

Non-Convex Phase Synchronization via Projected Gradient Ascent with Provable Estimation and Convergence Guarantees

Anthony Man-Cho So

Department of Systems Engineering & Engineering Management
The Chinese University of Hong Kong (CUHK)

(Joint Work with Huikang Liu and Man-Chung Yue)

**The Greater Bay Area Optimization Day
The Hong Kong Polytechnic University**

24 February 2018

Phase Synchronization

- **Goal:** Estimate an unknown phase vector $z^* \in \mathbb{T}^n = \{z \in \mathbb{C}^n : |z_1| = \dots = |z_n| = 1\}$ from the noisy measurements

$$C_{j\ell} = z_j^* \bar{z}_\ell^* + \Delta_{j\ell} \quad \text{for } 1 \leq j < \ell \leq n.$$

Here, we assume that the noise satisfies $\Delta_{jj} = 0$ and $\Delta_{j\ell} = \bar{\Delta}_{\ell j}$. Then, we can write $C = z^* z^{*H} + \Delta$ with $C_{jj} = 1$ and $C, \Delta \in \mathbb{H}^n$.

- Various applications
 - lensless imaging
 - clock synchronization in wireless networks
 - localization on circle
 - ranking of items based on noisy pairwise comparisons

Phase Synchronization

- Consider the following natural least-squares formulation:

$$\hat{z} \in \arg \min_{z \in \mathbb{T}^n} \sum_{1 \leq j < \ell \leq n} |C_{j\ell} - z_j \bar{z}_\ell|^2 = \arg \max_{z \in \mathbb{T}^n} \{f(z) = z^H C z\}. \quad (\text{QP})$$

- Note that if \hat{z} is an optimal solution, then so is $e^{i\theta} \hat{z}$ for any $\theta \in [0, 2\pi)$. Thus, we can at most identify z^* up to a global phase.
- In the Gaussian noise setting (i.e., $\Delta = \sigma W$, where $\sigma^2 > 0$ is the noise power and W is a Wigner matrix), any optimal solution to (QP) is a maximum likelihood estimator of z^* .

Phase Synchronization

- Given the estimation problem (QP), we are interested in the following:
 - Is \hat{z} (optimal solution) close to z^* (ground truth)?
 - Can we find \hat{z} efficiently?
- Recall that the value $z^H C z$ is invariant under multiplication of a common phase (unit-modulus complex number).
 - This motivates the following distance measure for two points w, z on \mathbb{T}^n :

$$d_2(w, z) = \min_{\theta \in [0, 2\pi)} \|w - e^{i\theta} z\|_2.$$

- **Fact: [Bandeira-Boumal-Singer'17]** (Estimation Error of \hat{z})

$$d_2(\hat{z}, z^*) \leq \frac{4\|\Delta\|_{\text{op}}}{\sqrt{n}}.$$

- Thus, we shall mainly focus on the second question.

Phase Synchronization

- Given an arbitrary C , Problem (QP) is NP-hard [**Toker-Özbay'98**].
- However, we have a highly structured C . Perhaps we can find \hat{z} efficiently?
- A natural approach: SDP relaxation

$$\min_{Z \in \mathbb{H}^n} \text{Tr}(CZ) \quad \text{subject to} \quad \text{diag}(Z) = e, \quad Z \succeq \mathbf{0}. \quad (\text{SDP})$$

- **Fact: [Bandeira-Boumal-Singer'17]** Under the Gaussian noise setting, if $\sigma = O(n^{1/4})$, then (SDP) admits a unique optimal solution \hat{Z} that is of rank-one.
- Despite its theoretical appeal, solving large instances of (SDP) is considered to be computationally expensive.

Can we develop fast methods for solving (QP) while still being able to establish some theoretical guarantee?

