Distributed optimization-based coordination and estimation for multi-agent systems

Tamás Keviczky

Delft Center for Systems and Control
Delft University of Technology
Delft, The Netherlands

Joint work with: B. / M. / K.H. } Johansson, and A. Simonetto

February 4, 2010

Motivation

- Distributed control and estimation in networks of systems
- Information and processing power is distributed among cooperating agents
- Global objective through local computations and interaction
- Design is local (on-line) as opposed to global (off-line)

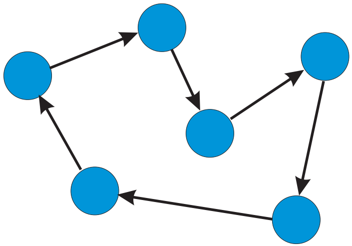

Online Optimization-based Multi-agent Coordination

1. Formulate an optimization problem for each agent that reflects the global objectives.

Online Optimization-based Multi-agent Coordination

1. Formulate an optimization problem for each agent that reflects the global objectives.
2. Incorporate local knowledge, exchange information locally. (cost function, models, constraints, neighboring states or measurements)

Online Optimization-based Multi-agent Coordination

1. Formulate an optimization problem for each agent that reflects the global objectives.
2. Incorporate local knowledge, exchange information locally. (cost function, models, constraints, neighboring states or measurements)
3. Solve and implement repeatedly.

Examples - 1

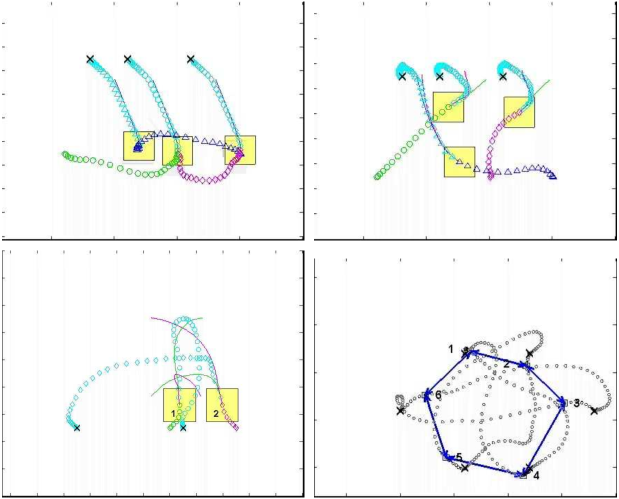

[Keviczky et al, 2004, 2006]

Examples - 2

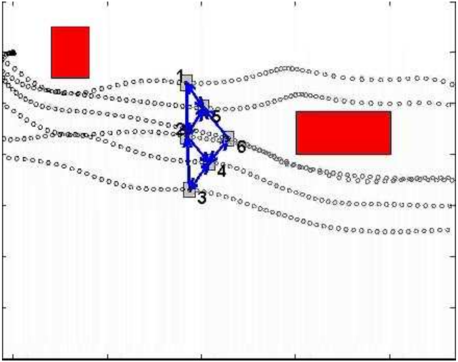

[Keviczky et al, 2008]

Examples - 3

- Moving agents to maximize λ_2 of a state-dependent graph Laplacian L.

$$\begin{aligned} \max_{x} \quad & \gamma \\ \text{subject to} \quad & \|x_i - x_j\|^2 \geq \rho \\ & P^\top L(x) P \geq \gamma I_{N-1} \end{aligned}$$

  where P is the matrix of orthonormal basis vectors spanning the subspace $\mathbf{1}^\perp = \{x \in \mathbb{R}^N | \mathbf{1}^\top x = 0\}$.

- Iterative SDP approach using discretization and linearization of state-dependent Laplacian entries.

[Kim - Mesbahi, 2006]

Examples - 3

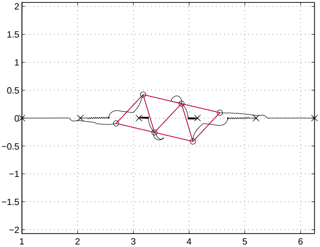

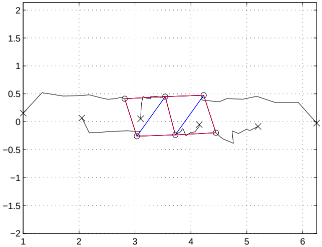

Centralized (black)

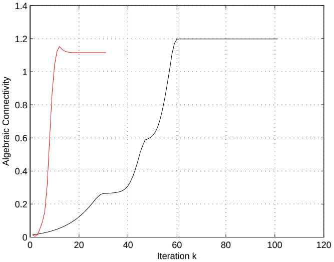

Distributed (red)

[Simonetto - Keviczky, 2010]

Typical Issues

These practical results are attractive, yet...

