PROXIMITY WITHOUT CONSENSUS IN ONLINE MULTI-AGENT OPTIMIZATION

Alec Koppel, Brian M. Sadler, and Alejandro Ribeiro Contact E-mail: {akoppel, aribeiro}@seas.upenn.edu, brian.m.sadler6.civ@mail.mil

Penn — University of Pennsylvania

Multi-Agent Optimization

- Agents $i = 1, \dots, N$ in network $\mathcal{G} = (V, \mathcal{E})$ sequentially observe signals $\boldsymbol{\theta}_{i,t}$
  - $\Rightarrow$ Select $\mathbf{x}_i \in \mathbb{R}^p$ which are good w.r.t. global loss $\sum_i f_i$
- If decisions of $\mathbf{x}_i$ and $\mathbf{x}_j$ uncorrelated $\Rightarrow$ consensus optimization
- Consensus: agents try to minimize $\sum_i f_i$ with constraint $\mathbf{x}_i = \mathbf{x}_j$
- Consider estimation problems where decisions have correlation
  - $\Rightarrow \rho(\mathbf{x}_i, \mathbf{x}_j) \neq 0 \Rightarrow$ consensus yields worse estimation accuracy
- Goal: allow agents the leeway to select good actions w.r.t. global cost
  - $\Rightarrow$ incorporate the structure of locally observed information
  - $\Rightarrow$ avoid limitations of consensus constraints in collaborative learning
- Practical examples:
  - $\Rightarrow$ multi-target tracking problem in a sensor network
  - $\Rightarrow$ learning in robotic team with each platform in distinct domain
  - $\Rightarrow$ online source localization problems

Local and Global Estimators

- Associate to each node $i \in V$
  - $\Rightarrow$ m-strongly cvx. $f_i : \mathbb{R}^p \times \Theta_i \to \mathbb{R}$
  - $\Rightarrow$ estimator $\mathbf{x}_i \in \mathbb{R}^p$
  - $\Rightarrow$ random signal $\theta_i \in \Theta_i$
- Functions $f_i(\mathbf{x}_i, \theta_i)$ for different θ_i
  - $\Rightarrow$ obs. of distinct stochastic model

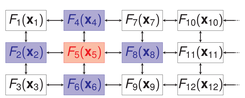

- Agent i wants to compute local estimate

$$\mathbf{x}_i^\star := \operatorname*{argmin}_{\mathbf{x}_i \in \mathbb{R}^p} F_i(\mathbf{x}_i) := \operatorname*{argmin}_{\mathbf{x}_i \in \mathbb{R}^p} \mathbb{E}_{\theta_i}[f_i(\mathbf{x}_i, \theta_i)] \ .$$

- Also aims to incorporate info. θ_j received at other nodes $j \neq i$
  - $\Rightarrow$ Could consider consensus constraint $\mathbf{x}_i = \mathbf{x}_j$ for all $j \in n_i$, all $i \in V$
  - $\Rightarrow$ Implicitly assumes distribution of θ_i same for all i

Network Proximity

- Generalization of consensus
- Nearby nodes' obs. θ_i and θ_j
  - $\Rightarrow$ similar, possibly unequal
  - $\Rightarrow$ e.g. estimate non-uniform field
- Introduce convex local proximity func.
  - $\Rightarrow h_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ with tolerance γ_{ij}
- Couple node i vars. to neighbors $j \in n_i$
  - $\Rightarrow$ E.g., $h_{ij}(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|^2 \leq \gamma_{ij}$

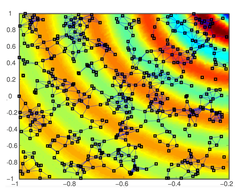

Figure: Example: source localization.

- This leads to constrained stochastic program

$$\mathbf{x}^\star := \operatorname*{argmin}_{\mathbf{x} \in \mathbb{R}^{Np}} \sum_{i=1}^{N} \mathbb{E}_{\theta_i}[f_i(\mathbf{x}_i, \theta_i)] \quad \text{s.t.} \quad h_{ij}(\mathbf{x}_i, \mathbf{x}_j) \leq \gamma_{ij}, \quad \text{for all } j \in n_i.$$

- Implicitly allows i to incorporate the relevant info. of neighbors
  - $\Rightarrow$ avoid incorporating info. from far away (weakly correlated) nodes
- We want to solve this opt. problem in distributed online settings
  - $\Rightarrow$ nodes don't know the dist. of random variable θ_i
  - $\Rightarrow$ observe local instantaneous functions $f_i(\mathbf{x}_i, \theta_i)$ sequentially.
- Distributed gradient descent and dual decomposition can't be used
  - $\Rightarrow$ they work only when the constraints $h_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ are linear.

