Distributed intelligence in mobile multi-agent networks

Usman A. Khan
Department of Electrical and Computer Engineering
Tufts University

IEEE Signal Processing Society Winter School
Concordia University
November 3, 2016

Tufts UNIVERSITY | School of Engineering
KTH ROYAL INSTITUTE OF TECHNOLOGY

Who am I

- **Usman A. Khan**
  - **Assistant Professor, Tufts**
- **Postdoc**
  - **U-Penn**
- **Education**
  - **PhD, Carnegie Mellon**
  - **MS, UW-Madison**
  - **BS, Pakistan**

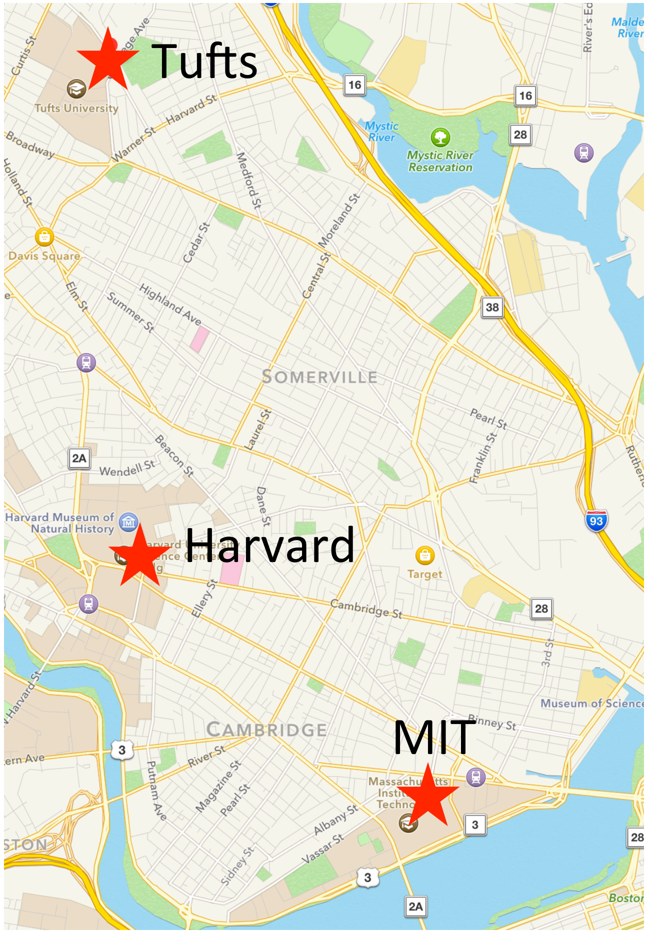

My Research Lab: Projects and demos

Research Team

PhD Students

Current
- Fakhteh Saadatniaki, Sep. 2014 to date
- Sam Safavi, Jan. 2013 to date
- Chengiuang Xi, Sep. 2012 to date

Alumni
- Mohammadreza Doostmohammadian, graduated May 2015

MS Students

Current
- Christopher Sacca, Sep. 2015 to date
- Dong Park, Sep. 2015 to date

Alumni
- Alexander Henry, graduated Aug. 2015, Adaptive methods for robotic path planning
- Michael Tran, graduated Aug. 2015, Distributed target tracking in a sensor network
- Dibeyandu Das, graduated Aug. 2015, Consensus with non-participating agents
- Anders Simpson-Wolf, graduated Dec. 2014, Privacy and differentially private methods
- Luke Grymek, graduated Aug. 2013, Coverage and surveillance with autonomous agents
- Gerald Solimini, graduated May 2012, Distributed path planning algorithms for UAVs
- Syed S. Akbar, graduated Dec. 2011, Object recognition on AR-Drone (UAV) platform
- Qiong Wu (Applied Mathematics), Summer 2012, Stochastic modeling of wind turbines

Undergraduates

Current
- Ryan Kortvelesy, ECE Freshman, Fall 2015 onwards
- Anuththari Gamage, ECE Sophomore, Fall 2015 onwards
- Syed M. Bukhari, ECE Junior, Summer 2015 onwards
- Terrence Tufuor, ECE Junior, Summer 2015 onwards

Alumni
- Dong Park, ECE Junior, Apr. 2014 to May 2015
- Cody Chen, ECE Junior, Jan. 2014 to May 2015
- Oghenefego Ahia, ECE Freshman, Summer 2014
- Pratik Chatrath, ECE Junior, visiting student from SVNIT, Gujarat, India, Summer 2014
- John Kelly, CS Senior, Jan. 2013 to May 2014
- Cornell Wilson, ECE Junior, May 2013 to Sep. 2013, Aerial robot navigation
- Karman Chu, ECE Junior, Jan. 2013 to Aug. 2013, Robotic networks
- Josh Pfosi, ECE Sophomore, Jan. 2013 to Aug. 2013, Robotic networks

- Senior Design Project, 2014-2015
  - D. Park, C. Chen, and B. Zhang: GPS based self-navigating UAV
- Senior Design Project, 2012-2013
  - M. Tran, R. Singh, A. Simpson-Wolf, and S. Stanneiwicz: Low-power aerial surveillance for engineering infrastructures
    Best project in ECE
- Senior Design Projects, 2011-2012
  - E. Zheng, F. Shaukat, T. Perkins, and Y. Garcia: GPS and compass integration on AR-Drone platform
    3rd Place Winners, IBM/IEEE Smarter Planet Challenge: Student Projects Changing the World
  - Robust hardware and software redesign for AR-Drone platform

- ME undergrads (graduated), Spring 2012: N. Stone, C. N. Bargar, J. Arena, W. Langford
- Kevin Morrissey (graduated), Distributed control of wind-farms, Spring 2012
- Hassan Oukacha, GPS-based autonomous navigation, Summer 2012
- Michael Tran, Decentralized target tracking, May 2011 to May 2013
- Jesse Weeks, Operational attributes and obstacle avoidance, Summer 2011 to August 2012
- Tyler Heck (Junior), Feasibility of object recognition algorithm, Summer 2011

