From Robotic Routing and Balancing to Stochastic Surveillance

Francesco Bullo

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

System and Control Department
Indian Institute of Technology Bombay, Jan 4, 2017

Acknowledgments

Jorge Cortés*
Sonia Martínez*
Emilio Frazzoli*
Marco Pavone*
Paolo Frasca*

G. Notarstefano
Anurag Ganguli
Ketan Savla
Ruggero Carli*
Sara Susca

Stephen Smith
Shaunak Bopardikar
Karl Obermeyer
Joey Durham
Vaibhav Srivastava

Fabio Pasqualetti
Rush Patel
Pushkarini Agharkar
Jeff Peters
Mishel George


Rush Patel, Northrop Grumman


Pushkarini Agharkar, Secrete Corporation


Mishel George, UCSB


Vaibhav Srivastava, Michigan State


Andrea Carron, ETH Zurich

NSF AFOSR ARO ONR DOE

New text "Lectures on Robotic Planning and Kinematics"

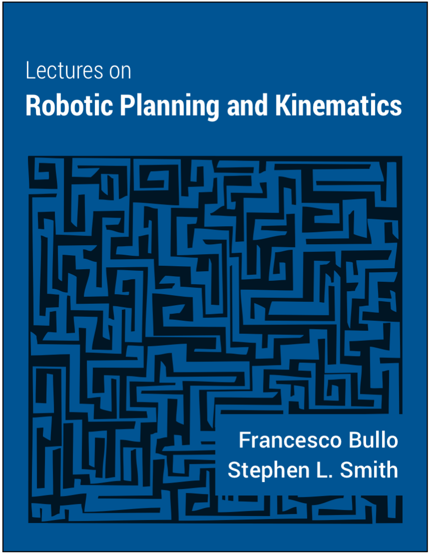

Lectures on Robotic Planning and Kinematics, ver .91
For students: free PDF for download
For instructors: slides and answer keys
`http://motion.me.ucsb.edu/book-lrpk/`

Robotic Planning:

1. Sensor-based planning
2. Motion planning via decomposition and search
3. Configuration spaces
4. Sampling and collision detetion
5. Motion planning via sampling

Robotic Kinematics:

6. Intro to kinematics
7. Rotation matrices
8. Displacement matrices and inverse kinematics
9. Linear and angular velocities

New text “Lectures on Network Systems”

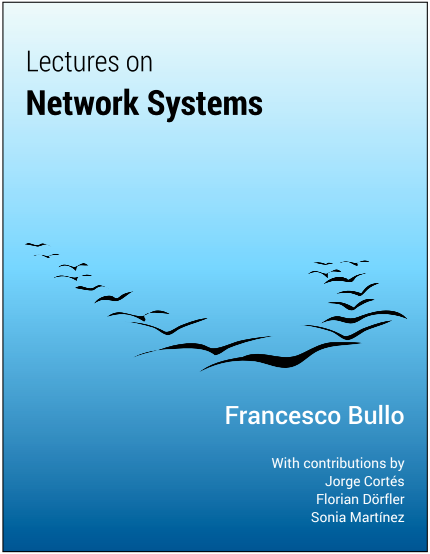

Lectures on Network Systems, ver .85
For students: free PDF for download
For instructors: slides and answer keys
`http://motion.me.ucsb.edu/book-lns/`

Linear Systems:

1. motivating examples from social, sensor and compartmental networks,
2. matrix and graph theory, with an emphasis on Perron–Frobenius theory and algebraic graph theory,
3. averaging algorithms in discrete and continuous time, described by static and time-varying matrices, and
4. positive and compartmental systems, described by Metzler matrices.

Nonlinear Systems:

5. formation control problems for robotic networks,
6. coupled oscillators, with an emphasis on the Kuramoto model and models of power networks, and
7. virus propagation models, including lumped and network models as well as stochastic and deterministic models, and
8. population dynamic models in multi-species systems.

Stochastic surveillance and dynamic routing

Design efficient vehicle control strategies to

1. search unpredictably
2. detect anomalies quickly
3. provide service to customers at known locations
4. perform load balancing among vehicles

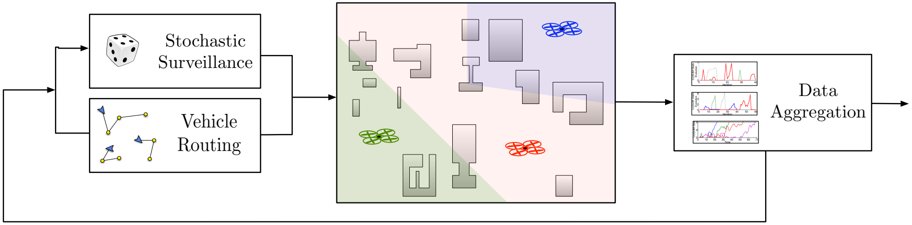

Outline

1. **vehicle routing**
2. load balancing and partitioning
3. stochastic surveillance