Lagrange Relaxation and Stochastic Approximation

- At each t we approximately enforce $h_{ij}(\mathbf{x}_i, \mathbf{x}_j) \leq \gamma_{ij}$ for all $j \in n_i$
  - $\Rightarrow$ Incentivize information exchange among nearby nodes
- Lagrange relaxation of constrained optimization problem:

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = \sum_{i=1}^{N} \mathbb{E}_{\theta_i}[f_i(\mathbf{x}_i, \boldsymbol{\theta}_i)] + \sum_{i,j \in n_i} \lambda_{ij} \left(h_{ij}(\mathbf{x}_i, \mathbf{x}_j) - \gamma_{ij} \right) ,$$

  - $\Rightarrow \lambda_{ij}$ associated with proximity constraint $h_{ij}(\mathbf{x}_i, \mathbf{x}_j) \leq \gamma_{ij}$
- Convex/concave function in the primal/dual variables, respectively
- Consider stochastic approximation of Lagrangian:

$$\hat{\mathcal{L}}_t(\mathbf{x}, \boldsymbol{\lambda}) = \sum_{i=1}^{N} f_i(\mathbf{x}_i, \boldsymbol{\theta}_{i,t}) + \sum_{i,j \in n_i} \lambda_{ij} \left(h_{ij}(\mathbf{x}_i, \mathbf{x}_j) - \gamma_{ij} \right) .$$

  - $\Rightarrow$ replace average F_i with instantaneous loss f_i
  - $\Rightarrow$ instantaneous loss f_i evaluated at realization $\boldsymbol{\theta}_{i,t}$ of RVs $\boldsymbol{\theta}_i$
- Need dual set projections for bounded Lagrangian primal subgradients
  - $\Rightarrow$ project dual variable $\boldsymbol{\lambda}_{ij,t}$ onto a closed subset of $\mathbb{R}^M$, where $M = |\mathcal{E}|$

Stochastic Saddle Point Method

- Alternate primal descent and dual ascent steps on stochastic Lagrangian
  - $\Rightarrow$ Primal descent step: minimize local loss with proximity penalty term
  - $\Rightarrow$ Dual correction $\Rightarrow$ penalty coeff. associated with network proximity
- **Algorithm formulation:**

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \epsilon_t \nabla_{\mathbf{x}} \hat{\mathcal{L}}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)$$
$$\boldsymbol{\lambda}_{t+1} = \mathcal{P}_\Lambda[\boldsymbol{\lambda}_t + \epsilon_t \nabla_{\boldsymbol{\lambda}} \hat{\mathcal{L}}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)]$$

- **Decentralized estimation scheme:**

$$\text{Primal: } \mathbf{x}_{i,t+1} = \mathbf{x}_{i,t} - \epsilon_t \Big(\nabla_{\mathbf{x}_i} f_i(\mathbf{x}_{i,t}; \boldsymbol{\theta}_{i,t}) + \sum_{j \in n_i} (\lambda_{ij,t} + \lambda_{ji,t}) \nabla_{\mathbf{x}_i} h_{ij}(\mathbf{x}_{i,t}, \mathbf{x}_{j,t}) \Big)$$

$$\text{Dual: } \lambda_{ij,t+1} = \mathcal{P}_{\Lambda_{ij}} \Big[\lambda_{ij,t} + \epsilon_t \left(h_{ij}(\mathbf{x}_{i,t+1}, \mathbf{x}_{j,t+1}) - \gamma_{ij} \right) \Big] .$$

- Assume primal var. $\mathbf{x}_{i,t}$ and Lagrange multipliers $\lambda_{ij,t}$ kept by node i
- Primal and dual variables variables of distinct agents are decoupled.
- Updates require exchanges of information among neighboring nodes only.

