Distributed Particle Filters: Stability Results and Graph-based Compression of Weighted Particle Clouds

Michael Rabbat

Joint with: Syamantak Datta Gupta and Mark Coates (McGill), and Stephane Blouin (DRDC)

McGill

Decentralized Tracking

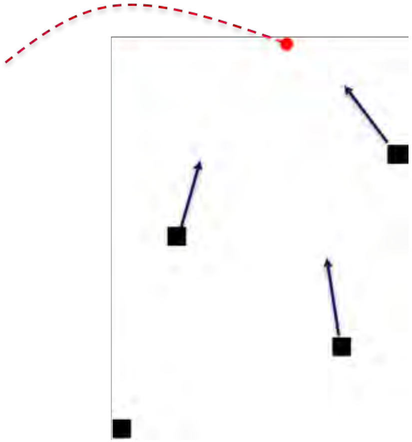

Multiple sensors
(e.g., bearings-only)

No central processing

Communicate, cooperate to
estimate object trajectory

Decentralized Particle Filtering (DPF)

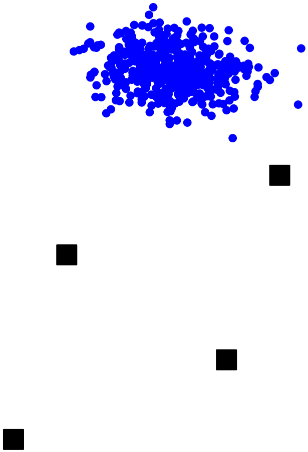

Weighted particle cloud representing possible object states at time t

$\{(x^i_{t-1}, w^i_{t-1}) \in \mathbb{R}^d \times \mathbb{R} \mid i = 1, \ldots, N\}$

Approximate posterior

$$\widehat{p}_{t-1}(x) = \frac{1}{N} \sum_{i=1}^{n} w^i_{t-1} \delta_{x^i_{t-1}}(x)$$

where $\delta_{x^i_{t-1}}(x)$ is Dirac delta mass at x^i_{t-1}

Decentralized Particle Filtering (DPF)

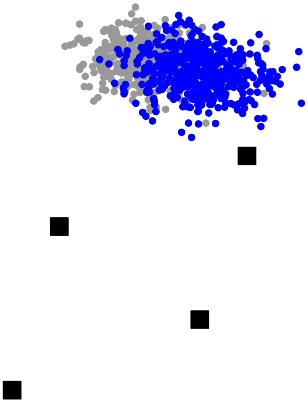

Predict new positions by sampling from the object dynamic model $p(x_t|x_{t-1}^i)$

Decentralized Particle Filtering (DPF)

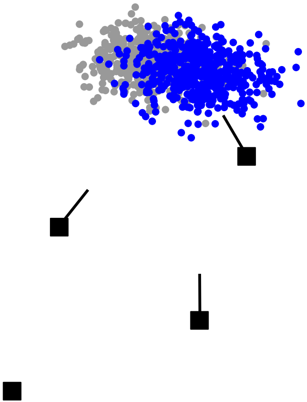

New observation y_t^j at sensor $j = 1, \dots, n$

Decentralized Particle Filtering (DPF)

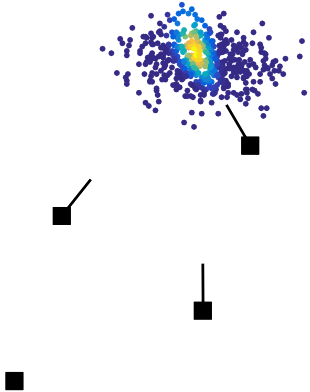

Update particle weight w_t^i proportional to

$$\log p(y_t^1, \ldots, y_t^n | x_{t|t-1}^i) = \sum_{j=1}^{n} \log p(y_t^j | x_{t|t-1}^i)$$

Data update/fusion requires communication

Consensus-Based Decentralized Particle Filtering

Survey: Hlinka, Hlawatsch, & Djuric, IEEE SP Mag 2013

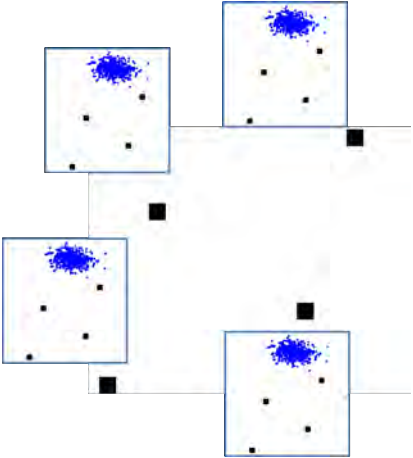

Use gossip/consensus algorithms for synchronization

Gossip Algorithms for Distributed Averaging

Seminal work: DeGroot 1974; Tsitsiklis, Bertsekas, & Athans, IEEE TAC 1986
Survey: Dimakis, Kar, Moura, Rabbat, & Scaglione, Proc. IEEE 2010

Communication topology encoded in $n \times n$ doubly-stochastic matrix P with $P_{j,\ell} > 0$ iff nodes j and ℓ communicate directly.