- decomposition is mostly ad-hoc

Typical Issues

These practical results are attractive, yet...

- decomposition is mostly ad-hoc
- no a priori guarantees on performance

Typical Issues

These practical results are attractive, yet...

- decomposition is mostly ad-hoc
- no a priori guarantees on performance
- analysis is performed a posteriori

Typical Issues

These practical results are attractive, yet...

- decomposition is mostly ad-hoc
- no a priori guarantees on performance
- analysis is performed a posteriori
- design assumptions often very conservative

Typical Issues

These practical results are attractive, yet...

- decomposition is mostly ad-hoc
- no a priori guarantees on performance
- analysis is performed a posteriori
- design assumptions often very conservative

Objective

Development of systematic design tools to distribute and coordinate the global optimization problem among agents.

Decomposition Methods for Distributed Optimization

See in Stephen Boyd's talk...

- Primal decomposition
- Dual decomposition
  (Lagrangian relaxation)
- Penalty function method
- Proximal point method
  (Augmented Lagrangian)
- Auxiliary problem principle

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.
- Finite number of coordinating iterations (a priori tunable). (Rate analysis for real-time implementation.)

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.
- Finite number of coordinating iterations (a priori tunable). (Rate analysis for real-time implementation.)
- Pre-specified distance from optimality. (Bounded suboptimality.)

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.
- Finite number of coordinating iterations (a priori tunable). (Rate analysis for real-time implementation.)
- Pre-specified distance from optimality. (Bounded suboptimality.)
- Approximate primal solutions from dual decomposition. (Primal feasibility of each iterate.)

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.
- Finite number of coordinating iterations (a priori tunable). (Rate analysis for real-time implementation.)
- Pre-specified distance from optimality. (Bounded suboptimality.)
- Approximate primal solutions from dual decomposition. (Primal feasibility of each iterate.)
- Distributed, on-line testing of termination criteria. (With a priori complexity bound.)

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.
- Finite number of coordinating iterations (a priori tunable). (Rate analysis for real-time implementation.)
- Pre-specified distance from optimality. (Bounded suboptimality.)
- Approximate primal solutions from dual decomposition. (Primal feasibility of each iterate.)
- Distributed, on-line testing of termination criteria. (With a priori complexity bound.)
- Heterogeneous subsystems and constraints.

Desiderata for Decomposition and Coordination Mechanisms

- Arbitrary connected (time-varying, delayed) communication graph.
- Finite number of coordinating iterations (a priori tunable). (Rate analysis for real-time implementation.)
- Pre-specified distance from optimality. (Bounded suboptimality.)
- Approximate primal solutions from dual decomposition. (Primal feasibility of each iterate.)
- Distributed, on-line testing of termination criteria. (With a priori complexity bound.)
- Heterogeneous subsystems and constraints.
- Non-convex coupling constraints.

Simplified Problem Formulation

$$\underset{x}{\text{minimize}} \quad f(x) = \sum_{i=1}^{N} f_i(x)$$

$$\text{subject to} \quad x \in \mathcal{X},$$

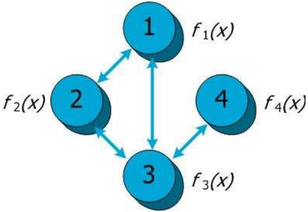

- $f_i \ : \ \mathbb{R}^M \to \mathbb{R}$ nondifferentiable convex functions.
- $\mathcal{X} \subseteq \mathbb{R}^M$ nonempty, closed, and convex set.
- Computations should be performed in a distributed fashion.

Simplified Problem Formulation

$$\underset{x}{\text{minimize}} \quad f(x) = \sum_{i=1}^{N} f_i(x)$$
$$\text{subject to} \quad x \in \mathcal{X},$$

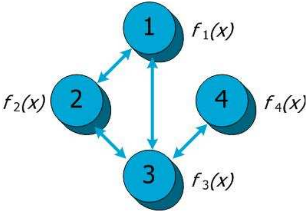

- $f_i \; : \; \mathbb{R}^M \to \mathbb{R}$ nondifferentiable convex functions.
- $\mathcal{X} \subseteq \mathbb{R}^M$ nonempty, closed, and convex set.
- Computations should be performed in a distributed fashion.
- Information exchange is only allowed through edges of an N-node undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$.