Technical Conditions

- The network $\mathcal{G}$ is symmetric and connected with diameter D.
- Lagrangian has Lipschitz continuous gradients w.r.t. primal and dual vars.

$$\|\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) - \nabla_{\mathbf{x}} \mathcal{L}(\tilde{\mathbf{x}}, \boldsymbol{\lambda})\| \leq L_{\mathbf{x}} \|\mathbf{x} - \tilde{\mathbf{x}}\| ,$$
$$\|\nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) - \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \tilde{\boldsymbol{\lambda}})\| \leq L_{\boldsymbol{\lambda}} \|\boldsymbol{\lambda} - \tilde{\boldsymbol{\lambda}}\| .$$

- Primal & projected dual gradient of the Lagrangian are bounded by $G_{\mathbf{x}}$, $G_{\boldsymbol{\lambda}}$

$$\|\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})\| \leq G_{\mathbf{x}} \ , \ \|\tilde{\nabla}_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})\| \leq G_{\boldsymbol{\lambda}} \ .$$

- Bounded conditional second moments of primal/dual stoch. grad.:

$$\max \Big(\mathbb{E}[\|\nabla_{\mathbf{x}} \hat{\mathcal{L}}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 \,|\, \mathcal{F}_t], \ \mathbb{E}[\|\tilde{\nabla}_{\boldsymbol{\lambda}} \hat{\mathcal{L}}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 \,|\, \mathcal{F}_t] \Big) \leq \sigma^2.$$

  - $\Rightarrow \mathcal{F}_t \supseteq \{\mathbf{x}_u, \boldsymbol{\lambda}_u, \boldsymbol{\theta}_u\}_{u=1}^t$ is sigma algebra measuring alg. hist. to time t.
- $\boldsymbol{\lambda}^\star \Rightarrow$ dual optimal set. Some optimal multipliers lie in projection set:

$$\boldsymbol{\lambda}^\star \cap \Lambda \neq \emptyset$$

.

Convergence Results

Theorem: *The saddle pt. sequence $(\mathbf{x}_t, \boldsymbol{\lambda}_t)$ run with diminishing stepsize rules*

$$\sum_{t=1}^{\infty} \epsilon_t = \infty \qquad \sum_{t=1}^{\infty} \epsilon_t^2 < \infty$$

converges to a primal-dual optimal pair $(\mathbf{x}^\star, \boldsymbol{\lambda}^\star)$ in expectation as

$$\lim_{t \to \infty} \mathbb{E}\|\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}_t, \boldsymbol{\lambda}_t)\| = 0 \ ,$$

and the dual iterates asymptotically achieve the feasibility condition

$$\lim_{t \to \infty} \mathbb{E}\|\tilde{\nabla}_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}_{t+1}, \boldsymbol{\lambda}_t)\| = 0 \ .$$

Theorem: *The Lagrangian $\mathcal{L}(\mathbf{x}_t, \boldsymbol{\lambda}_t)$ evaluated at the saddle pt. sequence $(\mathbf{x}_t, \boldsymbol{\lambda}_t)$ converges to a nbhd. of its value at a primal-dual optimal pair $\mathcal{L}(\mathbf{x}^\star, \boldsymbol{\lambda}^\star)$ when we use a constant stepsize $\epsilon_t = \epsilon < 1/(2m)$*

$$\liminf_{t \to \infty} \mathcal{L}(\mathbf{x}_t, \boldsymbol{\lambda}_t) - \mathcal{L}(\mathbf{x}^\star, \boldsymbol{\lambda}^\star) \leq \frac{\epsilon L_{\mathbf{x}} \sigma^2}{4m}.$$

Moreover, the expected error sequence converges linearly to a nbhd.

$$\mathbb{E}[\mathcal{L}(\mathbf{x}_t, \boldsymbol{\lambda}_t) - \mathcal{L}(\mathbf{x}^\star, \boldsymbol{\lambda}^\star)] \leq (1 - 2m\epsilon)^t \left[\mathcal{L}(\mathbf{x}_0, \boldsymbol{\lambda}_0) - \mathcal{L}(\mathbf{x}^\star, \boldsymbol{\lambda}^\star) \right] + \frac{\epsilon L_{\mathbf{x}} \sigma^2}{4m} \ .$$