Generalized Power Method for Phase Synchronization

Consider the following projected gradient-type scheme for solving (QP):

Algorithm 1 Generalized Power Method for Problem (QP) [Boumal'16]

```
1: input: objective matrix  $C \in \mathbb{H}^n$ , step size  $\alpha > 0$ , initial point  $z^0 \in \mathbb{T}^n$ 
2: for  $k = 0, 1, \dots$  do
3:   if termination criterion is met then
4:     return  $z^k$ 
5:   else
6:      $w^k \leftarrow (I + \frac{\alpha}{n}C) z^k$  // gradient ascent
7:      $z^{k+1} \leftarrow \frac{w^k}{|w^k|}$  // projection onto  $\mathbb{T}^n$ 
8:   end if
9: end for
```

- Here, $2Cz^k$ is the gradient of $z \mapsto z^H C z$ at z^k and

$$\left(\frac{v}{|v|} \right)_j = \begin{cases} \frac{v_j}{|v_j|} & \text{if } v_j \neq 0, \\ 1 & \text{otherwise} \end{cases}$$

is the projection of $v \in \mathbb{C}^n$ onto \mathbb{T}^n , which can be efficiently computed.

Generalized Power Method for Phase Synchronization

- When $\alpha \rightarrow \infty$, the update becomes $z^{k+1} \leftarrow \frac{Cz^k}{|Cz^k|}$, which resembles the **power method** for computing the dominant eigenvector of C , hence the name *generalized power method* (GPM) for Algorithm 1.
- Problem (QP) is non-convex in general. Thus, one expects that the **initial point** will affect the convergence performance of the algorithm.
- Consider the so-called **spectral initialization** $z^0 = v_C$, where $v_C = \frac{u}{|u|}$ and u is a dominant eigenvector of C .
- As it turns out, the spectral initialization has many nice properties.

Generalized Power Method for Phase Synchronization

- **Fact: [Boumal'16]** (Estimation Error of v_C)

$$d_2(v_C, z^*) \leq \frac{8\|\Delta\|_{\text{op}}}{\sqrt{n}}$$

(recall that $d_2(\hat{z}, z^*) \leq \frac{4\|\Delta\|_{\text{op}}}{\sqrt{n}}$).

- **Fact: [Boumal'16]** Under the Gaussian noise setting, if $\sigma = O(n^{1/6})$ and $\alpha = O(n^{1/2})$, then the GPM with spectral initialization will converge to an optimal solution to (QP).
- The above results do not resolve three key issues:
 - Are the estimation errors of the iterates $\{z^k\}$ generated by the GPM **monotonically decreasing**?
 - The convergence result applies to the noise regime $\sigma = O(n^{1/6})$, while the SDP tightness result applies to the **less restrictive** noise regime $\sigma = O(n^{1/4})$. Could the convergence result for the GPM apply to the latter noise regime?
 - Can we establish the **convergence rate** of the GPM?

Estimation Performance of the GPM

- **Theorem: [Liu-Yue-S.'17]** (Estimation Errors of the GPM Iterates) Suppose that

- $\|\Delta\|_{\text{op}} \leq \frac{n}{16}$,
- $\alpha \geq 6$,
- $z^0 = v_C$.

Then, the iterates $\{z^k\}_{k \geq 0}$ generated by the GPM satisfy

$$d_2(z^{k+1}, z^*) \leq \mu^{k+1} \cdot d_2(z^0, z^*) + \frac{\nu}{1 - \mu} \cdot \frac{8\|\Delta\|_{\text{op}}}{\sqrt{n}}$$

for some $\mu \in (0, 1)$ and $\nu > 0$.

- **Curious observation:** The above result does not assume the convergence of the GPM. It gives upper bounds on the estimation errors of the iterates, and the bounds are **decreasing**.
- Under the Gaussian noise setting, we have $\|\Delta\|_{\text{op}} = O(\sigma n^{1/2})$ with high probability. Thus, the above result applies to the noise regime $\sigma = O(n^{1/2})$.