AeroVironment Inc, "Raven"
unmanned aerial vehicle

iRobot Inc, "PackBot"
unmanned ground vehicle

Vehicle routing in dynamic stochastic environments

- customers appear sequentially randomly space/time
- robotic network *knows* locations and provides service
- Goal: distributed adaptive algos, delay vs throughput

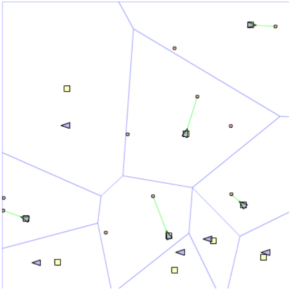

F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proceedings of the IEEE*, 99(9):1482–1504, 2011.

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`

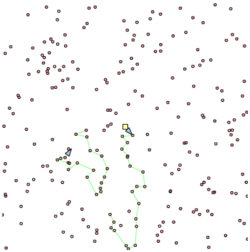

- shortest path is NP-hard, but effective heuristics available
- delay is optimal in light traffic
- delay is constant-factor optimal in high traffic

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

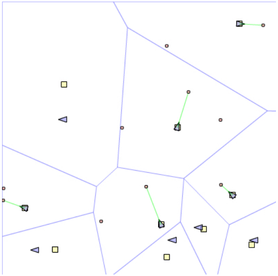

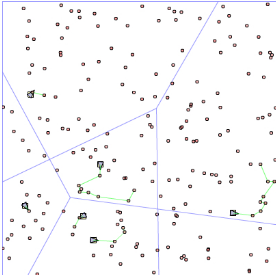

Outline

1. vehicle routing
2. **load balancing and partitioning**
3. stochastic surveillance

AeroVironment Inc, "Raven"
unmanned aerial vehicle

iRobot Inc, "PackBot"
unmanned ground vehicle

Load balancing via partitioning

ANALYSIS of cooperative distributed behaviors

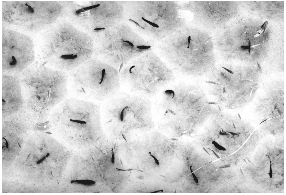

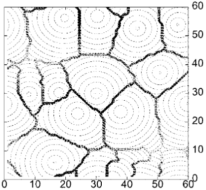

DESIGN of performance metrics

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?

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

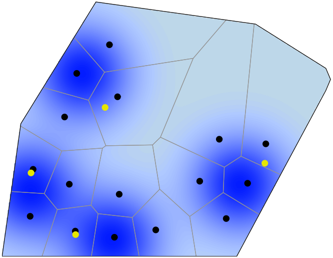

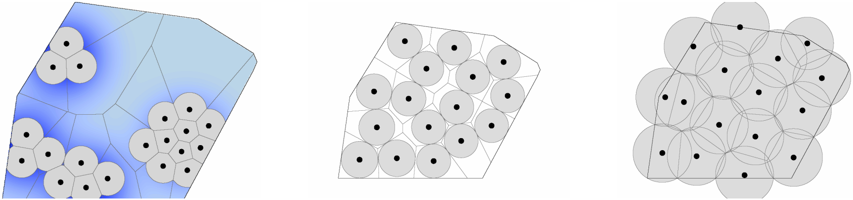

Area-center   Incenter   Circumcenter

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

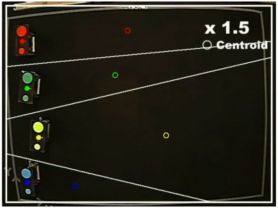

T. Hatanaka, M. Fujita, TokyoTech

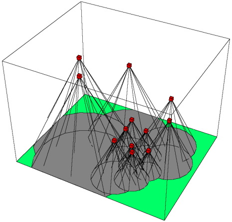

3D coverage

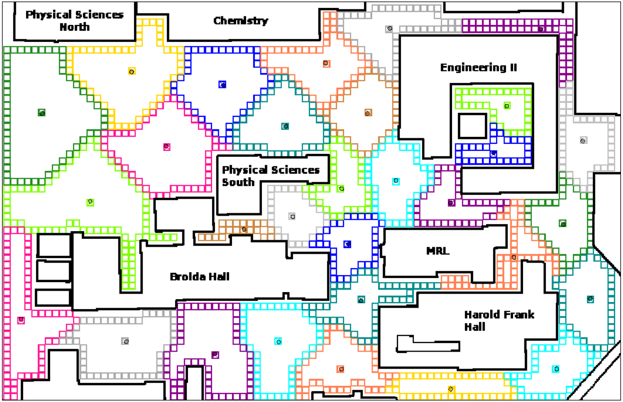

Outline

1. vehicle routing
2. load balancing and partitioning
3. **stochastic surveillance**