My Research Lab: Projects and demos

Inspecting leaks in NASA's lunar habitat

Aerial Formation Flying

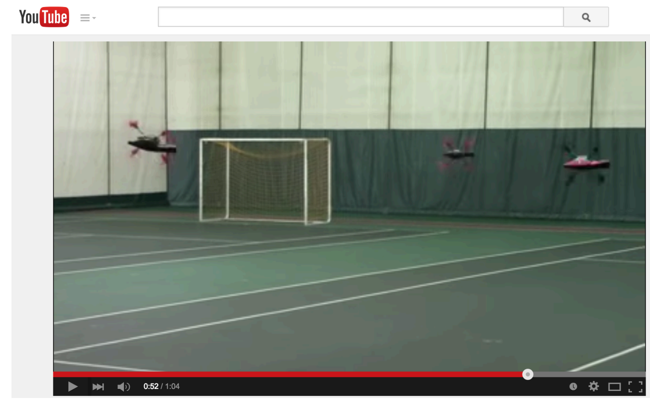

Inference in Social Networks

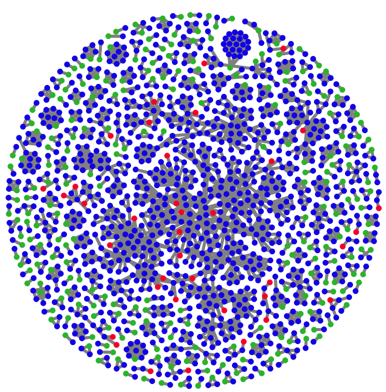

SHM over a campus footbridge

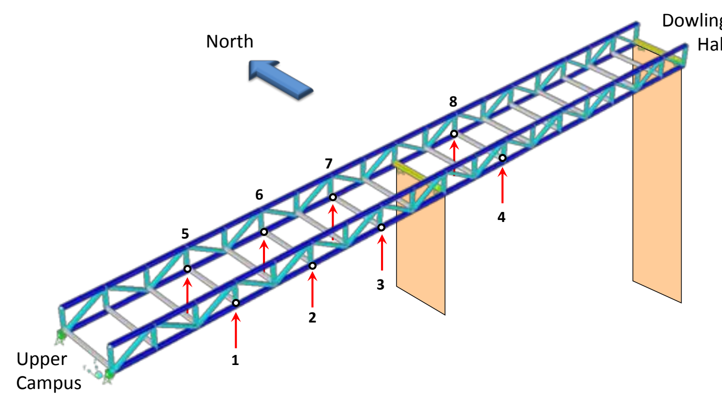

My Research Lab: Theory

Reza (2011-15):
Graph-theoretic estimation

Best paper, Journal cover

Xi (2012-16):
Optimization over directed graphs

Sam (2013-17):
Fusion in non-deterministic graphs

2 Best papers

Fakhteh (2014-):
Distributed estimation cont…d

My Research: In depth

- Distributed Intelligence in multi-agent systems
  - Estimation, optimization, and control over graphs (networks)

- Mobile → Dynamic
- Heterogeneous → Directed
- Autonomous → Non-deterministic

- Applications:
  - Cyber-physical systems, IoTs, Big Data
  - Aerial SHM, Power grid, Personal exposome
  - Indoor navigation (this talk)

How do we think about intelligence?

Intelligence: Conventional notion

- **Individual**
- **Droids**
- **Cyborg**
- **Cyborg with skin grafts**
  - **Terminator, T-800**
- **Todays robots/UAVs**

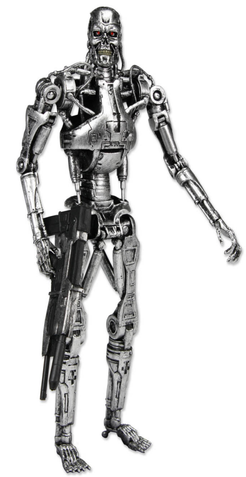

Intelligence: Conventional notion

- System
- SkyNet Central Core
  - Guarded by T1000000
- Cylon BaseStar
  - BSG
- Other examples
  - SCADA
  - Fusion centers in WSNs
  - Dispatch centers

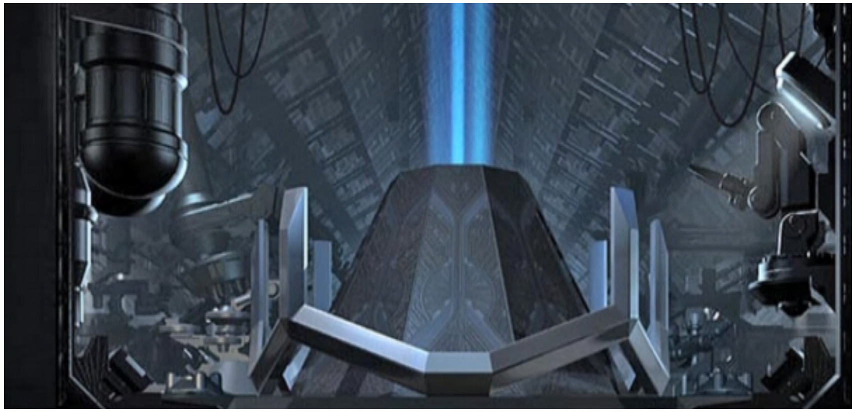

What can go wrong?

- Single point of failure
- Infiltration
- Cyber/terrorist threats
- Blackouts

How can we think about intelligence?

- **Individual level**
- **T1000**

- (something I am interested in)

How can we think about intelligence?

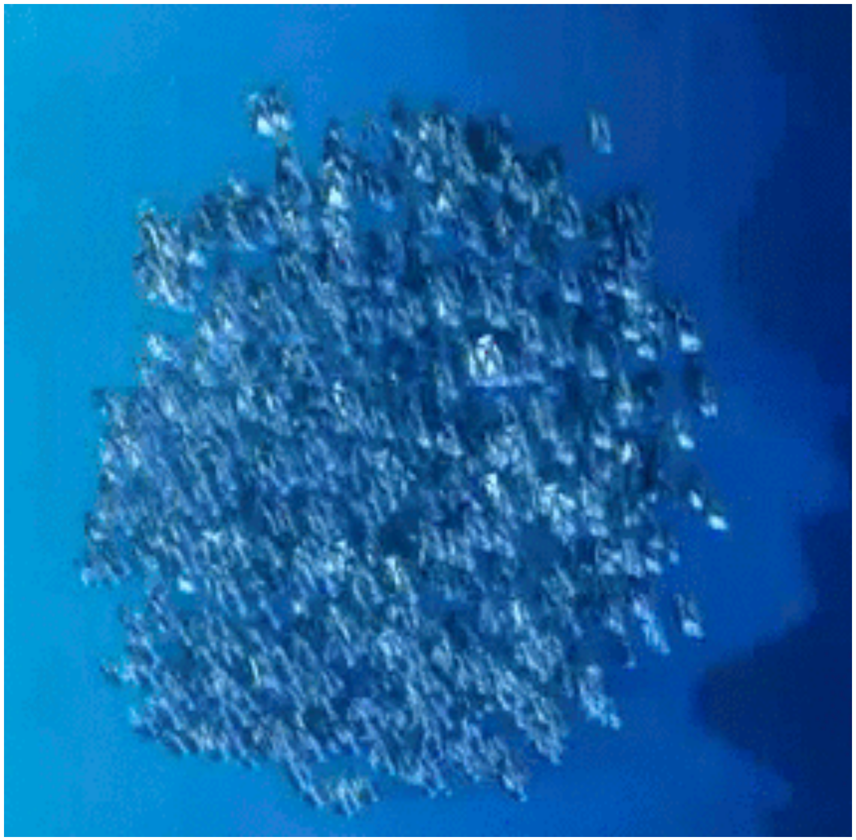

- **System level**
  - **Driving directions in *Finding Nemo***
  - ***(I wonder if Finding Dori is particularly harder than Finding Nemo)***
- **(Majority of) this talk**
  - **GPS-free navigation in mobile and autonomous systems**

An Example Project: Individual Intelligence

- Indoor GPS-free Aerial Navigation

- https://www.youtube.com/watch?v=0AuF_Xj_Xms
- General Problem with multiple robots