Node j begins with initial value $z_j(0)$ and repeats iterations

$$
\begin{aligned}
z_j(k+1) &= \sum_{\ell=1}^{n} P_{j,\ell} z_j(k) \\
&= P_{j,j} z_j(k) + \sum_{\ell \sim j} P_{j,\ell} z_\ell(k)
\end{aligned}
$$

Gossip Convergence Rates

If the communication topology is connected, then $\forall j = 1, \ldots, n$

$$\lim_{k \to \infty} z_j(k) = \frac{1}{n} \sum_{\ell=1}^{n} z_\ell(0)$$

and $|z_j(k) - \overline{z}| \leq \delta$ if

$$k \geq \frac{\log(\frac{1}{\delta}\sqrt{n}\max_\ell |z_\ell(0) - \overline{z}|)}{1 - \lambda_2(P)}.$$

Gossip Convergence Rates

If the communication topology is connected, then $\forall j = 1, \ldots, n$

$$\lim_{k \to \infty} z_j(k) = \frac{1}{n} \sum_{\ell=1}^{n} z_\ell(0)$$

and $|z_j(k) - \overline{z}| \leq \delta$ if

$$k \geq \frac{\log(\frac{1}{\delta}\sqrt{n}\,\max_\ell |z_\ell(0) - \overline{z}|)}{1 - \lambda_2(P)}.$$

ring

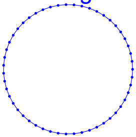

$$\frac{1}{1-\lambda_2} = \mathcal{O}(n^2)$$

random geometric graph

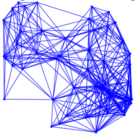

$$\frac{1}{1-\lambda_2} = \mathcal{O}\left(\frac{n}{\log(n)}\right)$$

expander

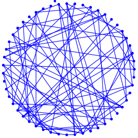

$$\frac{1}{1-\lambda_2} = \mathcal{O}(1)$$

Consensus-Based Decentralized Particle Filtering

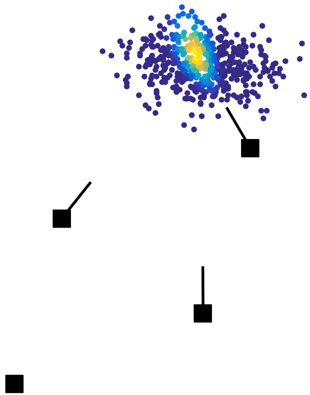

Update particle weight w_t^i proportional to

$$\log p(y_t^1, \ldots, y_t^n | x_{t|t-1}^i) = \sum_{j=1}^{n} \log p(y_t^j | x_{t|t-1}^i)$$

Data update/fusion requires communication

Naïve approach requires communication per particle

PF Convergence Results and Stability

- Strong results for stability of centralized particle filters [1–4]
- Sampling introduces an error.
- Gossip introduces an error.
- How do these errors propagate over time?

[1] F. LeGland and N. Oudjane, *Ann. App. Prob.* 2004
[2] D. Crisan and A. Doucet, *IEEE Trans. Sig. Proc.*, 2002
[3] P. Del Moral, Springer-Verlag, 2004
[4] N. Chopin, *Ann. Stat.* 2004

PF Stability: Assumptions

Bounded likelihoods:

- Let $L_t(x_t^i) = p(y_t^1, \ldots, y_t^n | x_t^i) = \prod_{j=1}^n p(y_t^j | x_t^i)$
- There exists $\epsilon_L \in (0, 1)$ s.t.

$$L_t(x_t^i) \geq \epsilon_L L_t(x_t^{i'}) \qquad \forall i, i' = 1, \ldots, N$$

PF Stability: Assumptions

Bounded likelihoods:

- Let $L_t(x_t^i) = p(y_t^1, \ldots, y_t^n | x_t^i) = \prod_{j=1}^n p(y_t^j | x_t^i)$
- There exists $\epsilon_L \in (0,1)$ s.t.

$$L_t(x_t^i) \geq \epsilon_L L_t(x_t^{i'}) \qquad \forall i, i' = 1, \ldots, N$$

Sufficient mixing dynamics:

