solve shortest-path problems?
2. do they adopt simpler efficient heuristics?

Outline

1. **Territory partitioning**
2. **Routing through known locations**
3. **Searching evaders**

Search and surveillance

Design motion strategies to search unpredictably and quickly

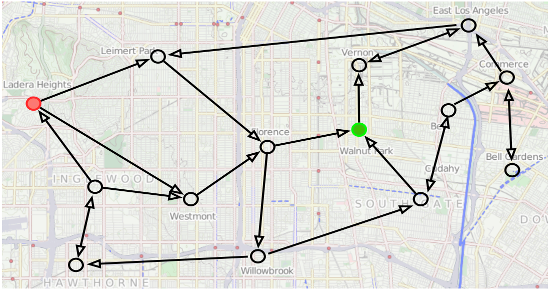

- **pursuer / predator**
- **evader / prey**

Design motion strategies to search unpredictably and quickly

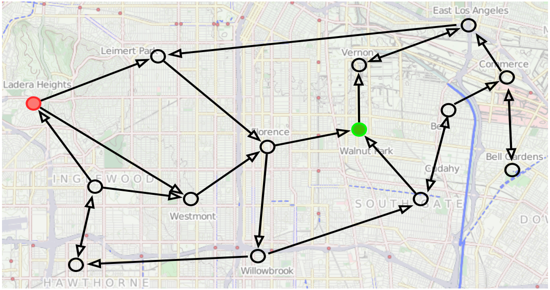

- **pursuer / predator**
- **evader / prey**

How many steps on average for predator to detect prey?
How to minimize? How to maximize?

Stochastic surveillance: Motivating example 2/2

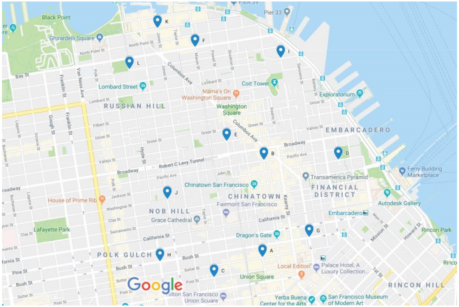

- San Francisco
- crime rate at 12 locations
- all-to-all driving times (quantized in minutes)

Stochastic surveillance: Motivating example 2/2

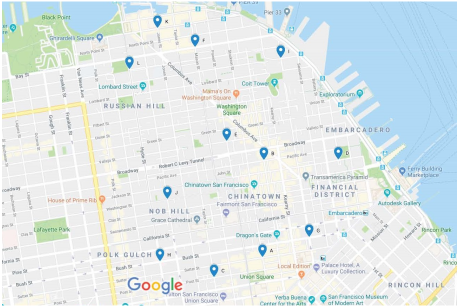

- San Francisco
- crime rate at 12 locations
- all-to-all driving times (quantized in minutes)

police:

- on patrol, moves around city

bank robber:

- robber picks bank
- attacks at time with minimum detection likelihood

Approach: Markov chains for routing and planning

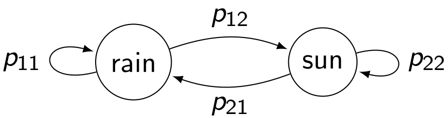

Advantages of adopting Markov chains:

1. rich behavior
2. finite-dimensional optimization problem
3. well-defined notion of unpredictability: entropy
4. well-defined notion of speed: hitting time

The entropy of a discrete random variable $X \in \{1, \ldots, k\}$ is

$$\mathbb{H}(X) = -\sum_{i=1}^{k} p_i \log p_i$$

Unbiased coin:	$\mathbb{P}[X = \text{Head}] = 0.5$	$\mathbb{H}(X) = \mathbf{0.69}$
Biased coin:	$\mathbb{P}[X = \text{Head}] = 0.75$	$\mathbb{H}(X) = \mathbf{0.56}$
Predictable coin:	$\mathbb{P}[X = \text{Head}] = 1$	$\mathbb{H}(X) = 0$

The entropy of what variable?

The entropy of what variable?

#1: sequence of random locations

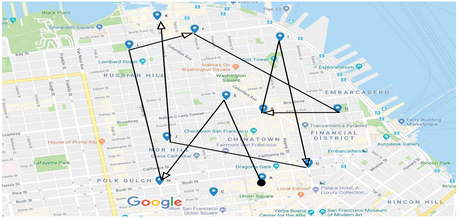

The entropy of what variable?

