

Efficient Detection of Long Consistent Cycles and its Application to Distributed Synchronization *

Shaohan Li
University of Minnesota
li000743@umn.edu

Yunpeng Shi
University of California, Davis
ypshi@ucdavis.edu

Gilad Lerman
University of Minnesota
lerman@umn.edu

Abstract

Group synchronization plays a crucial role in global pipelines for Structure from Motion (SfM). Its formulation is nonconvex and it is faced with highly corrupted measurements. Cycle consistency has been effective in addressing these challenges. However, computationally efficient solutions are needed for cycles longer than three, especially in practical scenarios where 3-cycles are unavailable. To overcome this computational bottleneck, we propose an algorithm for group synchronization that leverages information from cycles of lengths ranging from three to six with a time complexity of order $O(n^3)$ (or $O(n^{2.373})$ when using a faster matrix multiplication algorithm). We establish non-trivial theory for this and related methods that achieves competitive sample complexity, assuming the uniform corruption model. To advocate the practical need for our method, we consider distributed group synchronization, which requires at least 4-cycles, and we illustrate state-of-the-art performance by our method in this context.

1. Introduction

Structure from Motion (SfM) asks to recover the 3D structure of a stationary scene from multiple images taken by cameras from different orientations and locations. In the past decade, the global SfM pipeline has become increasingly popular due to its several advantages over the incremental pipelines [17, 31]. First of all, global SfM requires only one implementation of bundle adjustment, making it more efficient in computation. Second, it estimates camera poses in a global optimization framework which mitigates the drifting issue of the incremental pipelines. Despite the popularity of global SfM pipelines, the estimation of global camera poses (e.g., orientations) remains a highly challenging problem. For instance, estimating camera orientations from their relative measurements, often called rotation synchronization, is a highly nonconvex graph optimization problem. In typical scenarios of highly noisy or corrupted measurements

of relative orientations, many common solutions of rotation synchronization have poor accuracy and slow convergence.

Given these challenges, theoretical developments have demonstrated the critical role of cycle-consistency information in inferring corrupted measurements [25]. In practice, the consistency constraint on 3-cycles was utilized to estimate the error of each measured relative orientation. It also helped nonconvex iterative rotation synchronization solvers avoid spurious local minima and achieve significantly higher accuracy [39]. However, the usage of 3-cycles largely limits the application of these improved algorithms to other important scenarios. One scenario involves a sparse viewing graph lacking sufficient 3-cycles. This often occurs when the size of the graph is too large to densely measure the relative orientations on its edges, which could happen in certain case for the molecular orientation estimation in cryo-electron microscopy imaging. Another scenario is orientation estimation for each piece of jigsaw puzzles, where the graph is a 2D lattice and 3-cycle does not exist. Lastly, in distributed SfM, edges between any two clusters of nodes form a bipartite graph, and cycles of odd length do not exist. Our numerical results primarily emphasize the practical scenario of distributed SfM, which holds particular relevance for the broader computer vision community.

Despite the multiple critical applications of long-cycle consistency, inferring measurement noise from long cycles is challenging in both computation and theory. First of all, the number of cycles grows exponentially with the cycle length, and measuring cycle inconsistencies for each long cycle is computationally intractable. Moreover, developing theoretical guarantees for long-cycle inference methods is fundamentally more difficult than the 3-cycle case. Indeed, in a random graph setting, a set of longer cycles are more likely to share common edges, making their consistency score highly correlated. Therefore, new tools are required to handle the correlated empirical process.

In this work, we propose the first practical method, LongSync, for inferring edge corruption levels from long cycle consistency information. For this purpose, we carefully modify and vectorize the Cycle Edge Message Passing (CEMP) method [25]. This nontrivial modification drastically reduces its computational complexity when using longer cycles. Moreover,

*This work was supported by NSF award DMS-2152766.
 Supplementary code: <https://github.com/sli743/LongSync>

by employing a more delicate analysis and incorporating new tools from probability theory and combinatorics, we develop a significantly stronger exact recovery result with a general cycle length under a popular probabilistic model. The sample complexity in our theory is the lowest among all practical rotation synchronization methods. Although we limit our scope to the application of distributed SfM, our algorithm and theory applies to any finite-dimensional linear group in group synchronization, and not just $SO(3)$ in rotation synchronization.

1.1. Related Work

Earlier rotation synchronization methods [2, 10–13, 18, 33, 42, 43] seek to minimize a least squares energy function. They can be described as relaxed versions of the maximum likelihood estimator under the additive Gaussian noise model, but they are not robust in the presence of outliers or heavy-tailed noise. Nevertheless, in the case of global SfM, the initially estimated relative camera rotations can be severely corrupted due to the erroneous keypoint matches and the subsequent poor estimation of fundamental matrices.

To handle outliers, robust rotation synchronization methods either minimize a robust energy function or reweigh/trim the viewing graph based on the corruption levels of the edges. Wang and Singer [49] minimizes a corresponding ℓ_1 objective function using semidefinite programming (SDP) relaxation, which is slow in practice. Other energy minimization methods are typically nonconvex, which include the Weiszfeld algorithm [20] and the Riemannian subgradient method [29] for ℓ_1 minimization, and the iteratively reweighted least squares (IRLS) for minimizing general ℓ_p [5] and Geman-McClure [4] loss functions. However, all these methods heavily rely on good initialization. Sidhartha and Govindu [41] partially remedy the issue using adaptive Geman-McClure loss functions, but their approach remains sensitive to the initialized weights. Maunu and Lerman [30] propose to solve rotation synchronization by an iterative robust averaging method that utilizes Tukey depth, but they have not demonstrated effective performance for real SfM applications. Arrigoni et al. [1] applies a low-rank and sparse matrix decomposition method to $SO(3)$ and $SE(3)$ synchronization, but it is even less robust to outliers than IRLS.

Instead of employing a robust objective function, Shen et al. [36] and Zach et al. [51] uses the 3-cycle consistency constraint to detect and remove corrupted relative orientations. Lerman and Shi [25] take one step further to estimate the corruption level of each relative measurement by a novel cycle-edge message passing (CEMP) algorithm. They then use the estimated corruption levels to reweigh the graph and solve rotation synchronization using a weighted least squares method. This message passing procedure was further combined with IRLS to boost its accuracy in [39]. Particular versions and extensions of this procedure for permutation and partial permutation synchronization, which are relevant to the matching component of SfM, were discussed in [27, 40].

However, all the previously mentioned cycle-based methods [25, 27, 36, 38–40, 51] only use 3-cycles in practice, limiting their application for distributed synchronization. Indeed, the standard distributed synchronization often requires “stitching” local solutions by synchronizing the relative rotations between clusters. Each of these inter-cluster rotations is estimated by “averaging” the edges between the two clusters. These edges form a bipartite graph, and the minimal cycle length is 4. As pointed in [51], the number of operations for computing long cycle consistency information scales exponentially with the cycle length. Therefore, none of the existing distributed rotation synchronization methods directly exploits long cycle information due to this computational challenge.