- Let $M_t(x_{t-1}, dx_t) = p(x_t | x_{t-1}) dx_t$
- Given $m > 0$, there exists $\epsilon_m \in (0,1)$ s.t., $\forall x_t, \overline{x}_t$,

$$\begin{aligned} M_{t,t+m}(x_t, \cdot) &= M_{t+1} M_{t+2} \ldots M_{t+m}(x_t, \cdot) \\ &\geq \epsilon_m M_{t,t+m}(\overline{x}_t, \cdot) \end{aligned}$$

Bounded gossip error: There exists $\delta > 0$ s.t.

$$\frac{\left|\log \widehat{L}_t(x_t^i) - \log L_t(x_t^i)\right|}{\left|\log L_t(x_t^i)\right|} \leq \delta \qquad \forall i = 1, \ldots, N$$

Note: δ related to number of gossip iterations

Bounded gossip error: There exists $\delta > 0$ s.t.

$$\frac{\left|\log \widehat{L}_t(x_t^i) - \log L_t(x_t^i)\right|}{\left|\log L_t(x_t^i)\right|} \leq \delta \qquad \forall i = 1, \ldots, N$$

Note: δ related to number of gossip iterations

Test function has bounded oscillations: There exists $C > 0$ s.t.

$$\mathrm{osc}(h_t) = \sup_{x,x'} \left|h_t(x) - h_t(x')\right| < C$$

E.g., $h_t(x_t^i) = \left\|x_t^i - x_t\right\|_2$
Then $\widehat{\pi}^N(h_t) = \frac{1}{N} \sum_{i=1}^N w_t^i \left\|x_t^i - x_t\right\|_2$

Distributed PF Stability

We have uniform stability in the weak-sense L_2 error sense:

$$\boxed{\sup_{t\geq 0} \mathbb{E}\left[\left|[\widehat{\pi}_t^N - \pi_t](h_t)\right|^2\right]^{1/2} \leq \epsilon_0 \left(\frac{D}{\sqrt{N}} + \delta\,|\log \epsilon_L|\right)}$$

where D and ϵ_0 are independent of t, and

$$\epsilon_0 = \frac{2m}{\epsilon_m^3 \epsilon_L^{(2m-1)}}.$$

S. Datta Gupta, M. Coates, and M. Rabbat, "Error propagation in gossip-based distributed particle filters," *IEEE TSIPN*, 2015.

The Challenge: Communication Overhead

- Need to synchronize all particle weights across sensors
- Run one gossip instance for each particle? High overhead!
- Especially since
  - Many particles have low weight
  - Weights typically vary “smoothly”

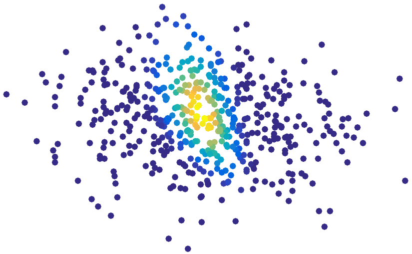

(Aside: Could communicate measurements... often worse)

A Graph-Based Approach

Idea: Transform coding with a basis adapted to the current particle cloud

- Use smoothness assumption: nearby particles have similar weight

Methodology:

- Fit a graph to the particle cloud (e.g., k-nearest neighbors)
- Use eigenbasis of Laplacian to compress log-likelihood weight vector defined over particles
- Gossip on few "low frequency" coefficients in Laplacian eigenbasis

Graph Laplacian Transform Coding

k-nn graph over particles

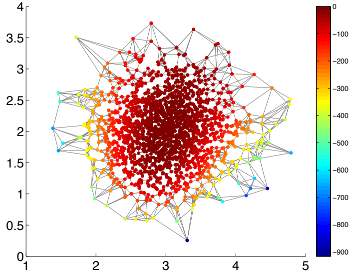

Adjacency matrix A
Degree matrix $D = \mathrm{diag}(\mathbf{1}^\mathsf{T} A)$
Laplacian $L = D - A = F \Lambda F^\mathsf{T}$

Laplacian eigenvectors F define a basis adapted to the particle distribution

Examine coefficients in this basis

$$\alpha = F^\mathsf{T} w$$

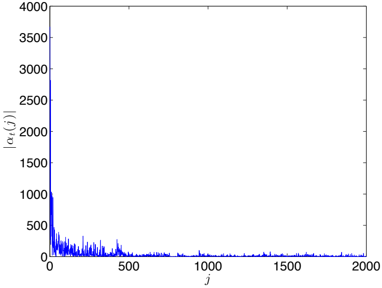

Graph Laplacian Transform Coding

Examine coefficients in this basis

$$\alpha = F^{\mathsf{T}} w$$

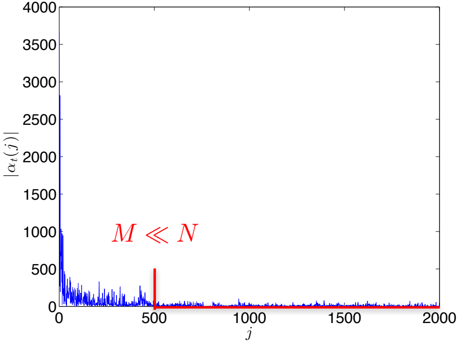

Exact weight representation

$$w = F\alpha = \sum_{j=1}^{N} \alpha_j f_j$$

Approximate by keeping "low frequency" coefficients (threshold rest to zero)