#1: sequence of random locations

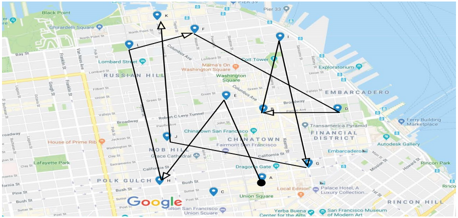

#2: sequence of return times

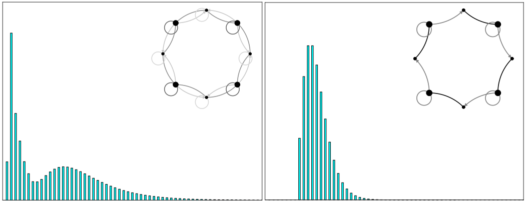

location entropy vs. return time entropy

1. MaxReturnEntropy

$$\max_P \mathbb{H}_{\text{return-time}}(P)$$

2. MaxLocationEntropy

$$\max_P \mathbb{H}_{\text{location}}(P)$$

3. MinCaptureTime: $\min_P \mathbb{E}[\text{capture time}(P)]$
   simplified intruder model: random attack location / time

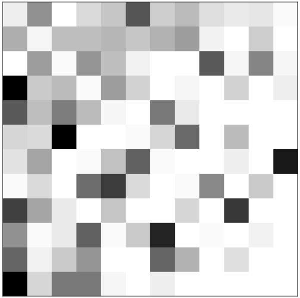

(a) MaxReturnEntropy

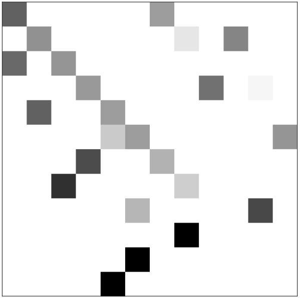

(b) MinCaptureTime

Pixel image of Markov chains: i^{th} row are transition probabilities out of i

- MinCaptureTime chain is close to “TSP + self weights”
- MaxReturnEntropy chain is dense, i.e., has higher entropy

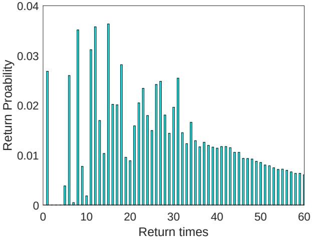

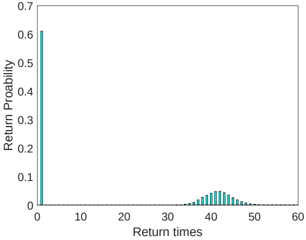

MaxReturnEntropy: $\mathbb{P}[0-10m] \approx 10\%$, $\mathbb{P}[10-20m] \approx 25\%$, $\mathbb{P}[20-30m] \approx 20\%$, ...

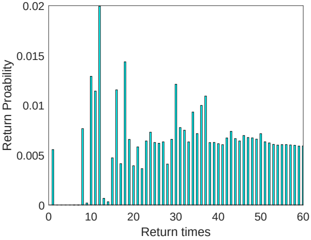

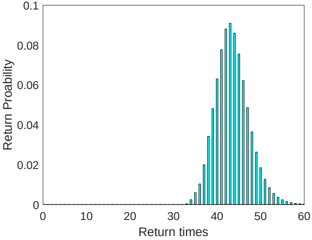

Rational intruder:

- Picks a node i to attack with probability π_i
- Collects the return time statistics of the pursuer
- Attacks when the pursuer is absent for s_i timesteps since last visit

$$s_i = \underset{0 \le s \le S_i}{\operatorname{argmin}} \left\{ \sum_{k=1}^{\tau} \mathbb{P}(T_{ii} = s + k \mid T_{ii} > s) \right\},$$

where τ is the attack duration and S_i is determined by the degree of impatience δ, i.e., $\mathbb{P}(T_{ii} \ge S_i) \le \delta$

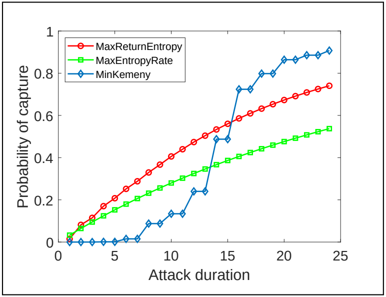

4x4 grid (unit travel times)

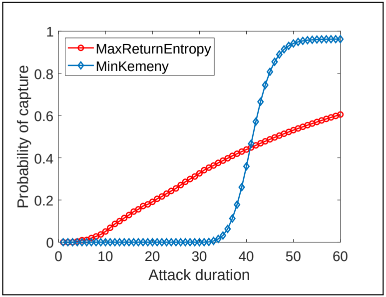

SF map

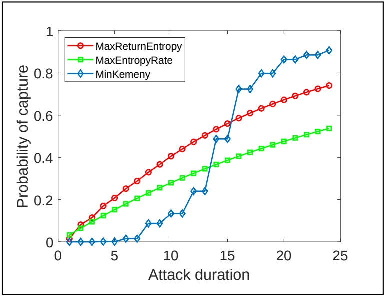

4x4 grid (unit travel times)

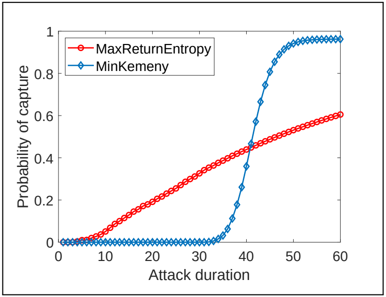

SF map

Lessons

- 4×4 grid: MaxReturnEntropy $>$ MaxLocationEntropy
- 4×4 grid: MaxReturnEntropy $>$ MinCaptureTime for short attacks
- SF: MaxReturnEntropy $>$ MinCaptureTime for short attacks

- search strategies by optimizing transition probabilities
- I am unaware of comparable animal behavior

Potential for future research and collaborations

1. how do animal play this search/hide games?
2. do they ever move unpredictably?