The earlier distributed methods for $SO(d)$ synchronization, such as [45] and [44], minimize a least squares energy and are not robust to outliers. A series of distributed SfM methods [14–16] implement incremental SfM algorithms for each cluster. However, these methods do not employ a standard rotation synchronization algorithm, as they require additional information such as the number of keypoint matches between images. Moreover, the incremental methods are slower since they require multiple rounds of global rotation synchronization. MultiSync [9] synchronizes the inter-cluster rotations directly using all inter-cluster edges among all clusters, by formulating a novel synchronization problem on a multi-graph. Although it utilizes a more unified formulation, its objective function is least squares which largely limits its robustness to outliers.

A recent and different type of methods for rotation synchronization use deep learning [21, 26, 34]. However, these methods are supervised and thus may not generalize well when switching datasets. Moreover, like many other previous methods, they lack theoretical guarantees.

A common theoretical setting to assess the performance of rotation synchronization algorithms is the uniform corruption model (UCM) described in §4. We provide the best sample complexity for LongSync, even with only 3 cycles, among all previously established estimates for the UCM model.

1.2. Contributions of This Work

- We propose the first practical algorithm that infers edge corruption levels from long cycle consistency information. The computation complexity of our method is reduced from $O(n^c)$ to $O(n^3)$ (or possibly $O(n^{2.373})$) for cycle length $c \leq 6$ and $O(n^{\lfloor (c+3)/2 \rfloor})$ for $c > 6$.
- We establish sample complexity estimates for our method under the uniform corruption model, where we get closer to the information theoretic bound than any other existing work. Our proof requires delicate analysis and it also improves previous estimates for the CEMP algorithm.
- We introduce a new graph partition and graph preprocessing method that utilizes our inference method, and demonstrate the effectiveness of our pipeline in boosting the performance of distributed synchronization.

- Extensive numerical experiments demonstrate the outstanding performance of our method.

2. Problem Formulation and Preliminaries

Assume a graph $G = ([n], E)$ where $[n]$ is the set of nodes indexed by $\{1, 2, \dots, n\}$ and E is the set of edges. Given a mathematical group \mathcal{G} , each graph node is assigned an underlying ground truth group element \mathbf{R}_i^* , where $\mathbf{R}_i^* \in \mathcal{G}$ and we use star superscript to emphasize the ground truth. For each edge $ij \in E$, we observe a relative group ratio $\mathbf{R}_{ij} \in \mathcal{G}$, whose clean counterpart is $\mathbf{R}_{ij}^* = \mathbf{R}_i^* \mathbf{R}_j^{*-1}$. **Group synchronization** aims to recover the ground truth group elements $\{\mathbf{R}_i^*\}_{i \in [n]}$ from the possibly noisy and corrupted measurements $\{\mathbf{R}_{ij}\}_{ij \in E}$. In this paper, we focus on the case of rotation synchronization, which is a special case of group synchronization with $\mathcal{G} = \text{SO}(d)$. For applications in camera orientation synchronization ($d = 3$), we estimate absolute rotations for each node $i \in [n]$ from measured relative rotations of edges in E . Note that since $\{\mathbf{R}_i^*\}_{i \in [n]}$ and $\{\mathbf{R}_i^* \mathbf{R}_0^*\}_{i \in [n]}$ generate the same set of relative rotations, one can only estimate $\{\mathbf{R}_i^*\}_{i \in [n]}$ up to a global rotation. The generalization to any linear groups is discussed in the supplementary material.

2.1. Notations and Definitions

We denote the adjacency matrix of graph G as \mathbf{A} , and form a pairwise observation matrix $\mathbf{R} \in \mathbb{R}^{dn \times dn}$ by stacking the \mathbf{R}_{ij} 's (for $ij \notin E$, set $\mathbf{R}_{ij} = \mathbf{0}_{3 \times 3}$):

$$\mathbf{R} := \begin{pmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} & \cdots & \mathbf{R}_{1n} \\ \mathbf{R}_{21} & \mathbf{R}_{22} & \cdots & \mathbf{R}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{R}_{n1} & \mathbf{R}_{n2} & \cdots & \mathbf{R}_{nn} \end{pmatrix}.$$

We list the matrix operations used in the paper. For matrices \mathbf{X} and \mathbf{Y} , the operations $\mathbf{X} \otimes \mathbf{Y}$, $\mathbf{X} \odot \mathbf{Y}$, $\mathbf{X} \oslash \mathbf{Y}$ respectively denotes the Kronecker product, Hardmard (element-wise) multiplication and Harmard division between \mathbf{X} and \mathbf{Y} . $\mathbf{X}^{\odot k}$ denotes the element-wise matrix k -power. For block matrices, $\langle \mathbf{X}, \mathbf{Y} \rangle_{\text{block}}$ denotes the blockwise inner product of \mathbf{X} and \mathbf{Y} , i.e. $\langle \mathbf{X}, \mathbf{Y} \rangle_{\text{block}}(i, j) = \langle \mathbf{X}[i, j], \mathbf{Y}[i, j] \rangle$, where $[i, j]$ refers to the corresponding block of the matrix.

2.2. Review of CEMP for c -Cycles

We assume the above setting of $\text{SO}(d)$ synchronization. Let \mathcal{D} be any bi-invariant metric on $\text{SO}(d)$. We assume a fixed number of cycles, c , and denote by N_{ij}^c the set of simple cycles of length c (or simple c -cycles) containing edge ij . CEMP [25] aims to estimate for each edge ij the corruption level

$$s_{ij}^* = \mathcal{D}(\mathbf{R}_{ij}, \mathbf{R}_{ij}^*), \quad (1)$$

from the set of cycle inconsistency measures

$$d_L = \mathcal{D}(\mathbf{R}_L, \mathbf{R}_{ij}^*) \quad (2)$$

where cycle $L = (ik_1, k_1k_2, \dots, k_{c-2}j, ji) \in N_{ij}^c$ and $\mathbf{R}_L := \mathbf{R}_{ik_1} \mathbf{R}_{k_1k_2} \cdots \mathbf{R}_{k_{c-2}j}$. The estimated s_{ij}^* can then be used for extracting a clean subgraph, or to implement a weighted least squares solver where higher weights are assigned to cleaner edges.

It is obvious that if all the edges in L are clean then $d_L = 0$. Moreover, due to bi-invariance of \mathcal{D} , the following holds true

$$d_L = s_{ij}^* \text{ whenever } L \in G_{ij}^c, \quad (3)$$

where G_{ij}^c is the set of good c -cycles with respect to ij , i.e. the set of cycles $L \in N_{ij}^c$ such that $ik_1, \dots, k_{c-2}j$ are clean. This gives a sufficient condition for d_L to be an exact estimator of s_{ij}^* .