Random Field Estimation

- $\boldsymbol{\theta}_{i,t} \in \mathbb{R}^q \Rightarrow$ the observation collected by sensor i at time t.
- Obs. related to signal via noisy linear transformation $\boldsymbol{\theta}_{i,t} = \mathbf{H}_i \mathbf{x}_i + \mathbf{w}_{i,t}$
- Signal $\mathbf{x}_i \in \mathbb{R}^p$ is contaminated w/ i.i.d Gaussian noise $\mathbf{w}_{i,t} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$
- Local MMSE estimator: $\mathbb{E}_{\theta_i} f_i(\mathbf{x}_i, \theta_i) = \|\mathbf{H}_i \mathbf{x}_i - \theta_i\|^2$
- May improve estimator quality using correlated info. of adjacent nodes
  - $\Rightarrow$ enforcing equality across the network would hurt estimator equality.
- This leads to the problem:

$$\mathbf{x}^\star := \operatorname*{argmin}_{\mathbf{x} \in \mathbb{R}^{Np}} \sum_{i=1}^{N} \mathbb{E}_{\theta_i} \Big[\|\mathbf{H}_i \mathbf{x}_i - \theta_i\|^2 \Big] \quad \text{s.t.} \quad (1/2)\|\mathbf{x}_i - \mathbf{x}_j\|^2 \leq \gamma_{ij}, \quad \text{for all } j \in n_i.$$

- Constraint $(1/2)\|\mathbf{x}_i - \mathbf{x}_j\|^2 \leq \gamma_{ij} \Rightarrow \mathbf{x}_i^\star$ close to $\mathbf{x}_j^\star$ of neighbors $j \in n_i$
  - $\Rightarrow$ Not so close to the estimates $\mathbf{x}_k^\star$ of nonadjacent nodes $k \notin n_i$.
- **Saddle point algorithm for the random field estimation problem:**

$$\mathbf{x}_{i,t+1} = \mathbf{x}_{i,t} - \epsilon_t \Big[2\mathbf{H}_i^T (\mathbf{H}_i \mathbf{x}_{i,t} - \boldsymbol{\theta}_{i,t}) + \sum_{j \in n_i} \left(\lambda_{ij,t} + \lambda_{ji,t} \right) \left(\mathbf{x}_{i,t} - \mathbf{x}_{j,t} \right) \Big].$$

$$\lambda_{ij,t+1} = \mathcal{P}_{\Lambda_{ij}} \Big[\lambda_{ij,t} + (\epsilon_t/2) \left(\|\mathbf{x}_{i,t+1} - \mathbf{x}_{j,t+1}\|^2 - \gamma_{ij} \right) \Big] \ .$$

- Decentralized estimation scheme for node i
  - $\Rightarrow$ gives preference to local and nearby information

Correlated Random Field Estimation

- Scalar case: $p = q = 1$,
- Scalar $\mathbf{H} = 1$, true signal $\mathbf{x} = \mathbf{1}$
- Algorithm run $T = 500$ iterations
- Sensors form grid network
  - $\Rightarrow$ 200×200 sq. meter region
- Distance-based correlation: $\rho(\mathbf{x}_i, \mathbf{x}_j) = e^{-\|l_i - l_j\|}$,
  - $\Rightarrow$ l_j is location of node j
- Sensors learn global information
  - $\Rightarrow$ prefer to nearby nodes' info.
- Saddle pt. > LMMSE estimator
- Benefit of saddle pt.
  - $\Rightarrow$ larger in larger regions
  - $\Rightarrow$ lower SNR settings

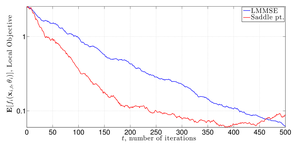

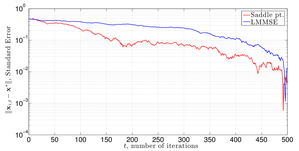

Decentralized Online Source Localization

- Consider an N sensor array in some deployed environment $\mathcal{A} \subset \mathbb{R}^p$
- $\mathbf{l}_i \in \mathbb{R}^p \Rightarrow$ position of sensor i in a deployed environment $\mathcal{A} \subset \mathbb{R}^p$
- Each node seeks location of source signal $\mathbf{x} \in \mathbb{R}^p$ through range obs.