AeroVironment Inc, "Raven"
unmanned aerial vehicle

iRobot Inc, "PackBot"
unmanned ground vehicle

Stochastic surveillance: Motivating Example

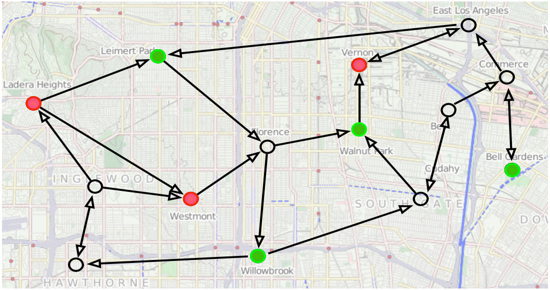

- **stationary anomalies / moving intruders**
- **pursuers**
- goal: when do they meet? how to optimize meeting time?
- assumption: both Markovian

Outline of Stochastic Surveillance

1. **Analysis: pursuer/evader meeting times**
2. Analysis/convex design:
   hitting time for reversible transitions with distances
3. Analysis/convex design: quickest detection
4. Analysis/SQP design: multiple pursuers

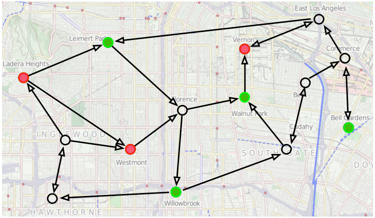

Single pursuer/evader expected first meeting time

$\mathcal{M}_{ij}(P_{\mathrm{p}}, P_{\mathrm{e}}) = \mathbb{E}$[first time pursuer starting @i meets evader starting @j]

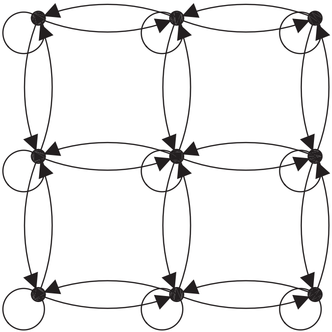

evader chain P_{e}

?

optimal
pursuer chain P_{p}?

Objective

Given evader chain P_{e}

$$\min_{\text{pursuer chain } P_{\mathrm{p}}} \mathbb{E}[\mathcal{M}_{ij}(P_{\mathrm{p}}, P_{\mathrm{e}})]$$

Walks in the Kronecker graph

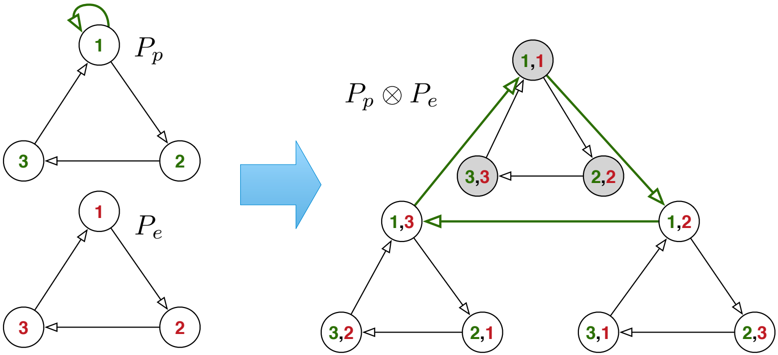

Thm 1: equivalent statements

(i) all $\mathcal{M}_{ij}$ are finite

(ii) from every (pursuer node, evader node) in Kronecker graph there is a walk to a common node

The Kronecker product of matrices $A \in \mathbb{R}^{n\times m}$ and $B \in \mathbb{R}^{q\times r}$ is an $nq \times mr$ matrix given by

$$A \otimes B = \begin{bmatrix} a_{1,1}B & \dots & a_{1,m}B \\ \vdots & \ddots & \vdots \\ a_{n,1}B & \ddots & a_{n,m}B \end{bmatrix}$$