Combined Consensus/Subgradient Scheme

- Goal is to use agreement protocols to relax communication constraints in distributed optimization schemes.

Modified subgradient iterations

$$x_i^{(k+1)} = \mathcal{P}_{\mathcal{X}} \left[\sum_{j=1}^{N} [W^{\varphi}]_{ij} \left(x_j^{(k)} - \alpha^{(k)} g_j(x_j^{(k)}) \right) \right]$$

with $W = I - \varepsilon L(\mathcal{G})$ Perron matrix corresponding to the communication graph.

[Johansson et al., 2008]

Main Iterations of the Algorithm

1. Perform local subgradient update on local variable x_i.
   (This is done in parallel for all nodes.)
2. Do φ consensus iterations with neighbors.
   (Can be interpreted as enforcing approximate equality constraints with neighboring variables.)
3. Repeat.

Convergence Analysis

- Establish a lower bound on the number of consensus steps φ that will ensure that the local variables will remain in a ball of radius $\beta^{(k)}$ of their average, from one iteration to the next.

$$\beta^{(k)} = \delta^{(k)}\beta_0, \quad \delta^{(k)} \geq \delta^{(k+1)}, \frac{\delta^{(k+1)}}{\delta^{(k)}} \geq \mu$$

$$\varphi \geq \frac{\log(\mu\beta_0) - log(4M\sqrt{N}(\beta_0 + \alpha_0 C))}{\log(\gamma)}$$

Convergence Analysis

- Establish a lower bound on the number of consensus steps φ that will ensure that the local variables will remain in a ball of radius $\beta^{(k)}$ of their average, from one iteration to the next.

$$\beta^{(k)} = \delta^{(k)} \beta_0, \quad \delta^{(k)} \geq \delta^{(k+1)}, \frac{\delta^{(k+1)}}{\delta^{(k)}} \geq \mu$$

$$\varphi \geq \frac{\log(\mu \beta_0) - log(4M\sqrt{N}(\beta_0 + \alpha_0 C))}{\log(\gamma)}$$

This bound does not depend on k!

Convergence Analysis

- Establish a lower bound on the number of consensus steps φ that will ensure that the local variables will remain in a ball of radius $\beta^{(k)}$ of their average, from one iteration to the next.

$$\beta^{(k)} = \delta^{(k)}\beta_0, \quad \delta^{(k)} \geq \delta^{(k+1)}, \frac{\delta^{(k+1)}}{\delta^{(k)}} \geq \mu$$

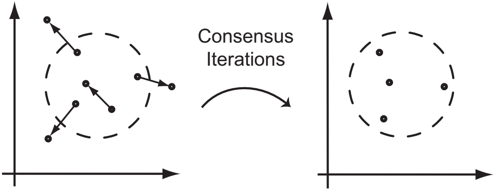

- Use approximate subgradient at the average value to account for different local subgradients.

Subgradient and ϵ-subgradient

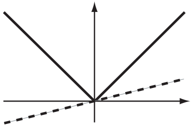

Definition

h is a subgradient for f at x if $h \in \partial f(x)$.

$$\partial f(x) = \left\{ g \in \mathbb{R}^M | f(y) \geq f(x) + g^\mathsf{T}(y - x), \forall y \in \mathbb{R}^M \right\}.$$

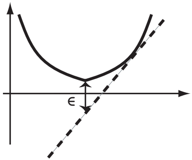

Definition

h is an ϵ-subgradient for f at x if $h \in \partial_\epsilon f(x)$.

$$\partial_\epsilon f(x) = \left\{ g \in \mathbb{R}^M | f(y) \geq f(x) + g^\mathsf{T}(y - x) - \epsilon, \forall y \in \mathbb{R}^M \right\}.$$

Convergence Result

Theorem (unconstrained case)

Under appropriate assumptions, the sequence $\{x_1^{(k)}, \ldots, x_N^{(k)}\}_{k=0}^{\infty}$ generated by the combined SG/consensus update with φ consensus iterations and $\left\|x_i^{(0)} - \bar{x}^0\right\| \leq \beta^{(0)}$, $\beta^{(k)} = \delta^{(k)}\beta_0$, $\alpha^{(k)} = \delta^{(k)}\alpha_0$, $\sum_{k=0}^{\infty} \alpha^{(k)} = \infty$ converges to the set of optimizers:

$$\lim_{k\to\infty} \mathrm{dist}_{\mathcal{X}^\star + \mathcal{B}_\eta}\left(x_i^{(k)}\right) = 0, \forall i = 1, \ldots, N,$$

with

$$\eta = \lim_{k\to\infty}\left(NC\alpha^{(k)} + \beta^{(k)}\right) + \max_{x\in\mathcal{X}_a} \mathrm{dist}_{\mathcal{X}^\star}(x),$$

$$a = \lim_{k\to\infty}\left(\frac{NC^2}{2}\alpha^{(k)} + 2N\beta^{(k)}C\right).$$

[Johansson et al., 2010]

Key Points, Remarks

- Distributed optimization for separable problems.