Distributed sensor localization

Distributed sensor localization

- **Localize M sensors with unknown locations in $\mathbf{R}^m$**
- **Sensors can only communicate in a neighborhood**
- **Only local distances in the neighborhood are available**
- **What is the minimal number of known locations required?**
  - **Called anchors, QR codes**
- **Where do we place them?**

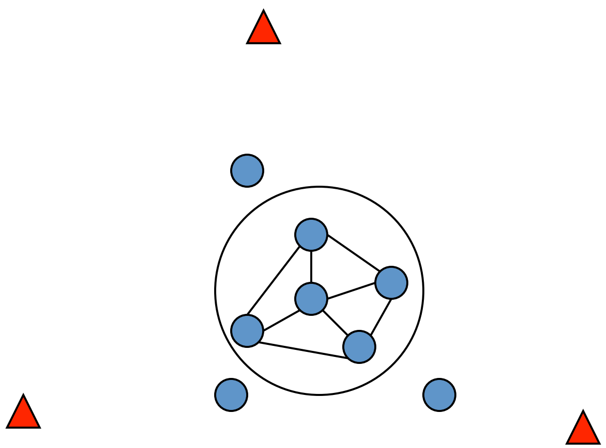

$m = 2$, plane

Sensor localization

- **Traditional (non-linear) multilateration scheme**
  - **(only distances to known locations are given)**
  - **Nonlinear**
  - **Coupled in coordinates**

- **Minimal anchors:** $m+1$

- **Placement: arbitrary**

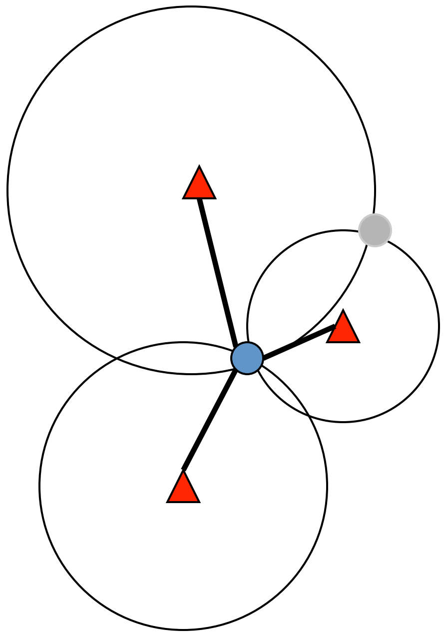

Distributed sensor localization

- **Can we iteratively build on the nonlinear approach?**
  - **each sensor iteratively updates its location**
  - **several sensors may not be able to talk to any anchor**
  - **no sensor may be able to talk to all of the $m+1$ anchors**

- **No, the iterations do not converge in general**

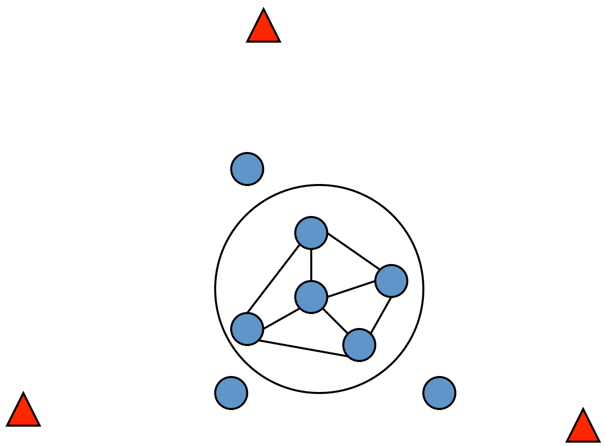

$m = 2$, plane

Distributed sensor localization

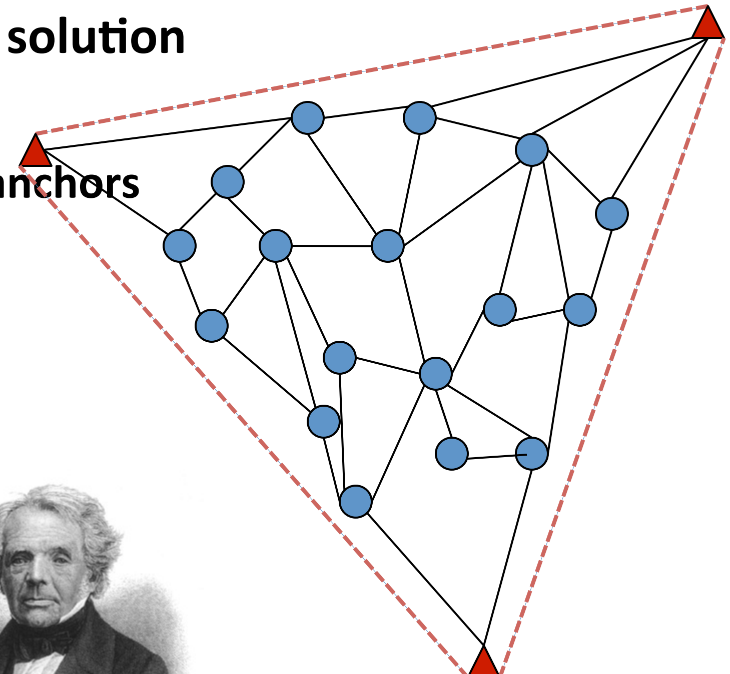

- The non-linear problem has a linear iterative solution
- Convexity arguments
  - Sensors lie in the convex hull of at least $m+1$ anchors
  - This condition can be relaxed
- Barycentric coordinates
  - August Ferdinand Möbius (1790 – 1868)

- Cayley-Menger determinants
  - Joseph-Louis Lagrange (1736 – 1813)
  - Arthur Cayley (1821 – 1895)
  - Karl Menger (1902 – 1985)

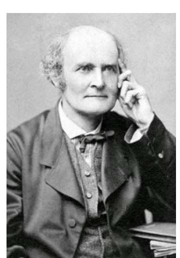

Barycentric coordinates: Main idea

- Linear-convex combination on a line (m=1)
- Need m+1 = 2 anchors

$$\mathbf{x} = \frac{5}{5+3}\cdot 1 + \frac{3}{5+3}\cdot 9$$

- Unknown location is within the convex hull of knowns
- Barycentric coordinates: Sum to 1 and are positive

- The idea is extendible to arbitrary dimensions
- What should replace distances?