$$\widehat{w} = \sum_{j=1}^{M} \alpha_j f_j$$

Case Study

- Two targets with correlated motion
- 50 time steps
- 8-dimensional state vector
- $N = 2000$ particles

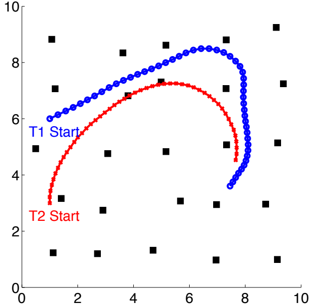

- 25 sensors, randomized grid topology
  - Only communicate with nearby neighbors
  - Vary number of gossip averaging iterations
- Sensors measure bearings only
  - Random number of detections per time step (avg $= 1/5$ / sensor / step)
  - Additive Gaussian noise, $\sigma = 5$ degrees
- Performance metric: total time-Averaged RMSE (ARMSE)

Communication-Accuracy Tradeoff

Varying number of coefficients in graph Laplacian approximation

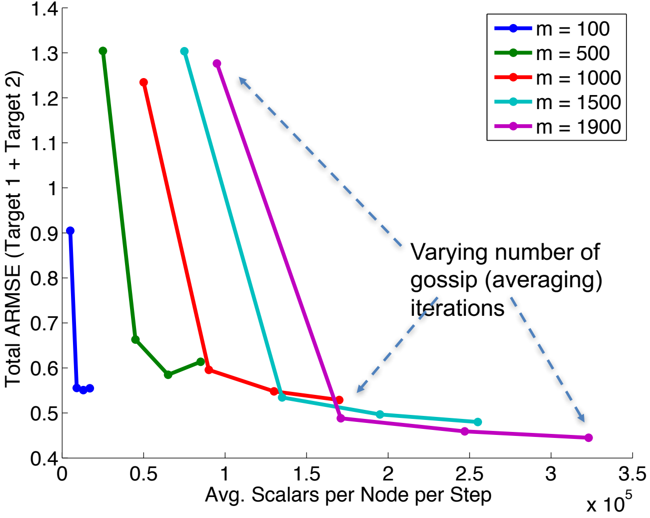

Comparison

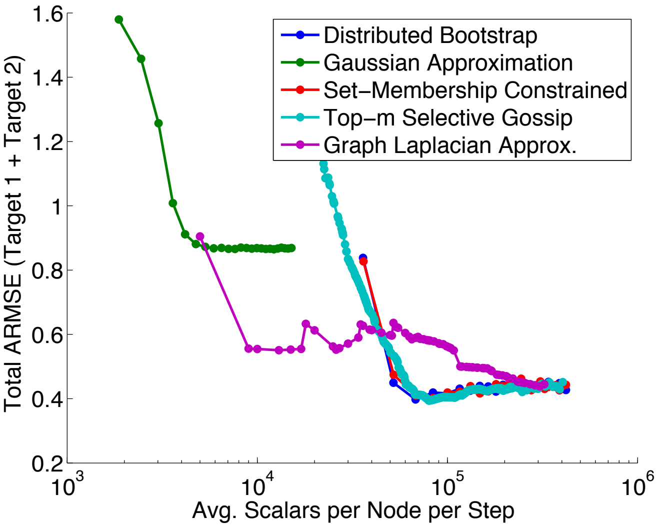

Total ARMSE (Target 1 + Target 2)
Avg. Scalars per Node per Step
Distributed Bootstrap
Gaussian Approximation
Set−Membership Constrained
Top−m Selective Gossip
Graph Laplacian Approx.
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
10^3
10^4
10^5
10^6

Conclusion

Summary:

- Stability bound for distributed particle filters
  - Depends on num. particles N, gossip error δ
- Graph-based approx. to reduce communication overhead
  - Graph captures particle proximity
  - Gives transform adapted to particle geometry

Ongoing work:

- Stability bounds for particle approximation methods
- Computationally-efficient graph-based approximation via clustering and non-linear reconstruction

michael.rabbat@mcgill.ca