To estimate the corruption levels of each edge ij , CEMP initializes the edge weight of each $ij \in E$ as $w_{ij}^{(0)} = 1$. It then iteratively updates the corruption level estimate as the following convex combination of d_L 's:

$$s_{ij}^{(t)} = \sum_{L \in N_{ij}^c} w_L^{(t)} d_L / z_{ij}^{(t)} \quad (4)$$

where $z_{ij}^{(t)} = \sum_{L \in N_{ij}^c} w_L^{(t)}$. The cycle weights $w_L^{(t+1)}$ are computed from the edge weights $w_e^{(t+1)} = e^{-\beta_t s_e^{(t)}}$:

$$w_L^{(t+1)} = \prod_{e \in L \setminus \{ij\}} w_e^{(t+1)} = \prod_{e \in L \setminus \{ij\}} e^{-\beta_t s_e^{(t)}}, \quad (5)$$

so that $w_L^{(t+1)}$ focuses on good cycles. The cycle weights and edge corruption levels are alternately updated and improved from each other. Interestingly, it is proved in [25] under two different corruption models that CEMP converges linearly to the ground truth corruption estimates under mild conditions for $c=3$. In practice, CEMP only uses 3-cycles for consideration of efficiency. For longer cycles, the complexity of CEMP scales exponentially with the cycle length c (which is discussed in §3.2), and the convergence guarantee of CEMP remains unknown.

3. Our method: LongSync

3.1. LongSync: Modification of CEMP

Our goal is to develop a scalable variant of CEMP for any fixed number of cycles, $c \geq 3$. The main computational bottleneck of step (4) in CEMP is that computing and summing the cycle inconsistency measures takes $\sum_{ij \in E} |N_{ij}^c| = O(n^c)$ operations and memory. Therefore, to develop a scalable algorithm, we aim to take weighted average over d_L without explicitly computing and storing each d_L . To achieve this, we propose the following specification and modification on CEMP:

- **Use Chordal distance on $\text{SO}(d)$.** We suggest the distance function

$$\mathcal{D}(\mathbf{R}_1, \mathbf{R}_2) = \sqrt{1 - \langle \mathbf{R}_1, \mathbf{R}_2 \rangle} / d.$$

This distance is proportional to the Chordal distance on $\text{SO}(3)$, which is the Euclidean distance between two rotations embedded in $\mathbb{R}^{d \times d}$.

- **Use weighted quadratic average for corruption level update.** Instead of updating the corruption level estimates by a weighted average of d_L , we use the weighted quadratic average of d_L , namely

$$s_{ij}^{(t)} = \sqrt{\sum_{L \in N_{ij}^c} w_L^{(t)} d_L^2 / z_{ij}^{(t)}} \quad (6)$$

where the update rule of cycle weights remains the same:

$$w_L^{(t+1)} = \prod_{e \in L \setminus \{ij\}} w_e^{(t+1)} = \prod_{e \in L \setminus \{ij\}} e^{-\beta_t s_e^{(t)}}. \quad (7)$$

As a result, $d_L^2 = d^2(\mathbf{R}_L, \mathbf{R}_{ij}) = 1 - \langle \mathbf{R}_L, \mathbf{R}_{ij} \rangle / d$ is linear in both \mathbf{R}_L and \mathbf{R}_{ij} . Therefore one can switch the order of d^2 and the weighted summation, so that the computation of $s_{ij}^{(t)}$ can be vectorized. Indeed, by this linearity and equations (4) and (2), and note that $z_{ij}^{(t)} = \sum_{L \in N_{ij}^c} w_L^{(t)}$, we obtain the following equation:

$$\begin{aligned} s_{ij}^{(t)} &= \left(\sum_{L \in N_{ij}^c} w_L^{(t)} d_L^2 / z_{ij}^{(t)} \right)^{1/2} \\ &= \left(\sum_{L \in N_{ij}^c} w_L^{(t)} \mathcal{D}^2(\mathbf{R}_L, \mathbf{R}_{ij}) / z_{ij}^{(t)} \right)^{1/2} \\ &= \left(\mathcal{D}^2 \left(\sum_{L \in N_{ij}^c} w_L^{(t)} \mathbf{R}_L, \mathbf{R}_{ij} \right) / z_{ij}^{(t)} \right)^{1/2} \\ &= \left(1 - \left\langle \sum_{L \in N_{ij}^c} w_L^{(t)} \mathbf{R}_L, \mathbf{R}_{ij} \right\rangle / d \sum_{L \in N_{ij}^c} w_L^{(t)} \right)^{1/2} \end{aligned} \quad (8)$$

Equation (8) can be vectorized using the trick of matrix power if we allow repeated nodes for each cycle. That is, one can stack the $s_{ij}^{(t)}$'s and $w_{ij}^{(t)}$'s into matrices $\mathbf{S}^{(t)}$ and $\mathbf{W}^{(t)}$, and vectorize (8) as

$$\mathbf{S}^{(t)} = \left(\mathbf{A} - \left\langle \left(\mathbf{W}^{(t)} \otimes \mathbf{1}_d \odot \mathbf{R} \right)^{c-1}, \mathbf{R} \right\rangle_{\text{block}} \odot d \left(\mathbf{W}^{(t)} \right)^{c-1} \right)^{\odot 1/2}. \quad (9)$$

Indeed, by (7) and the definition of \mathbf{R}_L , $\sum_{L \in C_{ij}^c} w_L^{(t)} \mathbf{R}_L$ is the ij -th block of $(\mathbf{W}^{(t)} \otimes \mathbf{1}_d \odot \mathbf{R})^{c-1}$, and $\sum_{L \in C_{ij}^c} w_L^{(t)}$ is the ij -th element of $\mathbf{W}^{(t)c-1}$, where C_{ij}^c is the set of c -cycles containing ij with possibly repeated nodes.

In the case of utilizing only simple cycles, (8) and (9) are not equivalent and we need to correct (9) to remove the cycles with repeated nodes, so that only simple cycles in N_{ij}^c remain. Let $g_c(\mathbf{W}, \mathbf{R})$ be the matrix valued function where $g_c(\mathbf{W}, \mathbf{R})(i, j) = \sum_{L \in N_{ij}^c} \left(\prod_{e \in L \setminus \{ij\}} w_e^{(t)} \right) \mathbf{R}_L$. Let $f_c(\mathbf{W})$ be the matrix valued function where $f_c(\mathbf{W})(i, j) = \sum_{L \in N_{ij}^c} \left(\prod_{e \in L \setminus \{ij\}} w_e^{(t)} \right)$. The following result holds:

Proposition 3.1. *The update rule of LongSync (8) is equivalent to the following matrix equations:*

$$\mathbf{S}^{(t)} = \left(\mathbf{A} - \left\langle h_c(\mathbf{W}^{(t)}, \mathbf{R}), \mathbf{R} \right\rangle_{\text{block}} \odot \mathbf{A} \right)^{\odot 1/2}, \quad (10)$$

where $\mathbf{W}^{(t+1)} = \mathbf{A} \odot \exp(-\beta_t \mathbf{S}^{(t)})$ and

$$h_c(\mathbf{W}^{(t)}, \mathbf{R}) := g_c(\mathbf{W}^{(t)}, \mathbf{R}) \odot (d \cdot f_c(\mathbf{W}^{(t)}) \otimes \mathbf{1}_d).$$

Here \exp denotes the elementwise exponential function.

We use equation (10) as the update rule of LongSync and propose the vectorized LongSync algorithm in algorithm 1.

