# Stochastic Bundle Adjustment for Efficient and Scalable 3D Reconstruction

Lei Zhou<sup>1</sup>[0000-0003-4988-5084]</sup>, Zixin Luo<sup>1</sup>[0000-0001-6946-2826]</sup>, Mingmin Zhen<sup>1</sup>[0000-0002-8180-1023]</sup>, Tianwei Shen<sup>1</sup>[0000-0002-3290-2258]</sup>, Shiwei Li<sup>2</sup>[0000-0003-0712-0059]</sup>, Zhuofei Huang<sup>1</sup>[0000-0002-8479-7568]</sup>, Tian Fang<sup>2</sup>[0000-0002-5871-3455]</sup>, and Long Quan<sup>1</sup>[0000-0001-8148-1771]</sup>

<sup>1</sup> Hong Kong University of Science and Technology {lzhouai,zluoag,mzhen,tshenaa,zhuangbr,quan}@cse.ust.hk <sup>2</sup> Everest Innovation Technology {sli,fangtian}@altizure.com

Abstract. Current bundle adjustment solvers such as the Levenberg-Marquardt (LM) algorithm are limited by the bottleneck in solving the Reduced Camera System (RCS) whose dimension is proportional to the camera number. When the problem is scaled up, this step is neither efficient in computation nor manageable for a single compute node. In this work, we propose a stochastic bundle adjustment algorithm which seeks to decompose the RCS approximately inside the LM iterations to improve the efficiency and scalability. It first reformulates the quadratic programming problem of an LM iteration based on the clustering of the visibility graph by introducing the equality constraints across clusters. Then, we propose to relax it into a chance constrained problem and solve it through sampled convex program. The relaxation is intended to eliminate the interdependence between clusters embodied by the constraints, so that a large RCS can be decomposed into independent linear sub-problems. Numerical experiments on unordered Internet image sets and sequential SLAM image sets, as well as distributed experiments on large-scale datasets, have demonstrated the high efficiency and scalability of the proposed approach. Codes are released at https://github.com/zlthinker/STBA.

Keywords: Stochastic bundle adjustment, Clustering, 3D reconstruction

# 1 Introduction

Bundle Adjustment (BA) is typically formulated as a nonlinear least square problem to refine the parameters of cameras and 3D points. It is usually addressed by the Levenberg-Marquardt (LM) algorithm, where a linear equation system called *Reduced Camera System* (RCS) [25, 15] must be solved in each iteration. However, when the problem is scaled up, solving the RCS has been a bottleneck which takes a major portion of computation time (see the first bar of Fig. 1). The dimension of the RCS is proportional to the camera number, and thus the  $\mathbf{2}$ 



Fig. 1. Per-iteration time of bundle adjustment w.r.t. the compute node number. The Levenberg-Marquardt (LM) algorithm is limited by the bottleneck when solving the reduced camera system (RCS). Our STBA splits the RCS into independent sub-problems, which achieves a speedup on a single-threaded compute node. Besides, STBA allows parallel and distributed computing with multiple compute nodes which further improves the efficiency and scalability.

increase of cameras would ramp up the computation and memory consumption, although methods have been proposed to use efficient linear solvers [2, 23, 13, 21, 34] and economize on matrix manipulations [2, 22, 34]. Furthermore, different to other operations such as Jacobian or gradient evaluations, this step is indivisible, making it hard to fit BA for parallel and distributed computing.

In order to accomplish efficient and scalable reconstructions, clustering has been adopted as a useful practice to decompose a large problem into smaller, more manageable ones. For example, a number of SfM approaches have been developed in a divide and conquer fashion, which first reconstruct the partitioned sub-maps independently and then merge the partial reconstructions together [29, 37, 38, 17]. Although these methods are able to produce the initial sparse reconstructions in an efficient and scalable way, a full bundle adjustment is still indispensable to optimize the camera and point parameters globally. Therefore, in the context of BA, the methods [16, 36] proposed to distribute the objectives of BA to the split sub-models and optimize the sum of the objectives under the distributed optimization frameworks [10, 5], which, however, involves extra costly inner iterations and thus makes the optimization over-complicated.

In this work, we follow the direction of exploiting the clustering methods and push forward the investigation on how to integrate a clustering scheme into the BA problem systematically. Instead of applying a fixed, and one-time partition at the pre-processing step, we derive a stochastic clustering-based strategy within each LM iteration so as to decompose the RCS for efficiency and scalability.

- First, we reformulate the quadratic programming problem of an LM iteration based on the clustering of the visibility graph. Such a formulation splits the problem into the most elementary structures, but meanwhile introduces additional equality constraints and raises the computational cost.
- Second, in order to make the above problem efficiently solvable, we propose to relax the constraints into chance constraints [24] and then solve it with sampled convex program [7]. The approach helps to eliminate the interdependence between different clusters by randomized constraint reduction, which

hence decomposes the RCS into independent linear sub-problems related to the clusters. In this way, an approximate step can be achieved efficiently.