Barycentric coordinates: Definition

- **Linear representation of coordinates**

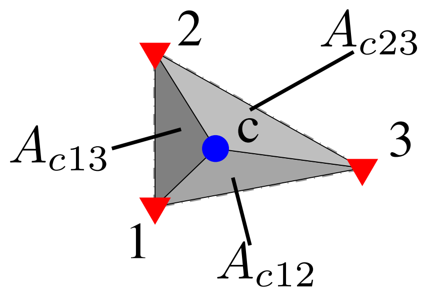

$$a_{c1} = \frac{A_{c23}}{A_{123}}$$

$$\mathbf{c} = a_{c1}\mathbf{c}_1 + a_{c2}\mathbf{c}_2 + a_{c3}\mathbf{c}_3$$

- **Decoupled in coordinates**
  - **Unique and between 0—1 (if within the convex hull)**
  - **Sum to 1**
- **How do we compute the areas or generalized volumes in $\mathbf{R}^m$?**

Cayley-Menger determinant

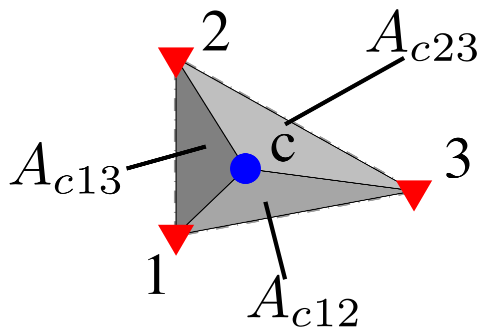

$$a_{c1} = \frac{A_{c23}}{A_{123}}$$

- **Computes the generalized volumes of an m-simplex**
  - **Computation from local distances alone (six distances)**

$$A^2_{\Theta_l} = \frac{1}{s_{m+1}} \begin{vmatrix} 0 & \mathbf{1}^T_{m+1} \\ \mathbf{1}_{m+1} & \mathbf{Y} \end{vmatrix} \qquad s_m = \frac{2^m (m!)^2}{(-1)^{m+1}}$$

  - **Y is a 3x3 matrix containing pairwise squared distances**

Distributed Sensor Localization

- **Recipe:**
- **Each sensor finds three neighbors such that it lies in their convex hull**
- **How?**
- **Finds BC from CM determinants and local distances**
- **Update its coordinates using the linear equation and coordinates from (appropriate) neighbors**

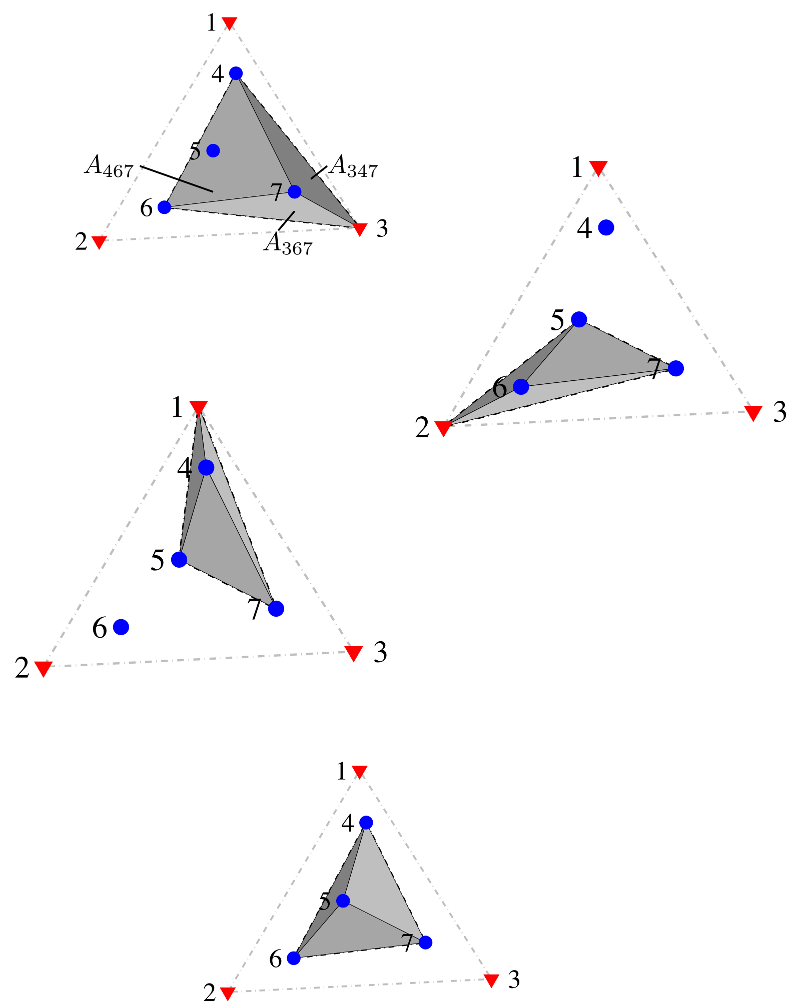

Tufts UNIVERSITY School of Engineering

KTH VETENSKAP OCH KONST ROYAL INSTITUTE OF TECHNOLOGY

Triangulation

- **Test to find a triangulation set**
- **Convex hull inclusion test based on the following observation**

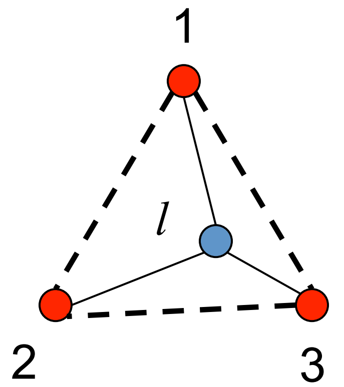

$$A_{l12} + A_{l13} + A_{l23} = A_{123}$$

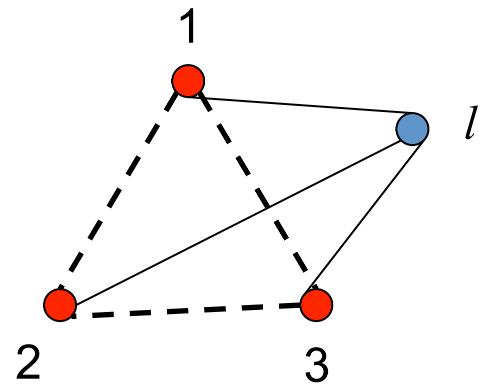

$$A_{l12} + A_{l13} + A_{l23} > A_{123}$$

- **The test becomes**

- **Node *l* is inside if the sum = total**
- **Node *l* is outside if the sum > total**

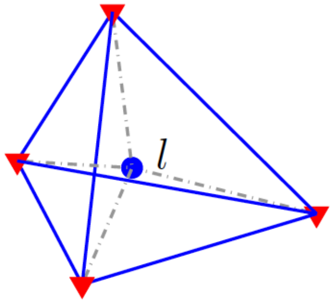

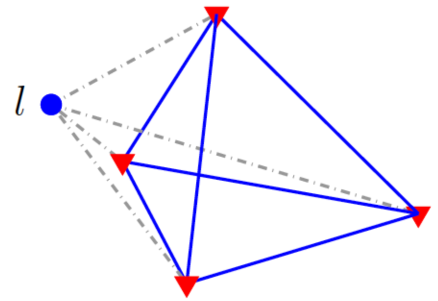

Distributed sensor localization

- Let $\mathbf{u}_k$ and $\mathbf{x}_l$ be the coordinates of kth anchor and lth sensor, respectively

- We have the following update:

$$\begin{aligned}
\text{Anchors:} \quad \mathbf{u}_k(t+1) &= \mathbf{u}_k(t) = \mathbf{u}_k^*, \\
\text{Sensors:} \quad \mathbf{x}_l(t+1) &= \sum_{j \in \Theta_l} \upsilon_{lj} \mathbf{x}_j(t), \\
&= \sum_{j \in \Omega_{\Theta_l}} \upsilon_{lj} \mathbf{x}_j(t) + \sum_{k \in \kappa_{\Theta_l}} \upsilon_{lk} \mathbf{u}_k^*
\end{aligned}$$