Algorithm 1 (LongSync)

Input: pairwise rotation matrix $\mathbf{R} \in \mathbb{R}^{dn \times dn}$, adjacency matrix $\mathbf{A} \in [0, 1]^{n \times n}$, cycle length c , positive parameters $\{\beta_t\}_{t \geq 1}$, time step T

$\mathbf{W}^{(0)}(i, j) \leftarrow \mathbf{A}$

for $t = 0 : T$ **do**

$$\mathbf{S}^{(t)} \leftarrow \left(\mathbf{A} - \left\langle h_c(\mathbf{W}^{(t)}, \mathbf{R}), \mathbf{R} \right\rangle_{\text{block}} \odot \mathbf{A} \right)^{\odot 1/2} \quad (11)$$

$$\mathbf{W}^{(t+1)} \leftarrow \mathbf{A} \odot \exp(-\beta_t \mathbf{S}^{(t)}) \quad (12)$$

end for

Output: edge weights $\mathbf{W}^{(T+1)}$, corruption levels $\mathbf{S}^{(T)}$

We claim that g_c and f_c can be computed with a sequence of matrix operations, thus greatly reducing the time and space consumption of LongSync compared to its original form. For $c \leq 6$, the time complexity of computing g_c and f_c is $O(r(dn))$, where $r(n)$ is the complexity for multiplying two $n \times n$ matrices; for $c \geq 7$ the time complexity is at most $O(n^{\lfloor (c+3)/2 \rfloor})$. This claim is proved in the supplementary material. We list the formula for g_c and f_c for $c = 3, 4, 5, 6$ inspired by [35, 46]. The formula for $c = 6$ is moved to the supplementary material due to the space limit. For $c \geq 7$ the formula becomes extremely complicated. We remark that in practice, the cycles of length greater than 6 are often not used.

We finally remark that although Algorithm 1 only utilizes cycles of a fixed length, one can easily generalize it to incorporate cycles of different lengths. Indeed, the equation (10) could use a convex combination of h_c 's that corresponds to different values of c . That is, for a preselected set of cycle lengths C , the equation (11) in Algorithm 1 is replaced by

$$\mathbf{S}^{(t)} = \left(\mathbf{A} - \left\langle \sum_{c \in C} \lambda_c h_c(\mathbf{W}^{(t)}, \mathbf{R}), \mathbf{R} \right\rangle_{\text{block}} \odot \mathbf{A} \right)^{\odot 1/2} \quad (13)$$

where the coefficients λ_c satisfies $\sum_{c \in C} \lambda_c = 1$ to ensure a convex combination. Here each λ_c is user-specified to reflect the importance of the cycles of length c . However, the optimal

choice of these parameters under certain statistical model remains unclear.

For simplicity, in the experiments we only use a fixed cycle length to avoid choosing λ_c . We have observed that such simple choice still yields satisfying accuracy in camera orientation estimation on both synthetic and real data. We refer the readers to §5 and 6 for more details.

c	Formula of $f_c(\mathbf{W})$	Formula of $g_c(\mathbf{W}, \mathbf{R})$
3	\mathbf{W}^2	\mathbf{P}^2
4	$\mathbf{W}^3 - \text{d}(\mathbf{W}^2)\mathbf{W}$ $-\mathbf{W}\text{d}(\mathbf{W}^2) + \mathbf{W}^{\odot 3}$	$\mathbf{P}^3 - \text{d}_{\text{block}}(\mathbf{P}^2)\mathbf{P}$ $-\mathbf{P}\text{d}_{\text{block}}(\mathbf{P}^2) + \mathbf{P}^{\odot 3}$
5	$\mathbf{W}^4 - \text{d}(\mathbf{W}^3)\mathbf{W}$ $-\text{d}(\mathbf{W}^2)\mathbf{W}^2$ $-\mathbf{W}^2\text{d}(\mathbf{W}^2)$ $-\mathbf{W}\text{d}(\mathbf{W}^2)\mathbf{W}$ $+3\mathbf{W}^{\odot 2}\mathbf{W}^2$ $+\mathbf{W}\mathbf{W}^{\odot 3} + \mathbf{W}^{\odot 3}\mathbf{W}$	$\mathbf{P}^4 - \text{d}_{\text{block}}(\mathbf{P}^3)\mathbf{P}$ $-\text{d}_{\text{block}}(\mathbf{P}^2)\mathbf{P}^2$ $-\mathbf{P}^2\text{d}_{\text{block}}(\mathbf{P}^2)$ $-\mathbf{P}\text{d}_{\text{block}}(\mathbf{P}^2)\mathbf{P}$ $+3\mathbf{P}^{\odot 2}\mathbf{P}^2$ $+\mathbf{P}\mathbf{P}^{\odot 3} + \mathbf{P}^{\odot 3}\mathbf{P}$

Table 1. Formulas for f_c and g_c . Here we let $\mathbf{P} = (\mathbf{W} \otimes 1_d) \odot \mathbf{R}$ for shorter notation. $\text{d}(\mathbf{X})$ returns the diagonal of matrix \mathbf{X} , $\text{d}_{\text{block}}(\mathbf{X})$ returns the diagonal block matrix from the $d \times d$ diagonal blocks of matrix \mathbf{X} .

3.2. Computational Complexity

We derive the space and time complexity for LongSync, and demonstrate its advantages over CEMP. The initialization step involves setting the weights of all edges to 1, which takes time $O(|E|)$ and space $O(n^2)$. For each iteration, LongSync updates the matrices $\mathbf{S}^{(t)}$ and $\mathbf{W}^{(t)}$ with equations (11) and (12), respectively. Computing $\mathbf{W}^{(t+1)}$ involves two matrix subtractions, one scalar-matrix multiplication and one element-wise matrix exponential operation on $\mathbf{S}^{(t)} \in \mathbb{R}^{n \times n}$. Therefore the update step (12) takes at most $O(n^2)$ time and space. Equation (11), on the other hand, involves a sequence of matrix operations on $\mathbf{P}^{(t)} = (\mathbf{W}^{(t)} \otimes 1_d) \odot \mathbf{R} \in \mathbb{R}^{dn \times dn}$ and $\mathbf{W}^{(t)} \in \mathbb{R}^{n \times n}$, including matrix multiplications, element-wise multiplications and diagonal block selections. Computing $\mathbf{P}^{(t)}$ takes $O(d^2 n^2)$ memory and $O(d^2 n^2)$ time. The matrix operations on $\mathbf{P}^{(t)}$ take $O(K_c d^3 n^3)$ time and $O(d^2 n^2)$ space, and the matrix operations on $\mathbf{W}^{(t)}$ take $O(K_c n^3)$ time and $O(n^2)$ space, where K_c is the number of terms in the equation for f_c and g_c . Therefore, each iteration over t takes $O(K_c d^3 n^3)$ time and $O(d^2 n^2)$ space. To sum up, for LongSync, the time complexity is $O(K_c d^3 n^3)$ and the space complexity is $O(d^2 n^2)$. For $c=3,4,5,6$, the number K_c is equal to 1,3,9,32.