Key Property of Projection onto \mathbb{T}^n

- The proof of the theorem relies crucially on the following result:

Proposition: For any $w \in \mathbb{C}^n$ and $z \in \mathbb{T}^n$,

$$\left\| \frac{w}{|w|} - z \right\|_2 \leq 2\|w - z\|_2.$$

- Roughly speaking, it says that the projection operator onto \mathbb{T}^n is not too expansive (recall that projections onto closed convex sets are always non-expansive).

Convergence Rate of the GPM

- From an optimization-theoretic viewpoint, we are still interested in knowing whether the GPM will converge or not, and if so, what is the rate of convergence.
- To answer these questions, we need to first get a handle on the set of potential optimal solutions to (QP).

Optimality Conditions of Problem (QP)

- Recall

$$\hat{z} \in \arg \max_{z \in \mathbb{T}^n} \{f(z) = z^H C z\}. \quad (\text{QP})$$

- Viewing \mathbb{T}^n as a manifold, the **tangent space** to \mathbb{T}^n at $z \in \mathbb{T}^n$ is given by

$$T_z \mathbb{T}^n = \{w \in \mathbb{C}^n : \Re\{w_j \bar{z}_j\} = 0 \text{ for } j = 1, \dots, n\}.$$

The **projector** onto $T_z \mathbb{T}^n$ can be computed as

$$\Pi_{T_z \mathbb{T}^n}(w) = w - \text{Diag}(\Re\{z_j \bar{w}_j\})z.$$

This yields the **first-order optimality condition** of (QP):

$$\mathbf{0} = \text{grad } f(z) = \Pi_{T_z \mathbb{T}^n}(2Cz) = -2S(z)z,$$

where $S(z) = \text{Diag}(\Re((Cz)_j \bar{z}_j)) - C$. We call $\text{grad } f(z)$ the **Riemannian gradient** of f at $z \in \mathbb{T}^n$.

Optimality Conditions of Problem (QP)

- We can also write down the second-order optimality condition. To do this, we need the **Riemannian Hessian** $\text{Hess } f(z)$ of f at $z \in \mathbb{T}^n$, which is obtained by projecting the directional derivatives of $\text{grad } f(z)$ onto the tangent space to \mathbb{T}^n at $z \in \mathbb{T}^n$:

$$(\text{Hess } f(z))(w) = \Pi_{T_z \mathbb{T}^n}(D \text{grad } f(z))(w) = -\Pi_{T_z \mathbb{T}^n}(2S(z)w).$$

- Then, the **second-order optimality condition** is given by

$$w^H (\text{Hess } f(z))w = -2w^H S(z)w \leq 0 \quad \text{for all } w \in T_z \mathbb{T}^n.$$

- **Fact:** Every optimal solution to (QP) is a **second-order critical point**; i.e., it satisfies both the first- and second-order optimality conditions.
- **Fact: [Boumal'16]** Every second-order critical point $\tilde{z} \in \mathbb{T}^n$ satisfies

$$\left(\text{Diag}(|\tilde{C}\tilde{z}|) - \tilde{C} \right) \tilde{z} = (\text{Diag}(|C\tilde{z}|) - C) \tilde{z} = \mathbf{0},$$

where $\tilde{C} = C + \frac{n}{\alpha}I$, for any $\alpha > 0$.

Proximity to Second-Order Criticality

- The earlier discussion motivates us to define the following:

$$\Sigma(z) = \text{Diag}(|\tilde{C}z|) - \tilde{C}, \quad \rho(z) = \|\Sigma(z)z\|_2.$$

- Intuitively, the residual function ρ measures the proximity to the set of second-order critical points.
 - This suggests ρ can be used to measure the progress of the GPM.
 - **Challenge:** Although $\rho(\hat{z}) = 0$ for any optimal solution \hat{z} to (QP), it is not clear the converse holds.