BC

Triangulation set

Convergence Analysis

- **The update converges to exact sensor locations regardless of the initial conditions**
- **Under strongly-connected sensor to sensor graph and when each anchor can communicate to at least one different sensor**
- **The proof sketch is as follows**
- **(each row sums to 1 and has +ve elements)**

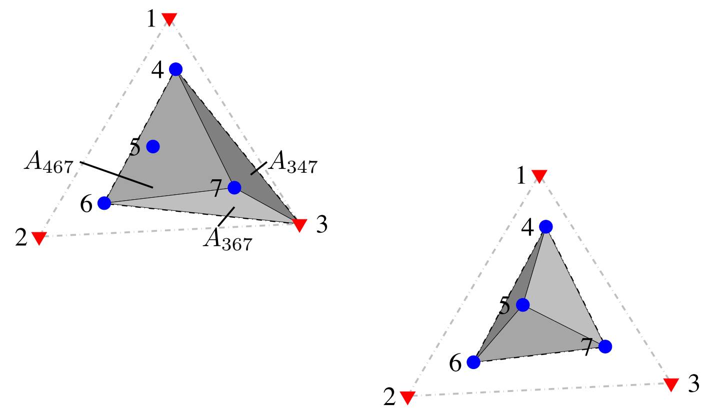

$$\begin{bmatrix} \mathbf{u}^{k+1} \\ \mathbf{x}^{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ \mathbf{B} & \mathbf{P} \end{bmatrix} \begin{bmatrix} \mathbf{u}^{k} \\ \mathbf{x}^{k} \end{bmatrix}$$

Convergence Analysis

- Comment on Absorbing Markov chains
- The iteration matrix is a stochastic matrix

(m+1) x (m+1) identity matrix

the anchors do not update

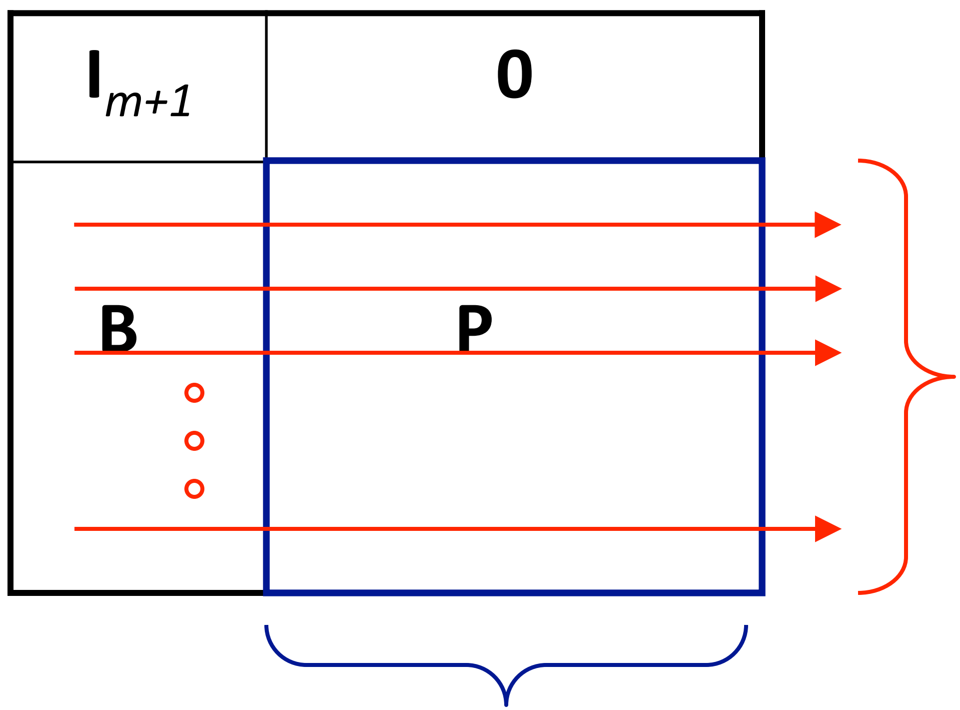

M x (M + m+1)

each row sums to 1

has exactly (m+1) non zeros in [0, 1]

we can show that P is Hurwitz

$\rho(\mathbf{P}) < 1,$

Convergence Analysis

$$\begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ \mathbf{B} & \mathbf{P} \end{bmatrix} \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ \mathbf{B} & \mathbf{P} \end{bmatrix} \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ \mathbf{B} & \mathbf{P} \end{bmatrix} \cdots$$

$$= \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ \Sigma\, \mathbf{P}^{k}\mathbf{B} & \mathbf{P}^{k} \end{bmatrix} \longrightarrow \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ (\mathbf{I}-\mathbf{P})^{-1}\mathbf{B} & \mathbf{0} \end{bmatrix}$$

Convergence Analysis

$$\begin{bmatrix} \mathbf{u}^{\infty} \\ \mathbf{x}^{\infty} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ (\mathbf{I}-\mathbf{P})^{-1}\mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}^{0} \\ \mathbf{x}^{0} \end{bmatrix}$$

- **Finally, we can show that $(\mathbf{I}-\mathbf{P})^{-1}\mathbf{B}\mathbf{u}^0$ are the exact sensor locations**

Localization – Simulations

- N=7 **node network in** 2**-d plane**
- M= 4 **sensors,** $K = m+1 = 3$ **anchors**