$$r_{i,t} = \|\mathbf{x} - \mathbf{l}_i\| + \varepsilon_{i,t}$$

- $\varepsilon_t = [\varepsilon_{1,t}; \cdots; \varepsilon_{N,t}] \Rightarrow$ unknown noise vector.
- Source localization $\Rightarrow$ wireless communications, geophysics, robotics
- The squared range-based least squares (SR-LS) problem is stated as

$$\mathbf{x}^\star := \operatorname*{argmin}_{\mathbf{x} \in \mathbb{R}^p} \sum_{i=1}^{N} \mathbb{E}_{\mathbf{r}_i} \Big(\|\mathbf{l}_i - \mathbf{x}\|^2 - r_i^2 \Big)^2 \ .$$

- Nonconvex $\Rightarrow$ to convexify, expand sq., add constraint, change of vars.
- In practical settings, SNR is higher for sensors nearer to the source
- This motivates the enforcement of the quadratic constraint

$$\|\mathbf{x}_i - \mathbf{x}_j\|^2 \leq \min\{\|\mathbf{x}_i - \mathbf{l}_i\|^2, \|\mathbf{x}_j - \mathbf{l}_j\|^2\} \text{ for all } j \in n_i$$

- Sensor i weights importance of sensors $j \in n_i$ by restricting its estimate $\mathbf{x}_i$
  - $\Rightarrow \ell_2$ ball centered at neighbors' estimate $\mathbf{x}_j$
  - $\Rightarrow$ radius of ℓ_2 ball $\Rightarrow$ pairwise min. of estimated distance to source
  - $\Rightarrow$ we enforce a convex approx. of proximity constraint (log-sum-exp)
- **Saddle point method for online source localization**

$$\mathbf{y}_{i,t+1} = \mathbf{y}_{i,t} - \epsilon_t \Big(2\mathbf{A}_{i,t}^T (\mathbf{A}_{i,t} \mathbf{y}_{i,t} - \mathbf{b}_{i,t}) + \sum_{j \in n_i} \lambda_{ij,t} \Big(\frac{e^{\|\mathbf{y}_{i,t} - \mathbf{l}_i\|^2} (\mathbf{y}_{i,t} - \mathbf{l}_i)}{e^{\|\mathbf{y}_{i,t} - \mathbf{l}_i\|^2} + e^{\|\mathbf{y}_{i,t} - \mathbf{l}_j\|^2}} + (\mathbf{y}_{i,t} - \mathbf{y}_{j,t}) \Big) \Big),$$

- Dual update at the link layer of the sensor network

$$\lambda_{ij,t+1} = \mathcal{P}_{\Lambda_{ij}} \Big[\lambda_{ij,t} + \epsilon_t g(\mathbf{y}_{i,t}, \mathbf{y}_{j,t}) \Big]$$

- New method for online source localization in a sensor network

Numerical Results

- $N = 64$ sensors
  - $\Rightarrow$ deployed in a grid formation
  - $\Rightarrow$ 8×8 square in a planar 1000×1000 region
- Noise is zero-mean Gaussian
  - $\Rightarrow$ Variance $\sigma^2 = 2\|\mathbf{l}_i - \mathbf{x}^\star\|$
- $\mathbf{x}^\star$ located at avg. sensor location.
  - $\Rightarrow \propto$ to distance to source

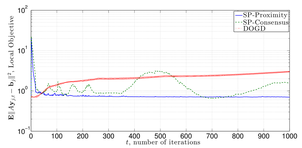

- **Consensus comparison**
- Proximity constrained saddle pt.
  - $\Rightarrow$ consensus saddle pt.
  - $\Rightarrow$ distributed gradient descent

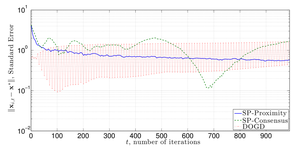

- Proximity constrained saddle pt.
  - $\Rightarrow$ outperforms consensus
  - $\Rightarrow$ best objective convergence
  - $\Rightarrow$ smallest standard error
  - $\Rightarrow$ dual domain convergence

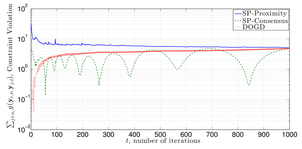

Conclusions

- We focus on online multi-agent optimization
  - $\Rightarrow$ generalize methods based on consensus constraints
  - $\Rightarrow$ motivated by cases where agents draws obs. from distinct dist.
- Consider stochastic extension of Arrow-Hurwicz saddle pt. method
- Establish convergence to primal-dual optimal pair
  - $\Rightarrow$ in diminishing and constant step-size schemes