Global Estimation with Local Communication

Ittay Eyal, Idit Keidar, Stacy Patterson, and Raphi Rom
Technion – Israel Institute of Technology

Background

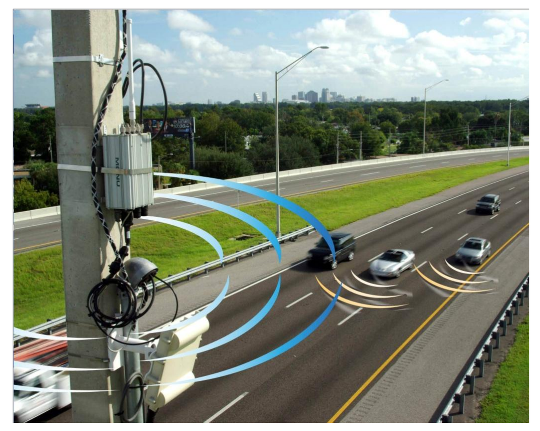

Modern vehicles have networking capabilities and many sensors. Can be used for environment monitoring.

Goal: Estimate a physical (continuous) phenomenon from samples over a vast area.

Collect all samples and analyze at a central site?

- Using 3G: Too expensive (money and power)
- Local network propagation: Bandwidth and latency too high.
- Centralized optimization: Computationally intractable.

Our approach: In-network estimation using **local** communication.

Solution Overview

Step 1: Divide region into sectors, with one node per sector.

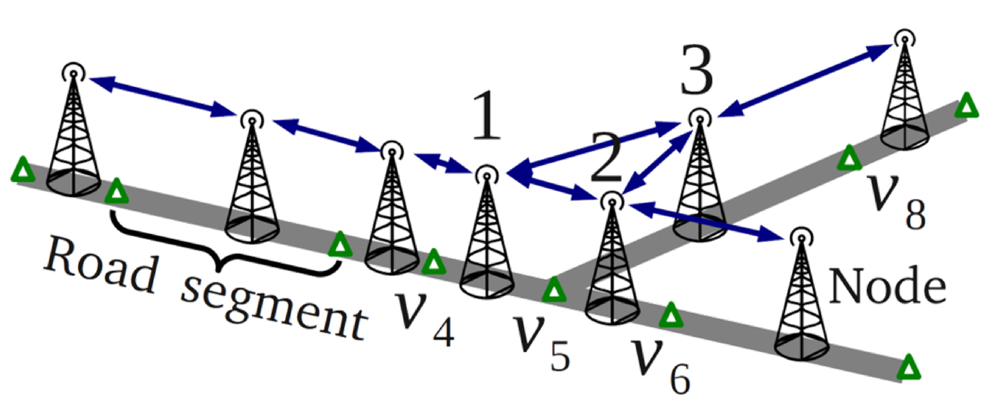

Step 2: Each node learns estimate in its region.

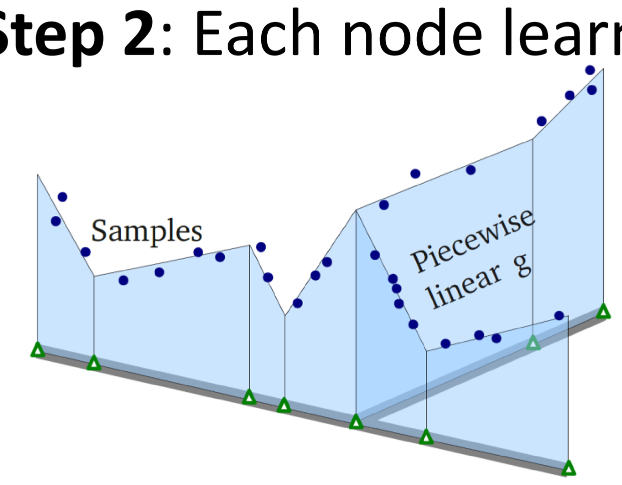

Phenomenon is continuous, therefore estimate should be continuous.

Each node generates estimate based on local samples and continuity constraints.

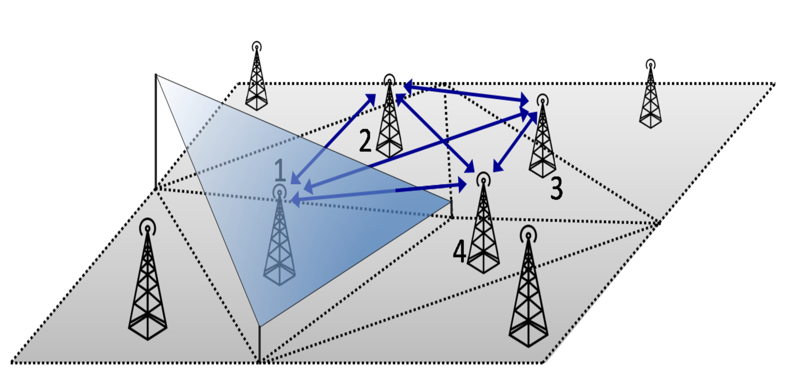

The result should be the globally optimal estimate!