References

1. S. Martínez, J. Cortés, and F. Bullo. Motion coordination with distributed information.
IEEE Control Systems, 27(4):75–88, 2007.
doi:10.1109/MCS.2007.384124

2. F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems.
Proceedings of the IEEE, 99(9):1482–1504, 2011.
doi:10.1109/JPROC.2011.2158181

3. X. Duan, M. George, and F. Bullo. Markov chains with maximum return time entropy for robotic surveillance.
IEEE Transactions on Automatic Control, 2019.
doi:10.1109/TAC.2019.2906473

Freely-downloadable textbooks

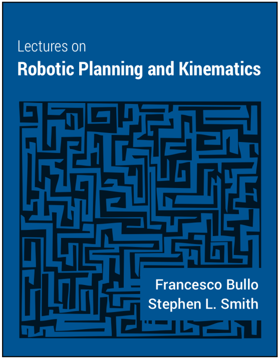

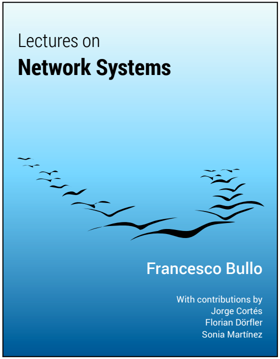

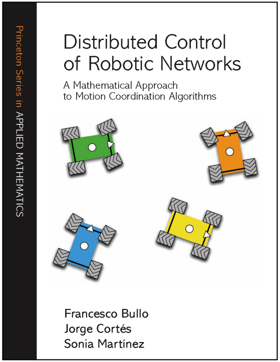

F. Bullo and S. L. Smith. *Lectures on Robotic Planning and Kinematics*.
Unpublished Manuscript, 2019.
URL: http://motion.me.ucsb.edu/book-lrpk

F. Bullo. *Lectures on Network Systems*.
Kindle Direct Publishing, 1.3 edition, July 2019, ISBN 978-1986425643.
URL: http://motion.me.ucsb.edu/book-lns

F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*.
Princeton University Press, 2009, ISBN 978-0-691-14195-4.
URL: http://www.coordinationbook.info

Lectures on Network Systems

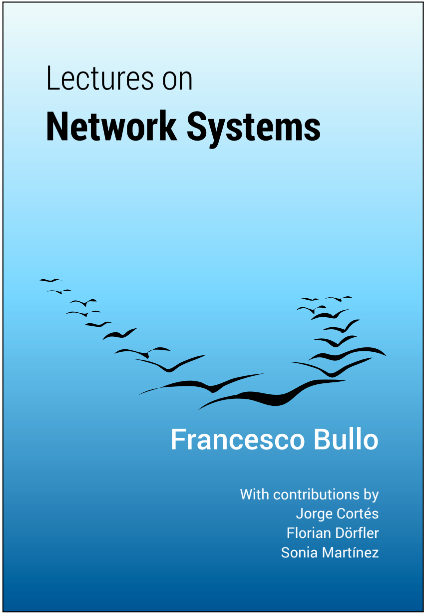

Lectures on Network Systems, Kindle Direct Publishing, 1.3 edition, 2019, ISBN 978-1-986425-64-3

1. Self-Published and Print-on-Demand at:
   https://www.amazon.com/dp/1986425649
2. PDF Freely available at
   http://motion.me.ucsb.edu/book-lns:
   For students: free PDF for download
   For instructors: slides, classnotes, and answer keys
3. incorporates lessons from my research experience:
   robotic multi-agent, social networks, power grids
4. now v1.3
   v2.0 will expand nonlinear coverage

316 pages
164 exercises, 205 pages solution manual
4.6K downloads Jun 2016 - Nov 2019
35 instructors in 16 countries

Conclusions

Robotic problems

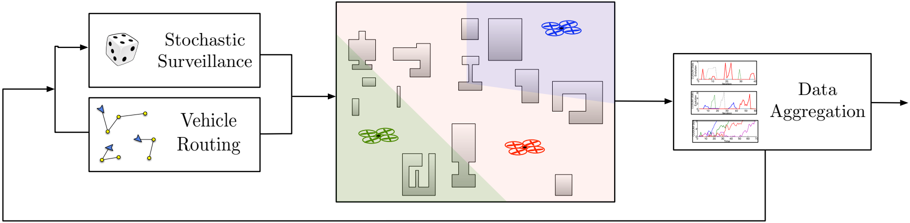

Potential collaborations on bioinspired coordination

1. optimal exploration-based territory partitioning
2. heuristics for routing through locations
3. unpredictability in animal motion