In comparison, we consider the initialization step of CEMP. For each edge $ij \in E$ and $L \in N_{ij}^c$, initializing CEMP involves computing and storing all the cycle inconsistency measures d_L using equation (2). For each $L \in N_{ij}^c$, computing d_L involves multiplying c rotations, which takes $O(cd^3)$ time and $O(d^2)$ space. This step is repeated for each $ij \in E$ and $L \in N_{ij}^c$, therefore the total time complexity is $O(cd^3 \sum_{ij \in E} |N_{ij}^c|)$ and

the total space complexity is $O(d^2 \sum_{ij \in E} |N_{ij}^c|)$. Since for each edge there are $(n-2)(n-3)\dots(n-c+1) = O(n^{c-2})$ cycle candidates, we know that $|N_{ij}^c| \sim O(n^{c-2})$ for each $ij \in E$ in the worst case scenario of a dense graph. Therefore the initialization of CEMP takes $O(cd^3 n^{c-2} |E|)$ time and $O(d^2 n^{c-2} |E|)$ space in the worst case. Given $c \geq 4$ and $|E| \sim n^2$, CEMP requires much more time and space than LongSync.

4. Theory for Uniform Corruption Model

In this section, we present the exact recovery guarantee of LongSync under the uniform corruption model (UCM). UCM is a popular probabilistic model that is widely adopted for synthetic experiments of many previous works on group synchronization [8, 25, 28, 29, 32, 42]. The model $\text{UCM}(n, p, q_g)$ assumes that G is an Erdős-Rényi graph with edge connection probability p . For each edge $ij \in E$, \mathbf{R}_{ij} is generated independently as follows:

$$\mathbf{R}_{ij} = \begin{cases} \mathbf{R}_{ij}^* & \text{w.p. } q_g; \\ \tilde{\mathbf{R}}_{ij} \sim \text{Haar}(\mathcal{G}) & \text{w.p. } 1 - q_g. \end{cases}$$

We also developed an exact recovery theory for a general model of adversarial corruption, which we include in section C.1 of the supplementary material. An informal version of our main result for UCM is stated in Theorem 4.1. Although the application of this paper is focused on rotation synchronization, the following theory for UCM is valid for any compact group \mathcal{G} , as explained in the supplementary material.

Theorem 4.1. *Let $0 < r < 1$, $0 < q < 1$, $0 < p \leq 1$, $\mathcal{G} = SO(3)$. Assume LongSync is applied with cycles of length c , $n/\log n \sim p^{-(c-1)/(c-2-\epsilon)} q_g^{-7(c-1)/3(c-2)}$ for some $\epsilon > 0$ and*

$$1/\beta_{t+1} = r/\beta_t \text{ for all } t \geq 1.$$

Then with appropriate choices of β_0, β_1, r , and high probability,

$$\max_{ij \in E} |s_{ij}^* - s_{ij}^{(t)}| \leq \frac{1}{2c\beta_t} \text{ for all } t \geq 1.$$

The major difficulty of proving Theorem 4.1 is the dependence in the cycle inconsistency measures for cycles in N_{ij}^c when $c \geq 4$. Unlike the 3-cycle case, the cycle inconsistency measure of a 4-cycle $L_1 = (ik_1, k_1k_2, k_2j)$ is correlated with that of $L_2 = (ik_1, k_1k_3, k_3j)$ under UCM. Therefore the key concentration inequalities for the proof cannot be concluded from the standard Chernoff bounds. To overcome this theoretical obstacle, we have integrated various mathematical techniques from [3, 7, 22–24, 48] to prove the theorem, whose details are included in the supplementary material.

Theorem 4.1 provides an upper bound of the sample complexity (the required graph size n) of LongSync for exact recovery of the ground truth solutions. This sample complexity is the closest to the information theoretic bound among all existing rotation synchronization methods. The comparison with previous works is summarized in Table 2.

Reference	Sample Complexity
[25] for CEMP	$O(p^{-2}q_g^{-28/3})$
[29] for ReSync	$O(p^{-2}q_g^{-7})$
Ours for CEMP	$O(p^{-2-\epsilon}q_g^{-14/3})$
Ours for LongSync ($c=3$)	$O(p^{-2-\epsilon}q_g^{-14/3})$
Ours for LongSync ($c=4$)	$O(p^{-1.5-\epsilon}q_g^{-3.5})$
Ours for LongSync (any c)	$O\left(p^{-\frac{c-1}{c-2-\epsilon}}q_g^{-\frac{7(c-1)}{3(c-2)}}\right)$
Ours for LongSync ($c \rightarrow \infty$)	$O(p^{-1-\epsilon}q_g^{-7/3})$
Information Theoretic Bound [6]	$O(p^{-1}q_g^{-2})$

Table 2. Comparison of the sample complexity requirement. Lower absolute values of the powers on p, q_g indicate better results. ϵ is an arbitrarily small positive real number.

5. Synthetic Data Experiment

We test LongSync on synthetic datasets generated with Uniform Corruption Model (UCM) and Uniform Bipartite Corruption Model (UBCM) respectively described in §5.1 and 5.2. For both models with their corresponding viewing graphs $G = ([n], E)$, we sample the ground truth absolute rotation matrices $\{\mathbf{R}_i^*\}_{i \in [n]}$ independently from the Haar measure on $SO(3)$, and we generate the observed relative rotations $\{\mathbf{R}_{ij}\}_{ij \in E}$ independently as follows:

$$\mathbf{R}_{ij} = \begin{cases} \mathbf{R}_i^* \mathbf{R}_j^* & \text{w.p. } q_g; \\ \tilde{\mathbf{R}}_{ij} \sim \text{Haar}(SO(3)) & \text{w.p. } 1 - q_g. \end{cases}$$

We use LongSync with cycle length c , $\beta_i = \min(2^i, 20)$ and $T = 10$ and record the edge weights. For UCM we set $c=4, 5$ and for UBCM we only use $c=4$ since no 5-cycles exist. For our method, we first build a weighted graph whose edge weights are estimated by LongSync. We then extract a maximum spanning tree (MST) of the resulting weighted graph. The resulting spanning tree is expected to be the cleanest possible spanning tree. To initialize our solution of absolute rotations, we first fix \mathbf{R}_1 as the identity rotation, and find the rest of \mathbf{R}_i 's by consecutively multiplying the relative rotations along the spanning tree using the formula $\mathbf{R}_i = \mathbf{R}_{ij} \mathbf{R}_j$. To refine our initialized solution, we apply IRLS with Geman-McClure [4] loss functions to minimize $\sum_{ij \in E} \rho_{GM}(d_{\angle}(\mathbf{R}_{ij}, \mathbf{R}_i \mathbf{R}_j^T))$, where d_{\angle} is the geodesic distance in $SO(3)$. We refer to this method as LongSync+IRLS.

To demonstrate the advantages of utilizing longer cycle information, we compare our method with IRLS initialized by other two different spanning trees. The first one is the random spanning tree, which uses no cycle information. The other one is the MST extracted from the CEMP-estimated weights. Note that CEMP only uses 3-cycle information. We refer to these methods as IRLS and CEMP+IRLS respectively.