Problem Formulation

Each node i learns $\theta_{u,i}, \theta_{v,i}$

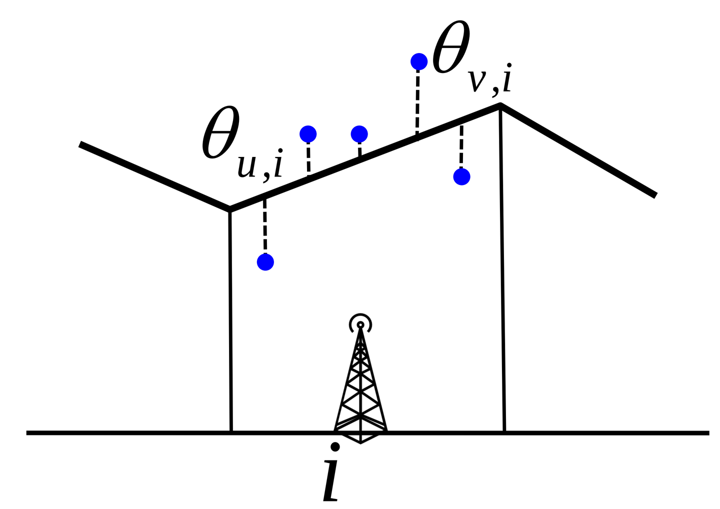

Local estimate error:

$C_i(Q_i, S_i) =$ Mean squared error in sector i

Global Optimization problem:

$$\begin{aligned} &\underset{\Theta}{\text{minimize}} && \sum_{i \in \mathcal{N}} C_i(\Theta_i; \mathcal{S}_i) \\ &\text{subject to} && \underbrace{\theta_{v,i} = \theta_{v,j},}_{\text{Continuity constraints.}} \quad \text{for } v \in \mathcal{V}, \quad i, j \in M(v), i \neq j \end{aligned}$$

Solving problem directly requires vertex-wise communication. ☹

The Dual Problem

Define the Lagrangian:

One Lagrange multiplier per constraint

$$\mathcal{L}(\Theta, \Lambda; \mathcal{S}) = \sum_{i \in \mathcal{N}} C_i(\Theta_i; \mathcal{S}_i) + \sum_{v \in \mathcal{V}} \sum_{\substack{i,j \in M(v) \\ i \neq j}} \lambda_{i,j}^{v} (\theta_{v,i} - \theta_{v,j})$$

Each Lagrange multiplier is shared by only two nodes. ☺

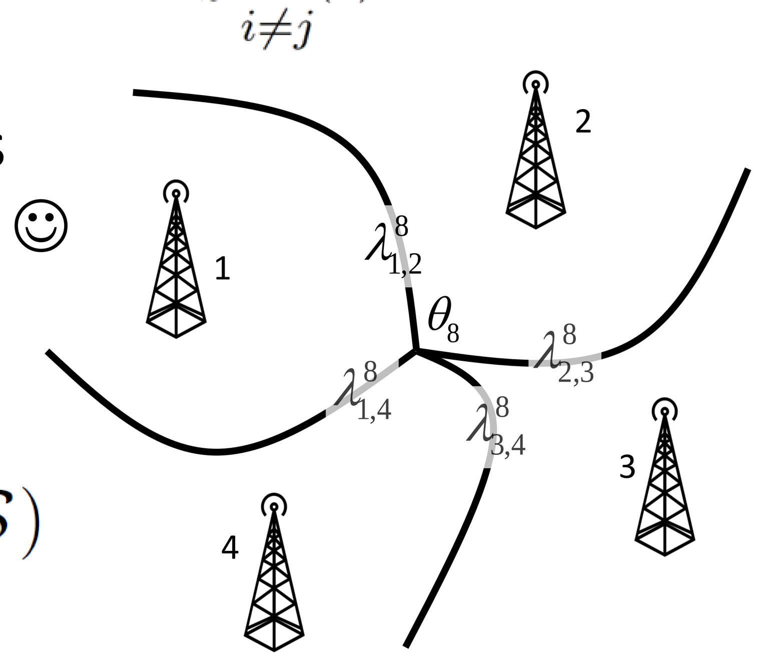

The dual function is:

$$q(\Lambda; \mathcal{S}) \triangleq \inf_{\Theta} \; \mathcal{L}(\Theta, \Lambda; \mathcal{S})$$

And the dual problem is:

$$\underset{\Lambda}{\text{maximize}} \quad q(\Lambda; \mathcal{S}) = \sum_{i \in \mathcal{N}} q_i(\Lambda_i; \mathcal{S}_i)$$

Solution to dual gives solution to primal:

$$\hat{\Theta} = \arg\min_{\Theta} \; \mathcal{L}(\Theta, \hat{\Lambda}; \mathcal{S})$$

Dual is unconstrained convex optimization problem.

Method of Coordinate Ascent

Solve an unconstrained convex optimization problem by repeatedly optimizing for one variable at a time.

Algorithm converges if order of updates follows **essentially cyclic rule**: every variable is updated at least once per bounded time period.

Our Distributed Algorithm

Solve the dual problem using novel, distributed implementation of coordinate ascent optimization.
Each Lagrange multiplier is a **shared variable**, of two neighbors.
Therefore, **only pairwise communication** per step ☺

Algorithm:

Each node maintains its current estimate: $\theta_{u,i}, \theta_{v,i}$

- Update one shared variable at a time.
- One leader per shared variable.
- Leader queues requests

Follower → **Leader**: Notify leader for update

Leader → **Follower**: Leader's update info — Follower: Update shared variable. Leader: Wait...

Follower → **Leader**: Follower's update info — Leader: Update shared variable and resume.

Algorithm properties:

- Simulates centralized coordinate ascent.
- With essentially cyclic variable update order after global stabilization time.
- Discontinuous global estimate until convergence.
- Accommodates sample updates.
- Accommodates churn.
- Quiescence at optimum.