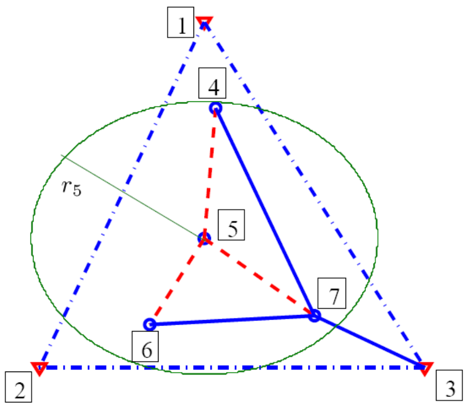

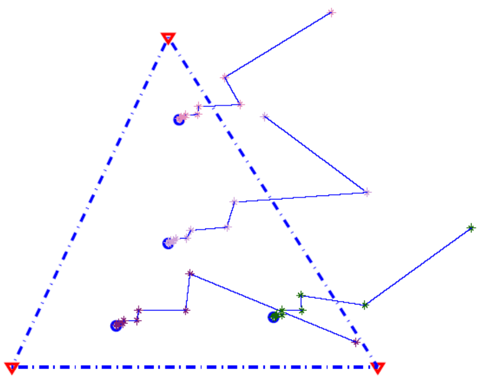

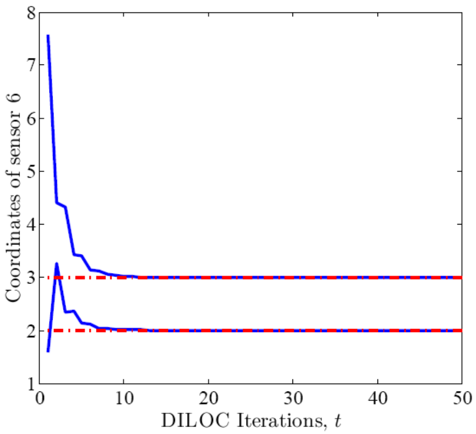

Simulations

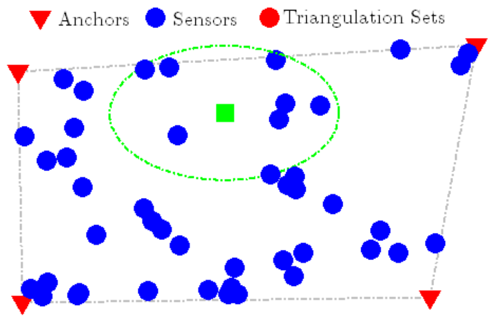

Distributed Sensor Localization, DILOC
Dynamic Topology and Convex Hulls

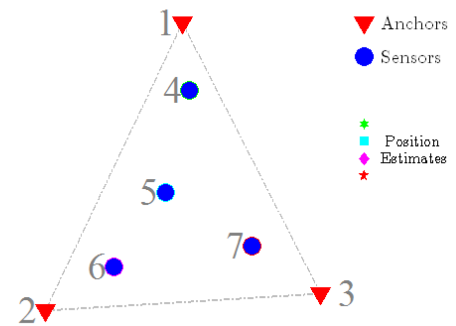

Distributed Sensor Localization, DILOC

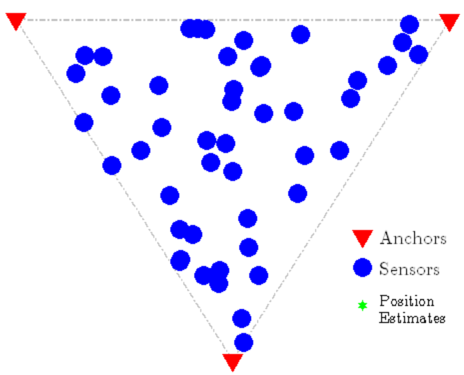

Distributed Sensor Localization, DILOC

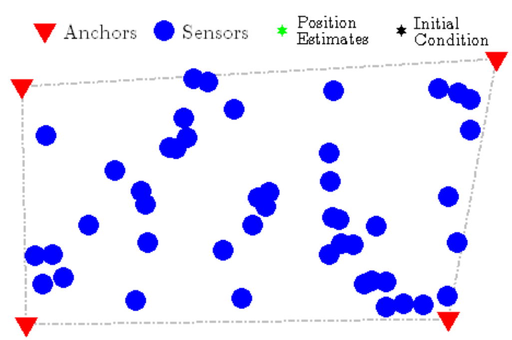

Distributed Sensor Localization, DILOC

Multimedia available here: http://ieeexplore.ieee.org/document/7180272/

Tufts UNIVERSITY | School of Engineering

KTH ROYAL INSTITUTE OF TECHNOLOGY

Robustness under imperfections

$$\mathbf{x}(t+1) = \left[1 - \alpha(t)\right]\mathbf{x}(t) + \alpha(t)\left\{\mathbf{E}_t \odot \mathbf{P}\left(\overline{\mathbf{d}}_t\right)\left[\mathbf{x}(t) + \mathbf{v}(t)\right] + \mathbf{E}_t \odot \mathbf{B}\left(\overline{\mathbf{d}}_t\right)\left[\mathbf{u} + \mathbf{v}(t)\right]\right\}$$

Graph randomness

Graph randomness

Barycentric matrices with imperfect distances

Stochastic approximation weights

Noise

Noise

Sensors state

Sensors state

Anchors state

- **The above algorithm converges a.s. to the exact sensor locations under some persistence conditions on the weights**
- Weights go to zero but not too fast

Distributed position tracking

- All of the nodes are mobile

- The agents are mobile and autonomous
- The graph is dynamic and non-deterministic

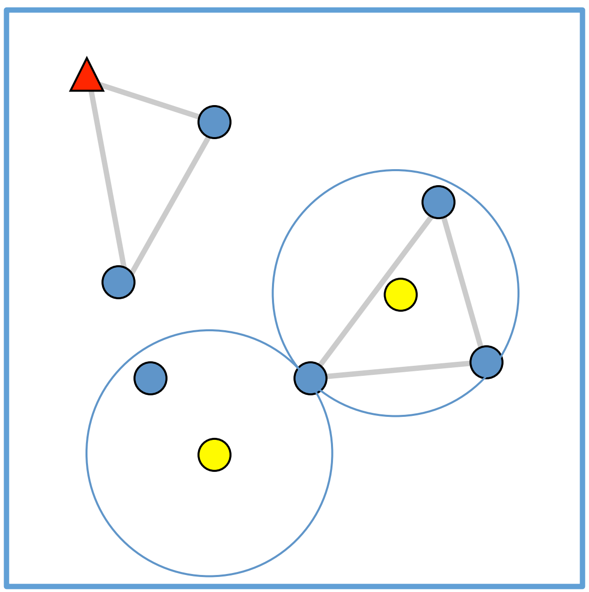

No update

All neighbors have unknown locations

At least one neighbor knows its location

Distributed position tracking

- New update:
- New Location = Old location + motion
- = Update with neighbors (if possible) + motion

- Recall the static update: Basically an LTI system

$$\begin{bmatrix} \mathbf{u}_{k+1} \\ \mathbf{x}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{m+1} & \mathbf{0} \\ \mathbf{B} & \mathbf{P} \end{bmatrix} \begin{bmatrix} \mathbf{u}_{k} \\ \mathbf{x}_{k} \end{bmatrix}$$

Distributed position tracking

- **Now we have an LTV system where the corresponding matrices are random**

$$\begin{bmatrix} \mathbf{u}_{k+1} \\ \mathbf{x}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{?} & \mathbf{0} \\ \mathbf{B}_{k} & \mathbf{P}_{k} \end{bmatrix} \begin{bmatrix} \mathbf{u}_{k} \\ \mathbf{x}_{k} \end{bmatrix} + \text{motion}$$

Distributed position tracking

- **New update:**

$$\begin{bmatrix} \mathbf{u}_{k+1} \\ \mathbf{x}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{?} & \mathbf{0} \\ \mathbf{B}_k & \mathbf{P}_k \end{bmatrix} \begin{bmatrix} \mathbf{u}_k \\ \mathbf{x}_k \end{bmatrix} + \text{motion}$$

- **Perfect locations follow:** $\mathbf{x}^*_{k+1} = \mathbf{P}_k \mathbf{x}^*_k + \mathbf{B}_k \mathbf{u}_k + \textbf{motion}$

- **Error:** $\mathbf{e}_{k+1} = \mathbf{P}_k \mathbf{e}_k$
- **where $\mathbf{P}_k$ is asymmetric, dynamic, non-deterministic**

Distributed position tracking

- Matrix form: e(k+1) = P(k) e(k)
- P_k randomly switches between no update, update with agents, update with anchors
- Consider the sequence: All agents update in, e.g., 6 steps
  - with the anchor, or
  - with an agent that has updated with the anchor