Since the solution of absolute rotations is determined up to a global rotation, we align our estimated rotation $\{\hat{\mathbf{R}}_i\}$ with $\{\mathbf{R}_i^*\}$ by $\mathbf{R}_{\text{align}}$ that minimizes the ℓ_1 rotation alignment error

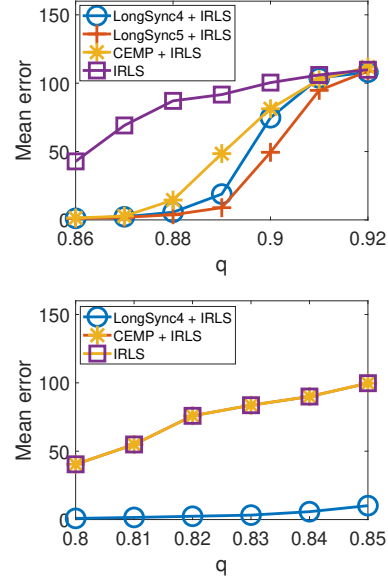


Figure 1. Average errors for IRLS, CEMP+IRLS and LongSync+IRLS with $c=4, 5$, using the uniform corruption (top) and uniform bipartite corruption (bottom) models. The mean errors are measured in degrees. LongSync4 and LongSync5 refer to LongSync with 4 and 5 cycles, respectively.

$\sum_{i \in [n]} \|\hat{\mathbf{R}}_i \mathbf{R}_{\text{align}} - \mathbf{R}_i^*\|_F$. We report the mean estimation error in degrees: $180 \cdot \sum_{i \in [n]} d_{\angle}(\hat{\mathbf{R}}_i \mathbf{R}_{\text{align}}, \mathbf{R}_i^*)/n$.

5.1. Uniform Corruption Model

For $\text{UCM}(n, p, q_g)$, we take $n=200$ and $p=1$ and corruption probability $q = 1 - q_g$ ranges from 0.86 to 0.92. We report the average mean estimation error from 20 trials of the uniform corruption model in the top panel of Figure 1.

We note that LongSync uniformly improves IRLS, and the mean error of LongSync decreases as the cycle length increases. When $q > 0.86$, the expected number of clean 3-cycles for each edge is less than 4, and therefore longer cycles are helpful. The numerical result aligns with our theory that using longer cycles may tolerate higher corruption with fixed graph size n .

5.2. Uniform Bipartite Corruption Model

For UBCM, we first generate the graph and relative rotations by $\text{UCM}(n, p, q_g)$ with $n=200$, $p=1$ and $q = 1 - q_g$ ranging from 0.8 to 0.85. Then we split the nodes into two clusters of equal size and remove the intra-cluster edges for both clusters. The resulting graph is bipartite, where only cycles of even lengths exist. We report the mean estimation error from 20 trials in the bottom panel of Figure 1.

We observe that LongSync with 4-cycles almost exactly recovers the rotations, while for other algorithms the rotation estimates are not even close to the ground truth.

6. Real Data Experiment

We test distributed synchronization with LongSync on the PhotoTourism dataset [50] to demonstrate its advantages in accuracy and speed over other baselines. PhotoTourism is a large scale dataset consisting of 15 sets of images taken for 3D reconstruction. The smallest dataset consists of 247 cameras, and the largest dataset consists of 5433 cameras. The input graph and initial pairwise rotation estimates are provided in the dataset. In the following, we first explain the common steps for distributed synchronization, and our improvement using LongSync. We then describe our graph processing method for filtering bad nodes and edges, which also is applicable to other baseline methods.

Steps in Distributed Synchronization:

1. **Graph partitioning.** The first step involves partitioning the graph $G = ([n], E)$ into K clusters $G_i = (V_i, E_i), i \in [K]$. In this paper we apply spectral clustering algorithm [37] on the adjacency matrix G , where $K = 0.6\sqrt{np}$ and $p = 2|E|/(n(n-1))$.
2. **Synchronization within clusters.** Run standard synchronization solvers for each cluster. In this work, we use the current state-of-the-art method MPLS [39]. Note that for each camera p in cluster k , one can only estimate the true rotation \mathbf{R}_p^* up to a global rotation \mathbf{R}_k . Namely, one only obtains $\hat{\mathbf{R}}^p \approx \mathbf{R}_p^* \mathbf{R}_k^{-1}$ where \mathbf{R}_k is unknown and is the same for all cameras in cluster k .
3. **Estimation of inter-cluster rotations.** To find \mathbf{R}_p^* of all cameras, one needs to solve \mathbf{R}_k for all clusters. Namely, one needs to rotate and stitch the solutions of all clusters so that they are in the same reference frame. To do this, it is common to first estimate the inter-cluster rotations $\mathbf{R}_{kl} := \mathbf{R}_k \mathbf{R}_l^{-1}$ between pairs of clusters k, l , and then synchronize these relative rotations. To estimate each \mathbf{R}_{kl} , we note that $\mathbf{R}_{kl} = \mathbf{R}_p^{*-1} \mathbf{R}_{pq}^* \mathbf{R}_q^*$ for each $p \in V_k$ and $q \in V_l$. Therefore, one can use the rotations in the set $S_{kl} := \{\hat{\mathbf{R}}_p^{-1} \mathbf{R}_{pq} \hat{\mathbf{R}}_q\}_{p \in V_k, q \in V_l}$ to approximate \mathbf{R}_{kl} . We remark that this step is crucial to the overall performance of distributed methods, and we compare the following methods for solving \mathbf{R}_{kl} :
 - **MultiSync [9]:** Run synchronization on a multi-graph where each edge kl is assigned a set of relative rotations $\{\hat{\mathbf{R}}_p^{-1} \mathbf{R}_{pq} \hat{\mathbf{R}}_q\}_{p \in V_k, q \in V_l}$. This combines the step 3 and 4 in a unified least squares formulation.
 - **Edge averaging using IRLS:** We initialize $\hat{\mathbf{R}}_{kl}$ with the quaternion ℓ_2 mean of the set S_{kl} and refine it using ℓ_1 -rotation averaging [19]. We refer to this method as IRLS in our comparison.
 - **Edge averaging using LongSync:** We first perform LongSync with 4-cycles to estimate the weights of these inter-cluster edges (there are no 3-cycles for a bipartite graph). We next initialize $\hat{\mathbf{R}}_{kl}$ as the quaternion weighted ℓ_2 mean of S_{kl} , using the edge weights from LongSync by their LongSync weights. Lastly, we refine the solution using [19].

4. **Synchronization of inter-cluster rotations.** This step is skipped for MultiSync. For other methods described in step 3, we find \mathbf{R}_k (up to a rotation) for each cluster k from the estimated $\{\mathbf{R}_{kl}\}_{k, l \in [K]}$ by MPLS.

5. **Rotation merging.** Finally, for each camera p in cluster k , the rotation estimate of p is given by $\mathbf{R}_p^{\text{final}} = \hat{\mathbf{R}}_p \mathbf{R}_k^{-1}$.

Next, we introduce our graph processing method to further boost the performance of all tested methods.

Extra Improvement by Graph Processing:

- **Spectral clustering with Jaccard Index.** For step 1, we use the Jaccard index matrix as the similarity matrix for spectral clustering, instead of the adjacency matrix. The $n \times n$ Jaccard index matrix \mathbf{A}_J is defined as follows:

$$\mathbf{A}_J(i, j) = \begin{cases} 0 & ij \notin E \\ \frac{|N_i \cap N_j|}{|N_i \cup N_j|} & ij \in E \end{cases} \quad (14)$$

where N_i and N_j denote the sets of neighboring nodes of node i and j , respectively. In this way, $\mathbf{A}_J(i, j)$ is higher for the pair ij contained in many 3-cycles, which is a more robust and nicely scaled statistics ($\in [0, 1]$) for measuring the local graph density around edge ij .