Key Points, Remarks

- Distributed optimization for separable problems.
- Use combination of
    - local subgradient updates
    - consensus iterations

  to relax information exchange assumptions.

Key Points, Remarks

- Distributed optimization for separable problems.
- Use combination of
  - local subgradient updates
  - consensus iterations

  to relax information exchange assumptions.
- Convergence analysis (unconstrained, constrained) using approximate subgradients.

Key Points, Remarks

- Distributed optimization for separable problems.
- Use combination of
  - local subgradient updates
  - consensus iterations

  to relax information exchange assumptions.
- Convergence analysis (unconstrained, constrained) using approximate subgradients.
- Closeness of iterates can be determined a priori.

Key Points, Remarks

- Distributed optimization for separable problems.
- Use combination of
    - local subgradient updates
    - consensus iterations

  to relax information exchange assumptions.
- Convergence analysis (unconstrained, constrained) using approximate subgradients.
- Closeness of iterates can be determined a priori.
- Trade-off between number of consensus steps and suboptimality for fixed stepsize.

Key Points, Remarks

- Distributed optimization for separable problems.
- Use combination of
  - local subgradient updates
  - consensus iterations

  to relax information exchange assumptions.
- Convergence analysis (unconstrained, constrained) using approximate subgradients.
- Closeness of iterates can be determined a priori.
- Trade-off between number of consensus steps and suboptimality for fixed stepsize.
- Convergence to optimal solution set with diminishing stepsizes.

Key Points, Remarks

- Distributed optimization for separable problems.
- Use combination of
  - local subgradient updates
  - consensus iterations

  to relax information exchange assumptions.
- Convergence analysis (unconstrained, constrained) using approximate subgradients.
- Closeness of iterates can be determined a priori.
- Trade-off between number of consensus steps and suboptimality for fixed stepsize.
- Convergence to optimal solution set with diminishing stepsizes.
- There are convergence results also for time-varying graphs and different local constraint sets (Nedic-Ozdaglar).

Illustration - Numerical Example

- Finite-time optimal rendezvous problem with three double integrator agents.
- Consensus matrix:

$$W = \begin{pmatrix} 0.75 & 0.25 & 0 \\ 0.25 & 0.5 & 0.25 \\ 0 & 0.25 & 0.75 \end{pmatrix}$$

Illustration - Numerical Example

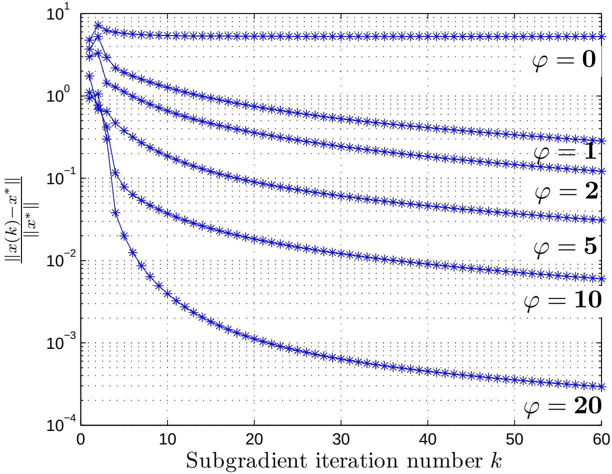

Receding Horizon Implementation: Optimal Rendezvous Example

- Consider N dynamic agents

$$x^i_{t+1} = A^i x^i_t + B^i u^i_t,$$
$$y^i_t = C^i x^i_t$$

- Polyhedral constraints

$$x^i_t \in \mathcal{X}^i, \quad u^i_t \in \mathcal{U}^i, \quad t \geq 0$$

Finite-time rendezvous

$$y^i_{T+k} = \theta, \quad \forall k \geq 0, \quad i = 1, \dots, N,$$
$$u^i_{T+k} = u^i_T, \quad \forall k \geq 0, \quad i = 1, \dots, N$$