Properties of the Kronecker product

Given the matrices A, B, C and D of appropriate dimensions,

(i) $(A \otimes B)$ is bilinear in A and B,

(ii) $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$,

(iii) $(B^\top \otimes A)\,\mathrm{vec}(C) = \mathrm{vec}(ACB)$,

where $\mathrm{vec}(C)$ is the vectorization of C by stacking of the columns

Walks in the Kronecker graph — or lack thereof

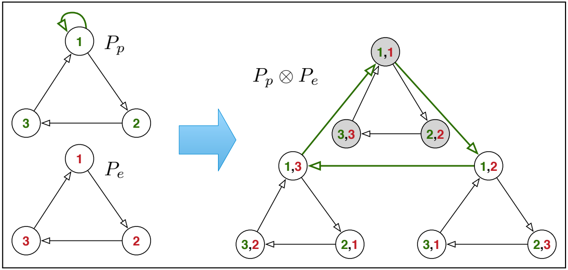

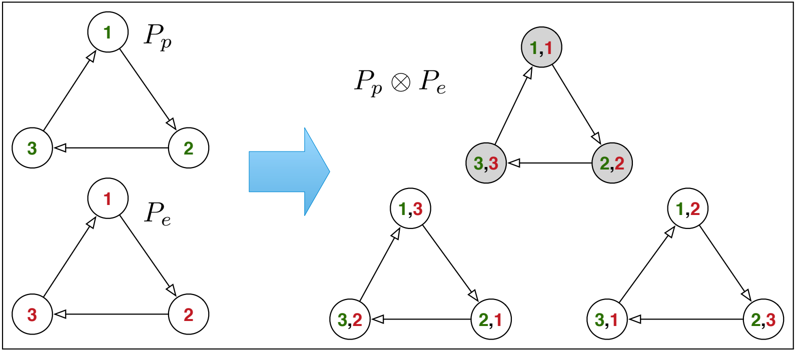

Sets of matrix pairs with all finite meeting times

$\mathcal{P}_{\text{one-ergodic}}$ = one of P_{p}, P_{e} is ergodic
$\mathcal{P}_{\text{SA-overlap}}$ = P_{p}, P_{e} have single absorbing classes, overlapping
$\mathcal{P}_{\text{MA-overlap}}$ = P_{p}, P_{e} have multiple absorbing classes, pairwise overlapping
$\mathcal{P}_{\text{finite}}$ = P_{p}, P_{e} satisfy conditions in Thm 1

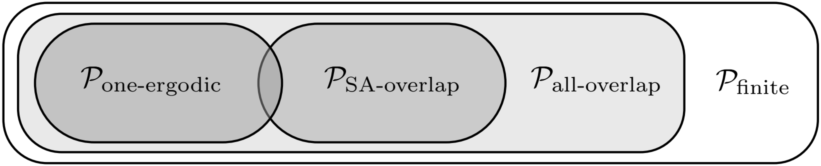

Thm 1: all $\mathcal{M}_{ij}$ are finite $\iff$ from every (pursuer node, evader node) there is a walk to a common node in Kronecker graph

Thm 2: Certain sets of matrix pairs have all $\mathcal{M}_{ij}$ finite

Closed-form expression

If all meeting times are finite,

$$\mathcal{M}_{ij}(P_{\mathrm{p}}, P_{\mathrm{e}}) = (\mathbb{e}_i \otimes \mathbb{e}_j)^\top (I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}})\, E)^{-1} \mathbb{1}_{n^2}$$

If P_{p}, P_{e} have stationary distributions π_{p}, π_{e} (i.e., $\mathcal{P}_{\text{SA-overlap}}$), then

$$\mathbb{E}[\mathcal{M}_{ij}(P_{\mathrm{p}}, P_{\mathrm{e}})] = (\pi_{\mathrm{p}} \otimes \pi_{\mathrm{e}})^\top (I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}})\, E)^{-1} \mathbb{1}_{n^2}$$

Thm 1: all $\mathcal{M}_{ij}$ are finite $\iff$ from every (pursuer node, evader node) there is a walk to a common node in Kronecker graph

Thm 2: Certain sets of pairs of matrices imply finiteness of all $\mathcal{M}_{ij}$

Thm 3: Closed-form expression for $\mathcal{M}_{ij}$ (matrix dimension n^2)

M. George, R. Patel and F. Bullo. The Meeting Time of Multiple Random Walks. *SIAM Journal on Matrix Analysis and Applications*, Submitted, Oct 2016.

Outline of Stochastic Surveillance

1. Analysis: pursuer/evader meeting times
2. **Analysis/convex design: hitting time for reversible transitions with distances**
3. Analysis/convex design: quickest detection
4. Analysis/SQP design: multiple pursuers

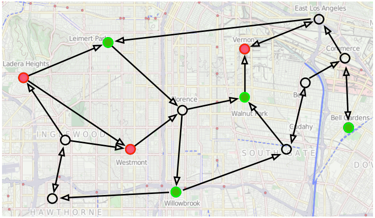

Meeting time for stationary evaders: Hitting time

Given a stationary evader with distribution π_{e},

$$\min_{P_{\mathrm{p}} \text{ with stationary } \pi_{\mathrm{p}}} \mathcal{H}(P_{\mathrm{p}}, \pi_{\mathrm{e}}) = \min_{P_{\mathrm{p}}} \mathbb{E}[\text{first time pursuer meets evader}]$$