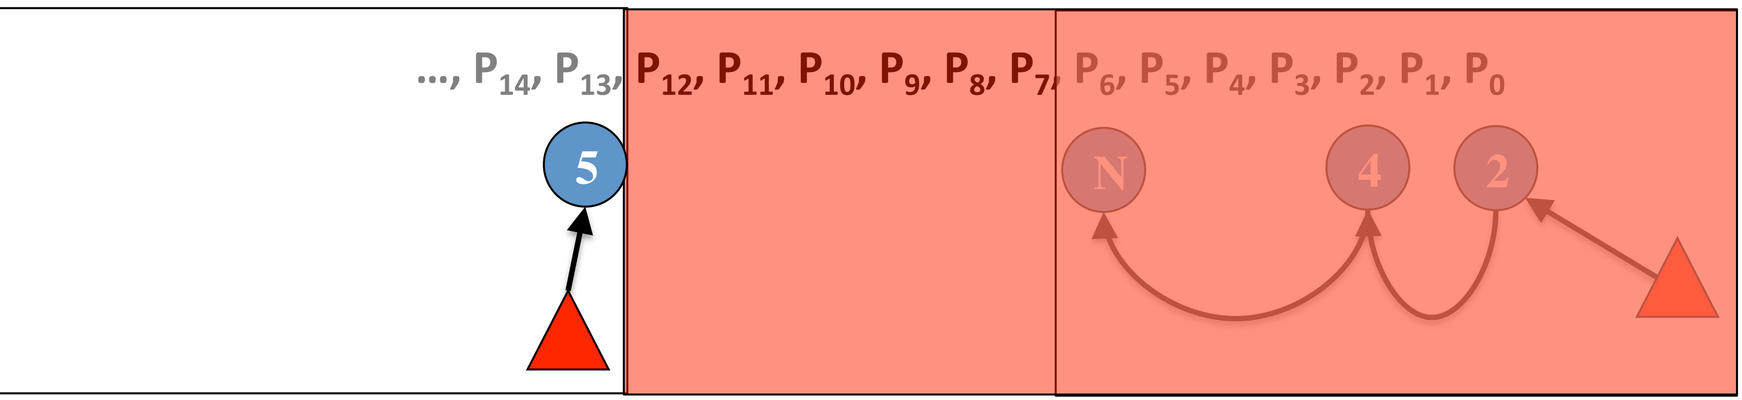

- The information cycle completes in 6 steps
- The next cycle starts at, e.g., time 13 (what happens between 6 and 13?)
- Slice: P12, P11, P10, P9, P8, P7, P6, P5, P4, P3, P2, P1, P0

Distributed position tracking

$\ldots, P_{14}, P_{13}, P_{12}, P_{11}, P_{10}, P_9, P_8, P_7, P_6, P_5, P_4, P_3, P_2, P_1, P_0$

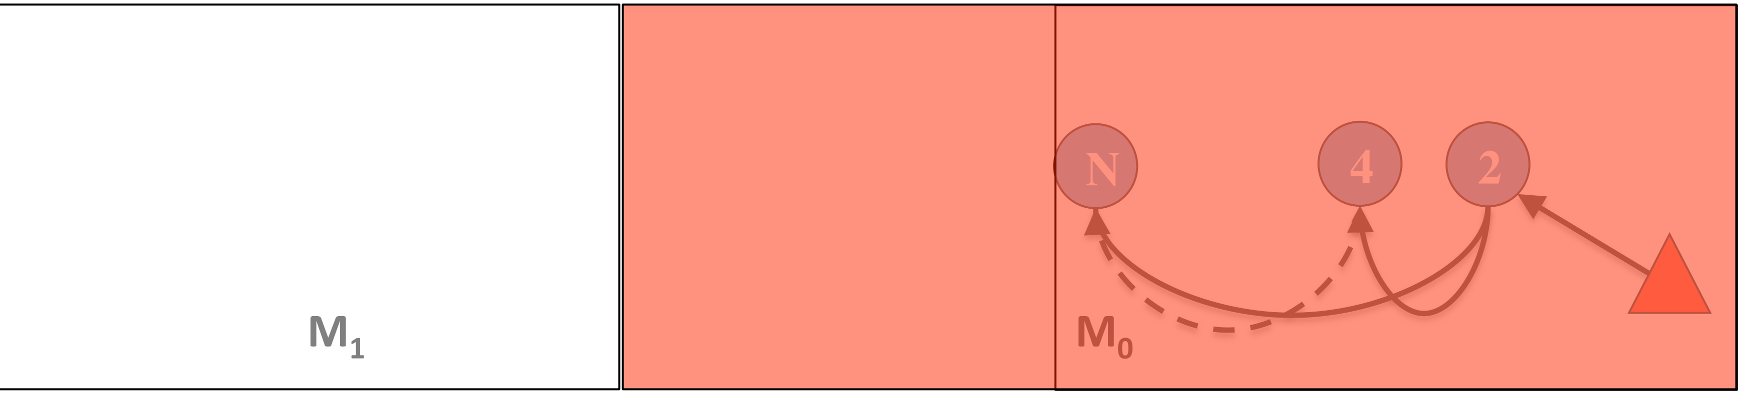

- Each slice contains
  - the system matrices such that one information cycle is completed, *and*
  - continues until the next cycle starts

- We have an alternate view: $e(\infty) = \ldots P_3 P_2 P_1 P_0 e(0) = \ldots M_3 M_2 M_1 M_0 e(0)$
- Instead of the product of system matrices, we study the product of slices

Tufts UNIVERSITY | School of Engineering

KTH VETENSKAP OCH KONST ROYAL INSTITUTE OF TECHNOLOGY

Distributed position tracking

- We have an alternate view: $e(\infty) = \ldots M_3\, M_2\, M_1\, M_0\, e(0)$
- Result 1: As the slice length goes to infinity, the two-norm goes to 1
- Result 2: If a slice completes in a finite time, then its two-norm is less than 1
- Result 3: If each slice completes in a fixed finite-time, then error goes to 0
  - (If an infinite subsequence completes in finite-time)
- Main Result 1: If the slice lengths grow at a certain rate, then the error goes to 0
- Main Result 2: The procedure works as long there is at least one anchor