Finite-Time Optimal Rendezvous (FTOR)
Problem Formulation

- Each agent has nontrivial, constrained LTI dynamics
- The rendezvous point $\theta \in \Theta$ is not fixed a priori, but chosen optimally

Cost function associated with the *i*-th system

$$
\begin{aligned}
V^i\left(x_k^i, u_k^i, \theta\right) = &\left(x_k^i - x_e^i(\theta)\right)^\mathsf{T} Q^i \left(x_k^i - x_e^i(\theta)\right) \\
&+ \left(u_k^i - u_e^i(\theta)\right)^\mathsf{T} R^i \left(u_k^i - u_e^i(\theta)\right)
\end{aligned}
$$

FTOR Problem Formulation

Optimization problem

$$
\begin{aligned}
\min_{U_t,\theta_t} \quad & \sum_{i=1}^{N} \sum_{k=0}^{T-1} V^i\left(x^i_{k,t}, u^i_{k,t}, \theta_t\right) \\
\text{subject to} \quad & x^i_{k+1,t} = A^i x^i_{k,t} + B^i u^i_{k,t}, \\
& y^i_{k,t} = C^i x^i_{k,t}, \\
& x^i_{k,t} \in \mathcal{X}^i, \quad k = 1, \ldots, T, \\
& u^i_{k,t} \in \mathcal{U}^i, \quad k = 0, \ldots, T-1, \\
& y^i_{T,t} = \theta_t, x^i_{T,t} = x^i_e(\theta_t), \\
& x^i_{0,t} = x^i_t, \quad i = 1, \ldots, N, \\
& \theta_t \in \Theta
\end{aligned}
$$

Primal Decomposition of FTOR

- Eliminate control inputs $u_t^i = k^i(x_t^i, \theta_t)$

$$f^i(x_t^i, \theta_t) = \min_{U_t^i} \sum_{k=0}^{T-1} V^i\left(x_{k,t}^i, u_{k,t}^i, \theta_t\right)$$
$$\text{subject to}\ \ \text{constraints}, \quad k = 1, \ldots, T-1$$

- Express the optimization problem as

$$f^*(x_t) = \min_{\theta_t} \sum_{i=1}^{N} f^i(x_t^i, \theta_t)$$
$$\text{subject to}\ \ \theta_t \in \Theta$$

- This belongs to problem class defined earlier

[Johansson et al, 2006]

Receding Horizon Implementation of FTOR

- Measure current state x_t^i
- Formulate Finite-Time Optimal Rendezvous Problem
- Solve using incremental subgradient method
  (distributed computations, sequential updates)
- Implement local solution corresponding to local rendezvous point

Numerical Example - Aerial Refueling Scenario

- Three aircraft need to rendezvous for refueling
- The rendezvous variable is altitude