The meeting time for a pursuer chain P_{p} and a stationary evader with distribution π_{e} is called the hitting time

Thm 4: Hitting time for stationary evader

$$\begin{aligned}\mathcal{H}(P_{\mathrm{p}}, \pi_{\mathrm{e}}) &= \lim_{P_{\mathrm{e}} \to I_n} \mathbb{E}[\mathcal{M}_{ij}(P_{\mathrm{p}}, P_{\mathrm{e}})] \\ &= (\pi_{\mathrm{p}} \otimes \pi_{\mathrm{e}})^\top \Big((I_{n^2} - P_{\mathrm{p}} \otimes I_n) \operatorname{diag}\big(\operatorname{vec}(I_n)\big) \Big)^{-1} \mathbb{1}_{n^2}\end{aligned}$$

SDP for hitting time of reversible chains

Thm 5: Convexity of hitting time

Given stationary distribution π_{e}, edge set E,

$$\text{minimize } \mathcal{H}(P_{\mathrm{p}}, \pi_{\mathrm{e}})$$

subject to

1. P_{p} is transition matrix with $\pi_{\mathrm{p}} = \pi_{\mathrm{e}}$
2. P_{p} is consistent with E
3. P_{p} is reversible

can be formulated as an SDP.

R. Patel, P. Agharkar and F. Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant. *IEEE Transactions on Automatic Control*, 60(12):3156-3157, 2015.

Application: Intruder detection

Intruders appear at random locations and persist for given life-time

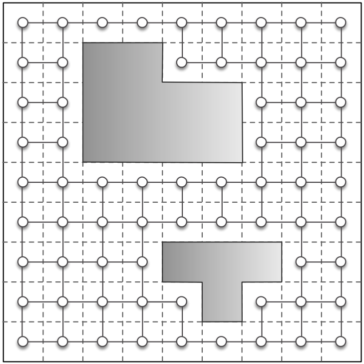

% Captures

Algorithm	**Mean**	**StdDev**	$\mathcal{H}$
Min $\mathcal{H}$	32.4%	2.1	207
FMMC*	29.8%	1.9	236
MHMC**	31.1%	2.1	231

*Fastest mixing Markov chain
**Metropolis-Hastings Markov chain

Weighted hitting time

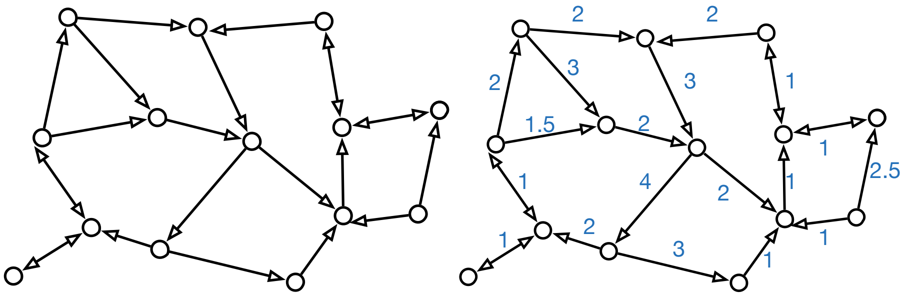

Hitting time can be computed for graphs with travel time matrix W

Thm 6: Weighted hitting time

$$\mathcal{H}_w(P_\mathrm{p}, \pi_\mathrm{e}, W) = (\pi_\mathrm{p} \otimes \pi_\mathrm{e})^\top \Big((I_{n^2} - P_\mathrm{p} \otimes I_n) \operatorname{diag}\big(\operatorname{vec}(I_n)\big) \Big)^{-1} \cdot \operatorname{vec}\big((P_\mathrm{p} \circ W) \mathbb{1}_n \mathbb{1}_n^T \big)$$

SDP for weighted hitting time of reversible chains

Thm 7: Convexity of weighted hitting time

Given stationary distribution π_{e}, edge set E with weights W,

$$\text{minimize } \mathcal{H}_w(P_{\mathrm{p}}, \pi_{\mathrm{e}}, W)$$

subject to

1. P_{p} is transition matrix with $\pi_{\mathrm{p}} = \pi_{\mathrm{e}}$
2. P_{p} is consistent with E
3. P_{p} is reversible

can be formulated as an SDP.

R. Patel, P. Agharkar and F. Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant. *IEEE Transactions on Automatic Control*, 60(12):3156-3157, 2015.