- **Refinement of intra-cluster edges and nodes.** For step 2, after the MPLS step, we perform CEMP with 3-cycles to estimate the corruption level of the intra-cluster edges for each cluster. We remove a camera if the number of neighboring ‘good’ edges, i.e. the edges with corruption level less than 0.1, is less than 4. The numbers 4 and 0.1 are chosen to balance the number of remaining cameras and the quality of intra-cluster rotation estimates. In order to eliminate the sparsely connected components inside the cluster, we use the Matlab built-in hierarchical spectral clustering function on the remaining cameras with the ‘cutoff’ and ‘depth’ parameters as 2 and 4, and we keep the largest component. The absolute rotations for the remaining cameras are estimated by MPLS. We remark that one could replace CEMP by LongSync with 3-cycles. However, we have not observed significant difference in the performance.

We respectively name MultiSync and IRLS with our new graph processing method as MultiSync(New) and IRLS(new). ‘‘LongSync’’ in our experiment refers to the full version of our algorithm: use LongSync weights for edge averaging in Step 3, with the graph processing step. We also compare with MPLS on the whole dataset, since it is a state-of-the-art non-distributed method, but we note that MPLS is significantly slower than all distributed methods. We report median error $180 \cdot \text{median}(\{d_{\angle}(\hat{\mathbf{R}}_i \mathbf{R}_{\text{align}}, \mathbf{R}_i^*)\}_{i \in [n]})$ of the tested methods on 14 datasets in Figure 2. We exclude the result of Gendarmenmarkt since all methods return large estimation errors in the figure. The full results, including that of mean error are included in the supplementary material.

In Figure 3, for each distributed method, we report the ratio (in percentage) between its total runtime on all datasets and that of the non-distributed MPLS. Namely, we compute




Figure 2. Median error for rotations for each dataset measured in degrees.




Figure 3. Runtime ratio and average median and mean error gaps between the distributed methods and MPLS on the entire graph.

$\sum_{d \in D} t_{\text{dist},d} / \sum_{d \in D} t_{\text{MPLS},d}$, where D is the set of 15 datasets, and each $t_{\text{dist},d}$ and $t_{\text{MPLS},d}$ is respectively the runtime of the distributed method and MPLS on data d . In the same figure, we present the mean/median error gap between each distributed method and MPLS. The mean and median error gap is respectively defined as $(\bar{e}_{\text{dist}} - \bar{e}_{\text{MPLS}}) / \bar{e}_{\text{dist}}$ and $(\hat{e}_{\text{dist}} - \hat{e}_{\text{MPLS}}) / \hat{e}_{\text{dist}}$, where \bar{e} and \hat{e} respectively denote the mean and median error over all cameras.

From Figure 2 and 3, our method outperforms other distributed methods on 13 out of 15 datasets. The most significant improvement in mean error and median error are respectively 28.6% and 46.4% (in Notre Dame) compared to the best performing method between IRLS(new) and MultiSync(new). The improvement is even more significant when comparing to the original version of these baseline methods without our graph processing method. The only two datasets without improvement are Gendarmenmarkt and Union Square. Our method is comparable to others on Union Square, and all methods return large errors on Gendarmenmarkt due to many repetitive patterns in its 3D scene.

The average mean and median error gap between our method and full MPLS are respectively 38.3% and 18.4%. Compared to the best performing method among others, our method reduces the average median error gap by 40.8%, and the average mean error gap by 9.6%. In terms of runtime, our method is uniformly faster than MultiSync and it is scalable on the largest dataset, taking less than 6% of the total runtime of full MPLS. In conclusion, our method significantly improves the result of distributed rotation synchronization without compromising runtime.

In the supplementary material, we further demonstrate the improvement by our new graph processing method, which significantly improves the results of LongSync (without extra graph processing) in 14 of the 15 datasets. On these 14 datasets, the average reduction on mean error is 59.2% and the average reduction on median error is 28.5%.

7. Conclusion

We propose LongSync, a robust and efficient algorithm for group synchronization. It modifies and vectorizes CEMP which enables efficient computation when using longer cycles. The theory we developed for LongSync is the strongest among all other existing results under UCM. Experiment shows that LongSync, together with our improved graph preprocessing method, achieves superior accuracy for distributed synchronization on large real datasets with competitive runtime. However, our method also has some limitations. First of all, in theory there is still a small gap of sample complexity from our method to the information theoretic one. Filling this gap is an open problem, which requires new tools and possibly more sophisticated analysis. Second, our graph preprocessing method is quite heuristic, and an automatic way of choosing parameters is needed. Our work also opens a door for some important future directions, including distributed partial permutation synchronization for multi-image matching, angular synchronization for Cryo-EM and Jigsaw Puzzles, and analysis of their algorithms.