Local Error Bound for Problem (QP)

- **Proposition: [Liu-Yue-S.'17]** (Local Error Bound for (QP)) Under the Gaussian noise setting, if $\sigma = O(n^{1/4})$ and $\alpha \geq 4$, then with high probability, we have

$$d_2(z, \hat{z}) \leq \frac{8}{n} \rho(z) \quad (\text{EB})$$

for any $z \in \mathbb{T}^n$ satisfying $d_2(z, z^*) \leq \frac{\sqrt{n}}{2}$ and any optimal solution \hat{z} to (QP).

- This shows that every second-order critical point of (QP) is optimal.
- The above error bound allows us to follow rather standard arguments to establish the convergence rate of the GPM.

Convergence Rate of the GPM

- **Proposition: [Liu-Yue-S.'17]** Under the Gaussian noise setting, suppose that
 - $\sigma = O(n^{1/4})$,
 - $\alpha = O(n^{1/4})$,
 - $z^0 = v_C$.

Then, for any optimal solution \hat{z} to (QP), the iterates $\{z^k\}_{k \geq 0}$ generated by the GPM possess the following properties:

- **(Sufficient Ascent)** There exists a constant $a_0 > 0$ such that

$$f(z^{k+1}) - f(z^k) \geq a_0 \cdot \|z^{k+1} - z^k\|_2^2. \quad (\text{A1})$$

- **(Cost-to-Go Estimate)** There exists a constant $a_1 > 0$ such that

$$f(\hat{z}) - f(z^k) \leq a_1 \cdot d_2(z^k, \hat{z})^2. \quad (\text{A2})$$

- **(Safeguard)** There exists a constant $a_2 > 0$ such that

$$\rho(z^k) \leq a_2 \cdot \|z^{k+1} - z^k\|_2. \quad (\text{A3})$$

Convergence Rate of the GPM

- By combining (EB) and (A1)–(A3), we have the following:

Theorem: [Liu-Yue-S'.17] (Linear Convergence of the GPM) Under the above setting, for any optimal solution \hat{z} to (QP), the iterates $\{z^k\}_{k \geq 0}$ generated by the GPM satisfy

$$\begin{aligned} f(\hat{z}) - f(z^k) &\leq (f(\hat{z}) - f(z^0)) \lambda^k, \\ d_2(z^k, \hat{z}) &\leq a (f(\hat{z}) - f(z^0))^{1/2} \lambda^{k/2} \end{aligned}$$

for some $a > 0$ and $\lambda \in (0, 1)$.

- The above result applies to the noise regime $\sigma = O(n^{1/4})$, which matches that required for the tightness of the SDP relaxation.

Extension to MIMO Detection

- The GPM is rather versatile and can be easily adapted to other problems.
- Take, for instance, MIMO detection, where we have the model

$$\mathbf{y} = \mathbf{H}\mathbf{x}^* + \boldsymbol{\nu}.$$

Here,

- $\mathbf{y} \in \mathbb{C}^m$ is the received signal vector,
 - $\mathbf{H} \in \mathbb{C}^{m \times n}$ is the channel matrix,
 - $\mathbf{x}^* \in \mathbb{C}^n$ is the transmitted symbol vector,
 - $\boldsymbol{\nu} \in \mathbb{C}^m$ is the noise vector.
- We assume that each symbol x_i is drawn from some discrete constellation \mathcal{S} , where \mathcal{S} is either the $(4u^2)$ -QAM constellation

$$\mathcal{Q}_u = \{z \in \mathbb{C} : \Re(z), \Im(z) = \pm 1, \pm 3, \dots, \pm(2u - 1)\}$$

or the MPSK constellation

$$\mathcal{S}_M = \{\exp(2\pi ik/M) : k = 0, 1, \dots, M - 1\},$$

Extension to MIMO Detection

- We are interested in the following maximum likelihood (ML) estimation problem:

$$\min_{\mathbf{x} \in \mathcal{S}^n} \{ F(\mathbf{x}) = \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 \}. \quad (\text{ML})$$