Distributed position tracking: Experiments

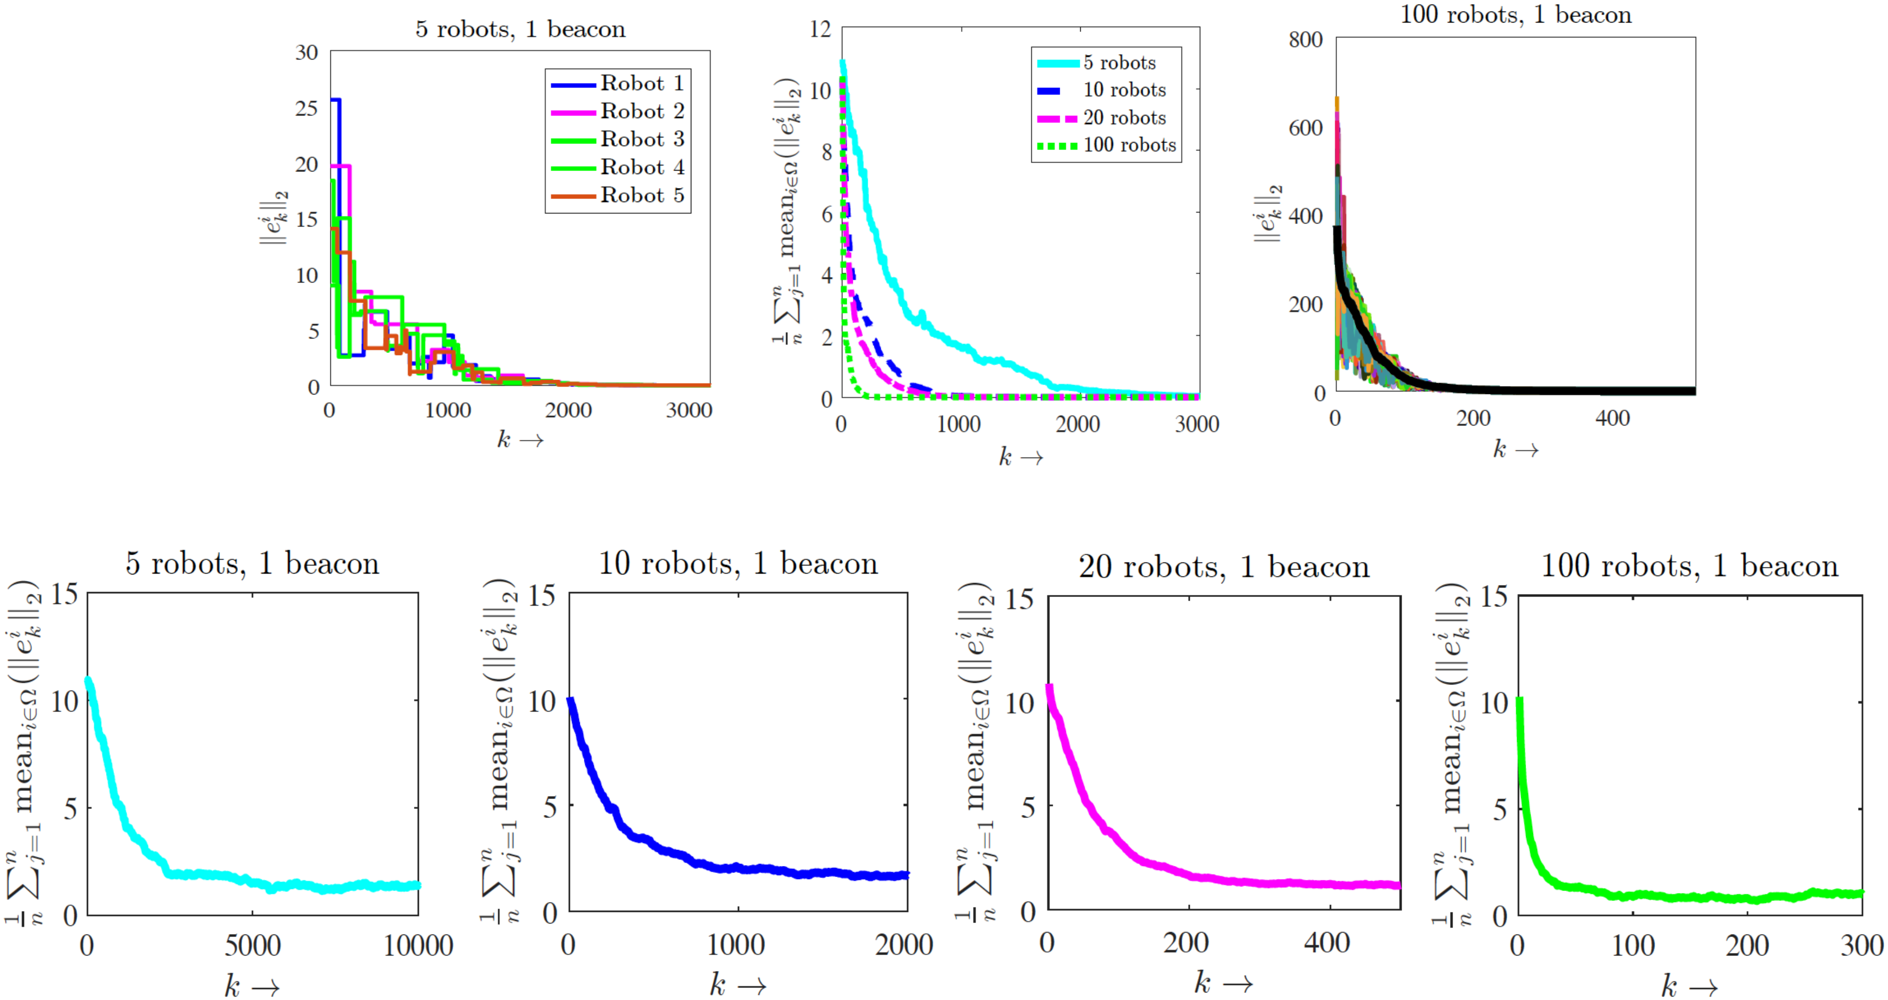

Trailer

SPARTN—Signal Processing and RoboTic Networks Lab at Tufts
https://www.youtube.com/watch?v=k6fOLbYj-5E

Tufts UNIVERSITY | School of Engineering

KTH VETENSKAP OCH KONST ROYAL INSTITUTE OF TECHNOLOGY

Structural Health Monitoring

- **Dowling Hall footbridge**

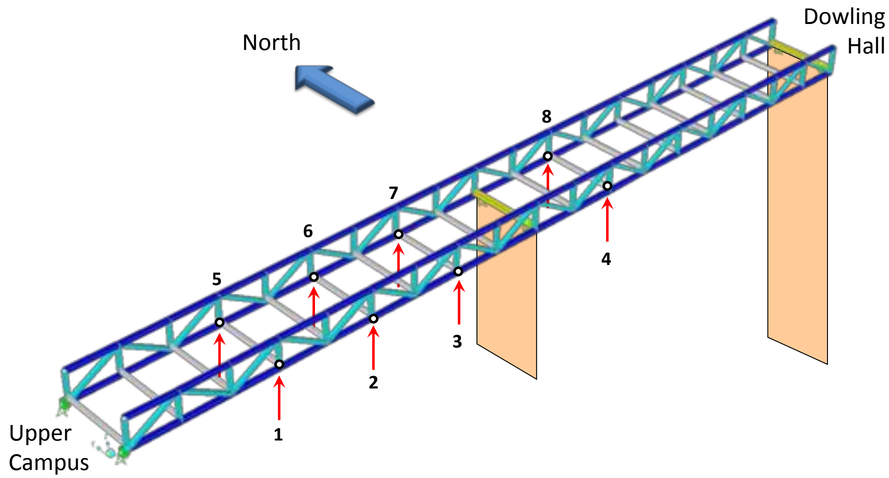

More Information

- **My webpage: http://www.eecs.tufts.edu/~khan/**

- **My email: khan@ece.tufts.edu**

- **My Lab's YouTube channel:**
  https://www.youtube.com/user/SPARTNatTufts/videos/