Minimum weighted hitting time: Results

Intruders appear at random locations and persist for given life-time

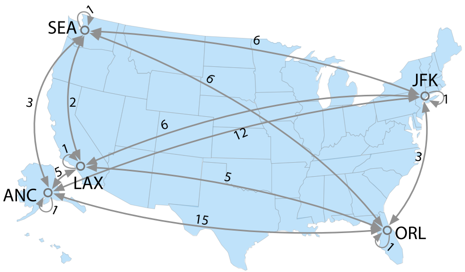

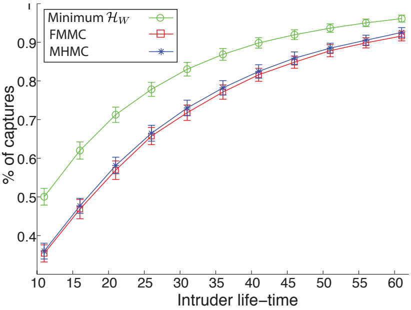

Outline of Stochastic Surveillance

1. Analysis: pursuer/evader meeting times
2. Analysis/convex design:
   hitting time for reversible transitions with distances
3. **Analysis/convex design: quickest detection**
4. Analysis/SQP design: multiple pursuers

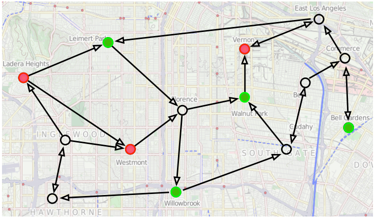

Quickest detection of anomalies

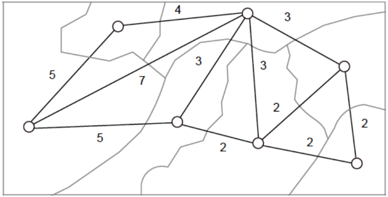

$f_k^0 \rightarrow$ nominal distribution

$f_k^1 \rightarrow$ anomalous distribution

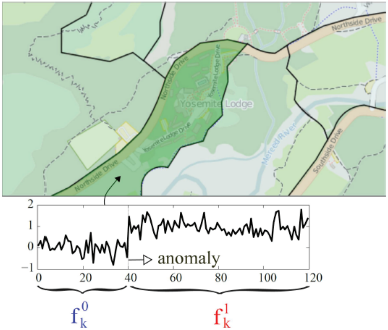

Given nominal/anomalous pdfs at locations,
travel times between nodes W,
spatial distribution of anomalies π_e,
compute and minimize detection time wrt monitoring agent chain P_a

$$\delta_{\text{avg}}(P_a, W, \pi_e, (f_k^0, f_k^1)) = \mathbb{E}[\text{average detection delay}]$$

Quickest detection: Single region

CUSUM algorithm

Given threshold η

1. set statistic $\Lambda = 0$
2. collect an observation y
3. update statistic
   $\Lambda = \max\left\{0, \Lambda + \log \frac{f_k^1(y)}{f_k^0(y)}\right\}$
4. if $\Lambda > \eta$: declare anomaly
5. else go to step 2.

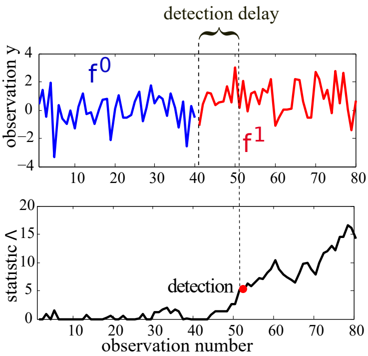

$\mathcal{D}_k$ = Kullback-Liebler divergence at location k
s_k = expected number of samples before detection at location k

$$s_k = \frac{e^{-\eta} + \eta - 1}{\mathcal{D}_k}$$

Quickest detection: Multiple regions = SDP

Ensemble CUSUM algorithm

1. Agent moves according to transition chain P_{a}, travel time matrix W
2. conducts N parallel CUSUM algorithms for each region k

Thm 8: Detection delay of ensemble CUSUM algorithm

detection delay at region k: $\delta_k = \sum_{i=1}^{n} (\pi_{\mathrm{a}})_i \mathcal{M}_{ik} + (s_k - 1)\mathcal{M}_{kk}$

Quickest detection: Multiple regions

Given priority of regions w_k, $\delta_{\text{avg}} = \sum_{k=1}^{n} w_k \delta_k$

Thm 9: Convexity of average detection delay

Given stationary distribution π_{e}, edge set E, travel matrix W and priority vector w

$$\min_{P_{\text{a}}} \ \delta_{\text{avg}}(P_{\text{a}}, \pi_{\text{e}}, W, w)$$

subject to

1. P_{a} is transition matrix with $\pi_{\text{a}} = \pi_{\text{e}}$
2. P_{a} is consistent with E
3. P_{a} is reversible

can be formulated as an SDP.