References

- [1] Federica Arrigoni, Beatrice Rossi, Pasqualina Fragneto, and Andrea Fusiello. Robust synchronization in $so(3)$ and $se(3)$ via low-rank and sparse matrix decomposition. *Computer Vision and Image Understanding*, 174:95–113, 2018. [2](#)
- [2] A. S. Bandeira, N. Boumal, and A. Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *arXiv preprint arXiv:1411.3272*, 2014. [2](#)
- [3] Olivier Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *C. R. Math. Acad. Sci. Paris*, 334(6):495–500, 2002. [5](#), [14](#), [16](#), [21](#)
- [4] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, pages 521–528, 2013. [2](#), [6](#)
- [5] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):958–972, 2018. [2](#)
- [6] Yuxin Chen, Changho Suh, and Andrea J. Goldsmith. Information recovery from pairwise measurements. *IEEE Trans. Inf. Theory*, 62(10):5881–5905, 2016. [6](#)
- [7] Keresztély Corradi and András Hajnal. On the maximal number of independent circuits in a graph. *Acta Mathematica Hungarica*, 14(3-4):423–439, 1963. [5](#), [14](#)
- [8] Mihai Cucuringu. Synchronization over Z_2 and community detection in signed multiplex networks with constraints. *J. Complex Networks*, 3(3):469–506, 2015. [5](#)
- [9] Andrea Porfiri Dal Cin, Luca Magri, Federica Arrigoni, Andrea Fusiello, and Giacomo Boracchi. Synchronization of group-labelled multi-graphs. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6433–6443, 2021. [2](#), [7](#)
- [10] Frank Dellaert, David M. Rosen, Jing Wu, Robert E. Mahony, and Luca Carlone. Shonan rotation averaging: Global optimality by surfing $so(p)^n$. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VI*, volume 12351 of *Lecture Notes in Computer Science*, pages 292–308. Springer, 2020. [2](#)
- [11] Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality, 2017.
- [12] Anders P. Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 127–135. Computer Vision Foundation / IEEE Computer Society, 2018.
- [13] Anders P. Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging with the chordal distance: Global minimizers and strong duality. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(1):256–268, 2021. [2](#)
- [14] Xiang Gao, Hainan Cui, Menghan Li, Zexiao Xie, and Shuhan Shen. Irav3: Hierarchical incremental rotation averaging on the fly. *IEEE Trans. Circuits Syst. Video Technol.*, 33(4):2001–2006, 2023. [2](#)
- [15] Xiang Gao, Hainan Cui, and Shuhan Shen. Incremental rotation averaging revisited and more: A new rotation averaging benchmark. *arXiv preprint, abs/2309.16924*, 2023.
- [16] Xiang Gao, Lingjie Zhu, Hainan Cui, Zexiao Xie, and Shuhan Shen. IRA++: distributed incremental rotation averaging. *IEEE Trans. Circuits Syst. Video Technol.*, 32(7):4885–4892, 2022. [2](#)
- [17] Thomas Goldstein, Paul Hand, Choongbum Lee, Vladislav Voroninski, and Stefano Soatto. Shapefit and shapekick for robust, scalable structure from motion. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VII*, pages 289–304, 2016. [1](#)
- [18] V.M. Govindu. Lie-algebraic averaging for globally consistent motion estimation. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I, 2004. [2](#)
- [19] R. Hartley, K. Aftab, and J. Trunpf. L1 rotation averaging using the weiszfeld algorithm. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3041–3048. IEEE, 2011. [7](#)
- [20] Richard I. Hartley, Khurram Aftab, and Jochen Trunpf. L1 rotation averaging using the weiszfeld algorithm. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 3041–3048, 2011. [2](#)
- [21] Xiangru Huang, Zhenxiao Liang, Xiaowei Zhou, Yao Xie, Leonidas J. Guibas, and Qixing Huang. Learning transformation synchronization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [22] Svante Janson. Poisson approximation for large deviations. *Random Structures & Algorithms*, 1(2):221–229, 1990. [5](#), [14](#), [16](#)
- [23] Svante Janson. Large deviations for sums of partly dependent random variables. *Random Structures & Algorithms*, 24(3):234–248, 2004. [18](#), [21](#)
- [24] Jeong Han Kim and Van H Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000. [5](#), [14](#), [18](#)
- [25] Gilad Lerman and Yunpeng Shi. Robust group synchronization via cycle-edge message passing. *Foundations of Computational Mathematics*, 22(6):1665–1741, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [13](#), [14](#), [16](#), [19](#), [21](#)
- [26] Heng Li, Zhaopeng Cui, Shuaicheng Liu, and Ping Tan. RAGO: recurrent graph optimizer for multiple rotation averaging. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 15766–15775. IEEE, 2022. [2](#)
- [27] Shaohan Li, Yunpeng Shi, and Gilad Lerman. Fast, accurate and memory-efficient partial permutation synchronization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15735–15743, June 2022. [2](#)
- [28] Shuyang Ling. Near-optimal performance bounds for orthogonal and permutation group synchronization via spectral methods. *arXiv preprint arXiv:2008.05341*, 2020. [5](#)
- [29] Huikang Liu, Xiao Li, and Anthony Man-Cho So. Resync: Riemannian subgradient-based robust rotation synchronization. In *Advances in Neural Information Processing Systems: 37th Annual Conference on Neural Information Processing Systems*, 2023. [2](#), [5](#), [6](#)
- [30] Tyler Maunu and Gilad Lerman. Depth descent synchronization in $SO(D)$. *Int. J. Comput. Vis.*, 131(4):968–986, 2023. [2](#)
- [31] Onur Özyesil and Amit Singer. Robust camera location estimation by convex programming. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2674–2683, 2015. [1](#)
- [32] Deepti Pachauri, Risi Kondor, and Vikas Singh. Solving the multi-way matching problem by permutation synchronization. In

- C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 1860–1868. Curran Associates, Inc., 2013. 5
- [33] Alvaro Parra, Shin-Fang Ch'ng, Tat-Jun Chin, Anders P. Eriksson, and Ian Reid. Rotation coordinate descent for fast globally optimal rotation averaging. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 4298–4307. Computer Vision Foundation / IEEE, 2021. 2
- [34] Pulak Purkait, Tat-Jun Chin, and Ian Reid. Neurora: Neural robust rotation averaging, 2020. 2
- [35] Ian C Ross and Frank Harary. On the determination of redundancies in sociometric chains. *Psychometrika*, 17(2):195–208, 1952. 4, 12
- [36] Tianwei Shen, Siyu Zhu, Tian Fang, Runze Zhang, and Long Quan. Graph-based consistent matching for structure-from-motion. In *European Conference on Computer Vision*, pages 139–155. Springer, 2016. 2
- [37] Jianbo Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. 7
- [38] Yunpeng Shi and Gilad Lerman. Estimation of camera locations in highly corrupted scenarios: All about that base, no shape trouble. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 2868–2876, 2018. 2
- [39] Yunpeng Shi and Gilad Lerman. Message passing least squares framework and its application to rotation synchronization. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020. 1, 2, 7
- [40] Yunpeng Shi, Shaohan Li, and Gilad Lerman. Robust multi-object matching via iterative reweighting of the graph connection laplacian. *Advances in Neural Information Processing Systems*, 2020-December, 2020. 2
- [41] Chitturi Sidhartha and Venu Madhav Govindu. It is all in the weights: Robust rotation averaging revisited. In *International Conference on 3D Vision, 3DV 2021, London, United Kingdom, December 1-3, 2021*, pages 1134–1143. IEEE, 2021. 2
- [42] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and computational harmonic analysis*, 30(1):20–36, 2011. 2, 5
- [43] A. Singer and Y. Shkolnisky. Angular synchronization by eigenvectors and semidefinite programming: Analysis and application to class averaging in cryo-electron microscopy, 2009. 2
- [44] Johan Thunberg, Florian Bernard, and Jorge Goncalves. Distributed methods for synchronization of orthogonal matrices over graphs. *Automatica*, 80:243–252, jun 2017. 2
- [45] Roberto Tron and René Vidal. Distributed image-based 3-d localization of camera sensor networks. In *Proceedings of the 48th IEEE Conference on Decision and Control, CDC 2009, December 16-18, 2009, Shanghai, China*, pages 901–908, 2009. 2
- [46] Anton Voropaev and Sergey Perepechko. The number of fixed length cycles in undirected graph. explicit formulae in case of small lengths. *Bulletin of PFUR. Series Mathematics, Information Sciences, Physics*, pages 5–11, 01 2012. 4
- [47] Anton Voropaev and Sergey Perepechko. The number of fixed length cycles in undirected graph. explicit formulae in case of small lengths. *Bulletin of PFUR. Series Mathematics, Information Sciences, Physics*, pages 5–11, 01 2012. 13
- [48] Van H Vu. A large deviation result on the number of small subgraphs of a random graph. *Combinatorics, Probability and Computing*, 10(1):79–94, 2001. 5, 14, 17
- [49] Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference*, 2013. 2
- [50] Kyle Wilson and Noah Snavely. Robust global translations with 1dsfm. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III*, pages 61–75, 2014. 7
- [51] Christopher Zach, Manfred Klopschitz, and Marc Pollefeys. Disambiguating visual relations using loop constraints. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, pages 1426–1433, 2010. 2