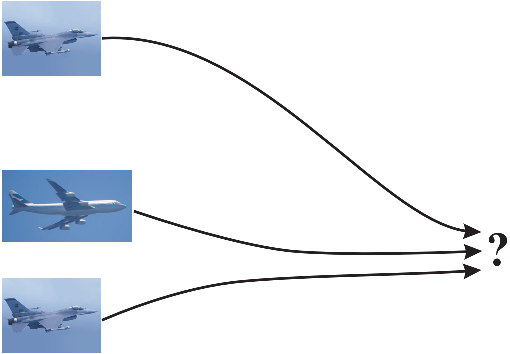

Numerical Example - Aerial Refueling Scenario

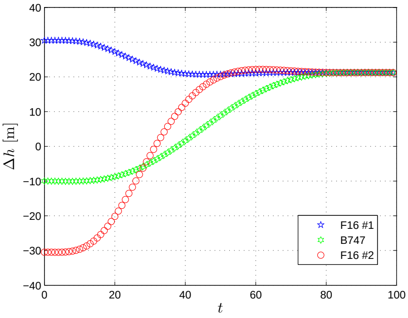

Centralized solution

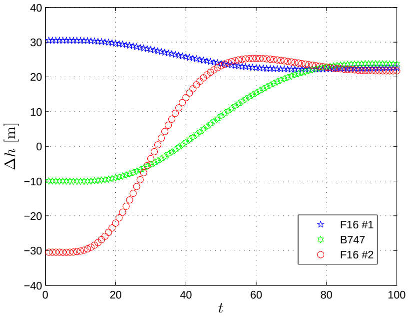

MPC implementation with 15 subgradient iterations

Numerical Example - Aerial Refueling Scenario

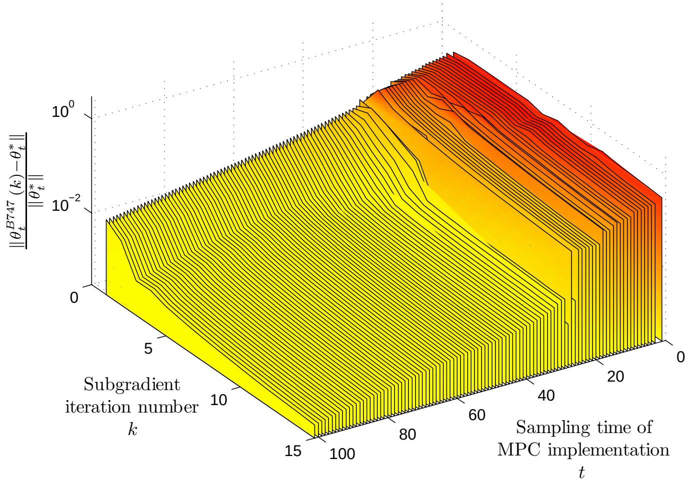

[Keviczky-Johansson, 2008]

Important Questions

What happens if the negotiations get interrupted?

1. How can we still guarantee stability?
   - additional constraints
   - global cost function (cooperation)
2. How can we still guarantee feasibility?
   - ensuring feasibility of each iterate
   - challenging with coupling constraints (dual decomposition)

Multi-agent Estimation Problems

- Weighted consensus processes, averaging widely used.
- Design choices fundamentally influence performance.

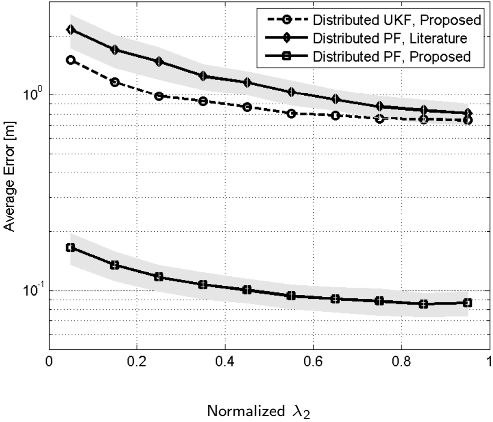

[Simonetto-Keviczky-Babuška, 2010]

Distributed Moving Horizon Estimation

Each node i in the sensor network solves a MHE problem and computes an estimate of the state, based on *neighboring* information $z_i(k)$:

$$\widehat{\phi}^*_{T,i} = \min_{\overline{x}_i, \{w_i(k)\}_{k=T-N}^{T-1}} \widehat{\phi}_{T,i}(\overline{x}_i; \{w_i(k)\}_{k=T-N}^{T-1})$$

with objective function

$$\widehat{\phi}_{T,i}(\overline{x}_i; \{w_i(k)\}_{k=T-N}^{T-1}) = \sum_{k=T-N}^{T-1} \Big(v_i^\top(k) R^{-1} v_i(k) + w_i^\top(k) Q^{-1} w_i(k) \Big)$$

$$+ (\overline{x}_i - \hat{x}_{i,c}(T-N))^\top \Pi^{-1}_{T-N,i} (\overline{x}_i - \hat{x}_{i,c}(T-N)) + \widehat{\phi}^*_{T-N,i}$$

Distributed Moving Horizon Estimation

- The arrival cost embeds a *consensus term*

$$\hat{x}_{i,c}(T-N) = \sum_{j \in \mathcal{N}_i} k_{ij} \hat{x}_j(T-N)$$

that is a weighted sum of the neighbors' estimates.

- The matrices $\Pi_{T-N,i}$ are calculated also using consensus:

$$\Pi_{T-N,i} = 2 \sum_{j \in \mathcal{N}_i} k_{ij}^2 \tilde{\Pi}_{T-N,j}$$

where each matrix $\tilde{\Pi}_{T-N,i}$ is defined by a recursive Riccati equation.

- Sufficient conditions for convergence are derived for interleaved consensus & moving horizon scheme.

[Farina et al, 2009]

Summary

- Online optimization-based multi-agent coordination.
- Various decomposition schemes for distributed optimization.
- Combined subgradient/consensus scheme.
  - Greater flexibility in the information exchange architecture via consensus process.
  - Parallel operation, local messaging, convergence analysis using approximate subgradients.
- Specific features are needed for applicability in real-time receding horizon control and estimation.