P. Agharkar and F. Bullo. Quickest detection over robotic roadmaps. *IEEE Transactions on Robotics*, 32(1):252-259, 2016.

Quickest detection: Example

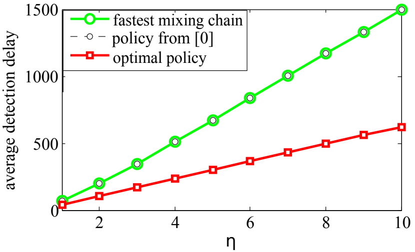

η = global CUSUM threshold

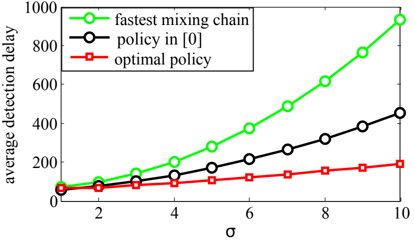

σ = variation in Kullback-Liebler divergence

V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *The International Journal of Robotics Research*, 32(12):1438-1458, 2013.

P. Agharkar and F. Bullo. Quickest detection over robotic roadmaps. *IEEE Transactions on Robotics*, 32(1):252-259, 2016.

Outline of Stochastic Surveillance

1. Analysis: pursuer/evader meeting times
2. Analysis/convex design:
   hitting time for reversible transitions with distances
3. Analysis/convex design: quickest detection
4. **Analysis/SQP design: multiple pursuers**

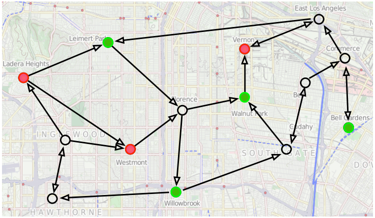

Multiple evaders and pursuers

Thm 10: Expected first meeting time among N pursuers and M evaders

$$\begin{aligned}\mathbb{E}[\mathcal{M}_{i_1\cdots i_N,\ j_1\cdots j_M}(P_p^{(1)},\ldots,P_p^{(N)},P_e^{(1)},\ldots,P_e^{(M)})]\\ = (\pi_p^{(1)}\otimes\cdots\otimes\pi_p^{(N)}\otimes\pi_e^{(1)}\otimes\cdots\otimes\pi_e^{(M)})\\ \cdot(I_{n^{N+M}}-(P_p^{(1)}\otimes\ldots\otimes P_P^{(N)}\otimes P_e^{(1)}\otimes\cdots\otimes P_e^{(M)})E_{(N,M)})^{-1}\mathbb{1}_{n^{N+M}}\end{aligned}$$

For N pursuers with single stationary evader, the group hitting time is

$$\begin{aligned}\mathcal{H}_N(P_p^{(1)},\ldots,P_p^{(N)},\pi_e) = (\pi_p^{(1)}\otimes\cdots\otimes\pi_p^{(N)}\otimes\pi_e)\\ \cdot(I_{n^{N+1}}-(P_p^{(1)}\otimes\cdots\otimes P_p^{(N)}\otimes I_n)E_{(N,1)})^{-1}\mathbb{1}_{n^{N+1}}\end{aligned}$$

Group hitting time

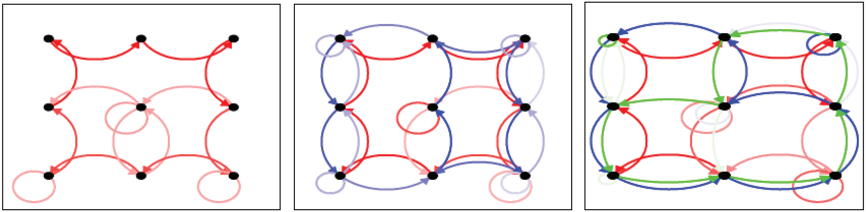

Random Walker(s)	Red	Blue	Green	H_N
One	6.8	–	–	6.8
Two	7.7	10.5	–	4.1
Three	7.0	15.9	16.9	2.9

- Optimizing transition matrices is nonlinear program, hence SQP
- Curse of dimensionality: system of equations $\mathcal{O}(n^{N+1})$ to be solved

R. Patel, A. Carron, and F. Bullo. The hitting time of multiple random walks. *SIAM Journal on Matrix Analysis and Applications*, 37(3):933-954, 2016.