- We can then write down the GPM for solving (ML) as follows:

Algorithm 2 Generalized Power Method for MIMO Detection

- 1: **input:** initial point $\mathbf{x}^0 \in \mathcal{S}^n$ and step sizes $\{\alpha_k\}_{k \geq 0}$
 - 2: **if** stopping criterion is not met **then**
 - 3: $\nabla F(\mathbf{x}^k) \leftarrow 2\mathbf{H}^*(\mathbf{H}\mathbf{x}^k - \mathbf{y})$
 - 4: $\mathbf{x}^{k+1} \leftarrow \Pi_{\mathcal{S}^n}(\mathbf{x}^k - \frac{\alpha_k}{m} \nabla F(\mathbf{x}^k))$
 - 5: $k \leftarrow k + 1$
 - 6: **end if**
-

- Note that for $\mathcal{S} = \mathcal{Q}_u$ and $\mathcal{S} = \mathcal{S}_M$, the projection $\Pi_{\mathcal{S}^n}(w)$ can be efficiently computed.

Analysis of the GPM for MIMO Detection

- Interestingly, we can also establish some convergence guarantee for the GPM.
- Key to the analysis is the following result, which concerns the not-too-expansiveness of the projection operator $\Pi_{\mathcal{S}^n}$:

Proposition: [Liu-Yue-S.-Ma'17] Consider the case where $\mathcal{S} = \mathcal{Q}_u$ or $\mathcal{S} = \mathcal{S}_M$. Let $\mathbf{z} \in \mathbb{C}^n$ and $\mathbf{x} \in \mathcal{S}^n$ be given. Then,

$$\|\Pi_{\mathcal{S}^n}(\mathbf{z}) - \mathbf{x}\|_2 \leq 2\|\mathbf{z} - \mathbf{x}\|_2.$$

Analysis of the GPM for MIMO Detection

- **Theorem:** [Liu-Yue-S.-Ma'17] Let $\{\mathbf{x}^k\}_{k \geq 0}$ be the sequence of iterates generated by the GPM with step sizes $\{\alpha_k\}_{k \geq 0}$ satisfying

$$\left\| \frac{2\alpha_k}{m} \mathbf{H}^* \boldsymbol{\nu} \right\|_{\infty} < \frac{1}{c} \quad \text{and} \quad \left\| \mathbf{I} - \frac{2\alpha_k}{m} \mathbf{H}^* \mathbf{H} \right\|_{\text{op}} \leq \beta < \frac{1}{4}, \quad (\text{P})$$

where $c = \frac{4}{\min_{s \neq s' \in \mathcal{S}} |s - s'|} < \infty$ (hence, we have $c = 2$ for $\mathcal{S} = \mathcal{Q}_u$ and $c = \frac{2}{\sin(\pi/M)}$ for $\mathcal{S} = \mathcal{S}_M$). Then, we have

$$\|\mathbf{x}^{k+1} - \mathbf{x}^*\|_2 \leq 4\beta \|\mathbf{x}^k - \mathbf{x}^*\|_2$$

for all $k \geq 0$.

- In particular, after at most $k^* = \left\lceil \ln \left(\frac{2}{c \|\mathbf{x}^0 - \mathbf{x}^*\|_2} \right) / \ln(4\beta) \right\rceil$ iterations, we have $\mathbf{x}^k = \mathbf{x}^*$ for all $k \geq k^*$; i.e., the GPM admits finite convergence.
- Under suitable probabilistic assumptions on \mathbf{H} and $\boldsymbol{\nu}$, one can show that condition (P) will be satisfied with high probability.

Closing Remarks

- There has been much recent interest in the design and analysis of fast methods for **structured non-convex** optimization problems.
- Many such problems arise in machine learning and signal processing.
- Similar to the convex case, **error bounds** play a fundamental role in the convergence analyses of these methods.

Thank You!