Group hitting time with partitioning

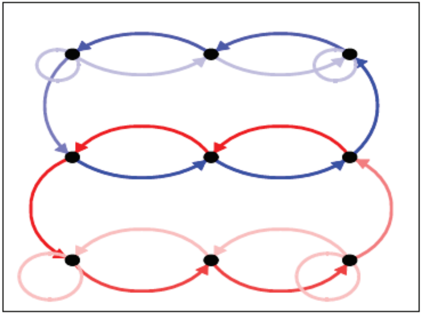

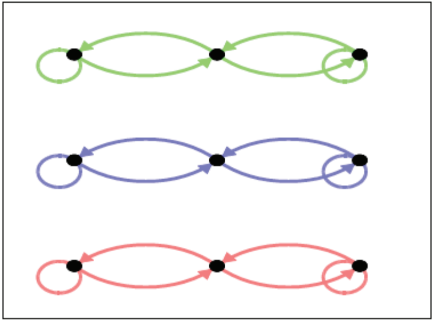

Random Walker(s)	H_N w/ Overlap	H_N w/ Partitioning
Two	4.1	3.6
Three	3.7	2.9

- Partitioning can lead to better group hitting times
- Complexity of problem can be reduced $\mathcal{O}(Nn_1n_2\dots n_N)$ where $n_1, n_2, \dots, n_N$ are size of partitions

Marostica case study

4 agents, 42 vertices and 56 edges: 2 minutes on 2.7Ghz, KNITRO solver

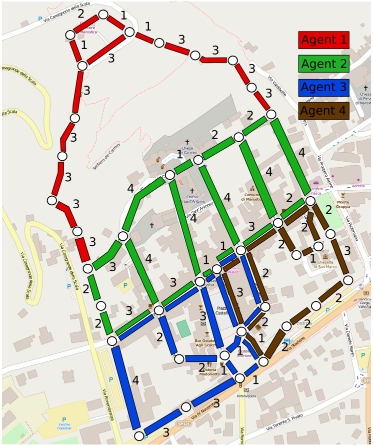

Marostica with travel distances and with pre-fixed partition

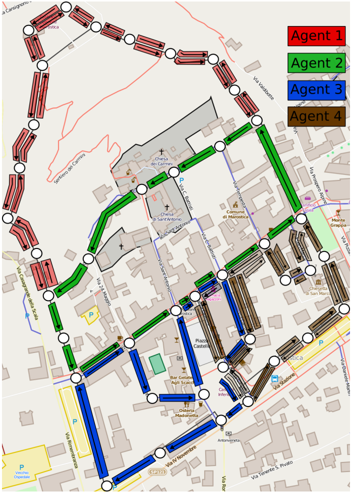

Optimized transitions ≈ edge transparency

A. Carron, R. Patel, and F. Bullo. Hitting time for doubly-weighted graphs with application to robotic surveillance. *European Control Conference*, Aalborg, Denmark, Jun 2016.

Publications

(1) V. Srivastava, F. Pasqualetti, and F. Bullo. **Stochastic surveillance strategies for spatial quickest detection.**
International Journal of Robotics Research, 32(12):1438–1458, 2013.

(2) R. Patel, P. Agharkar, and F. Bullo.
Robotic surveillance and Markov chains with minimal weighted Kemeny constant.
IEEE Transactions on Automatic Control, 60(12):3156–3157, 2015.

(3) P. Agharkar and F. Bullo.
Quickest detection over robotic roadmaps.
IEEE Transactions on Robotics, 32(1):252–259, 2016.

(4) R. Patel, A. Carron, and F. Bullo.
The hitting time of multiple random walks.
SIAM Journal on Matrix Analysis and Applications, 37(3):933–954, 2016.

(5) M. George, R. Patel, and F. Bullo.
The Meeting Time of Multiple Random Walks.
SIAM Journal on Matrix Analysis and Applications, Submitted, Oct 2016.

Conclusions

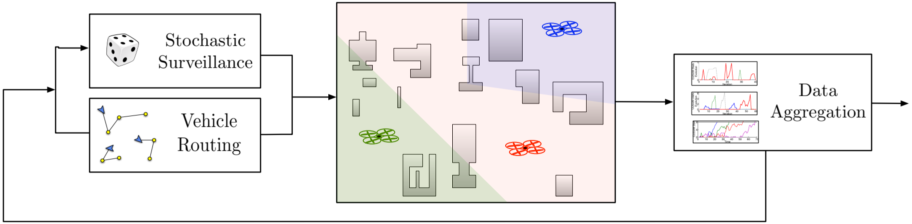

Summary

1. vehicle routing & environment partitioning
2. stochastic surveillance: analysis and design

Ongoing work on stochastic surveillance

1. multi-pursuer/evader: computational complexity
    1. optimize partitioning/covering for scalability
2. fast unpredicatable searchers
    1. optimizing lifted chains
    2. optimize canonical pairs and robotic interpretations