 Third, we present an add-on technique which helps to correct the approximate steps towards the steepest descent direction within a small trust region to improve the convergence.

Due to the stochastic process induced by the sampled convex program, we term our algorithm STochastic Bundle Adjustment (STBA), which brings the following tangible advantages. First, solving the split RCS in place of the original one has achieved a great speedup thanks to the reduced complexity. Second, the solving process can be made parallel and scalable to accommodate the growth of camera numbers, since all the sub-steps of a BA iteration can be decomposed. In Fig. 1, we visualize how the running time is reduced by distributing STBA over multiple compute nodes.

# 2 Related Works

Bundle adjustment (BA) is typically solved by the Levenberg-Marquardt (LM) algorithm [26], which approximately linearizes the error functions inside a local trust region and then solves a linear normal equation for an update step. SBA [25] first simplified the norm equation into a reduced camera system (RCS) through Schur complement by taking advantage of the special problem structure. After this, efforts were dedicated to solving the RCS faster in either exact or inexact ways. The exact solvers apply Cholesky factorization to the reduced camera matrix  $\mathbf{S}$ , while exploiting variable ordering [3, 12] and supernodal methods [30, 12] for acceleration. The inexact solvers are based on the Conjugate Gradient (CG) method [19] coupled with various preconditioners [2, 20, 23], which attains inexact solutions with better efficiency. Apart from the algorithmic improvements, [22, 34] presented well-optimized implementations of the LM solver to save the memory usage and exploit the CPU and GPU parallelism. However, despite the efforts above, solving a large and indivisible RCS will increasingly become the bottleneck of a BA solver when the problem is scaled up.

In order to make large-scale reconstruction tractable, the clustering methods are initially introduced into the structure from motion (SfM) domain. Basically, a divide-and-conquer strategy is applied, which first partitions a large scene into multiple sub-maps and then merges the partial reconstructions globally [29, 37, 38, 17]. In the formulation of these approaches, a reduced optimization problem other than the original BA problem is addressed, thus leading to a sub-optimal result. For example, [29, 38] factored out the internal variables inside the submaps and [37] registered all the cameras with motion averaging [9] without the involvement of points. In the realm of BA, [23] derived a block-diagonal preconditioner for the RCS from the clustering of cameras, but the clustering did not help to decompose the problem as it is done in the SfM algorithms [29, 37, 38, 17]. Instead, [16, 36] proposed to apply the distributed optimization frameworks like the Douglas-Rachford method [10] and ADMM [5] onto the empirically clusterized BA problems towards large scales. Although built upon a theoretical foundation, the methods required costly inner iterations and introduced a plethora of latent parameters during optimization.

## 3 Bundle Adjustment Revisited

In this section, we first revisit the bundle adjustment problem and its LM solution to give the necessary preliminaries and terminologies. Henceforth, vectors and matrices appear in boldface and  $\|.\|$  denotes the L2 norm.

A bundle adjustment problem is built upon a bipartite visibility graph  $\mathcal{G} = (\mathcal{C} \cup \mathcal{P}, \mathcal{E})$ . Here,  $\mathcal{C} = \{\mathbf{c}_i \in \mathbb{R}^d\}_{i=1}^m$  denotes the set of m cameras parameterized by d-dimensional vectors,  $\mathcal{P} = \{\mathbf{p}_i \in \mathbb{R}^3\}_{i=1}^n$  denotes the set of n 3D points, and  $\mathcal{E} = \{\mathbf{q}_i \in \mathbb{R}^2\}_{i=1}^q$  denotes the set of q projections. The objective is to minimize  $F(\mathbf{x}) = \|\mathbf{f}(\mathbf{x})\|^2$ , where  $\mathbf{f}$  denotes a 2q-dimensional vector of reprojection errors and  $\mathbf{x}$  concatenates camera parameters  $\mathbf{c} \in \mathbb{R}^{dm}$  and point parameters  $\mathbf{p} \in \mathbb{R}^{3n}$ , *i.e.*,  $\mathbf{x} = [\mathbf{c}^T, \mathbf{p}^T]^T$ .

The LM algorithm achieves an update step  $\Delta \mathbf{x}$  at each iteration by linearizing  $\mathbf{f}(\mathbf{x})$  as  $\mathbf{J}(\mathbf{x})\Delta\mathbf{x} + \mathbf{f}(\mathbf{x})$  in a trust region around  $\mathbf{x}$ , where  $\mathbf{J}(\mathbf{x}) = \nabla_{\mathbf{x}}\mathbf{f}(\mathbf{x}) = [\mathbf{J}(\mathbf{c}), \mathbf{J}(\mathbf{p})]$  is the Jacobian matrix. Then the minimization of  $F(\mathbf{x})$  is turned into

$$\min_{\Delta \mathbf{x}} \|\mathbf{J}(\mathbf{x})\Delta \mathbf{x} + \mathbf{f}(\mathbf{x})\|^2 + \lambda \|\mathbf{D}\Delta \mathbf{x}\|^2,$$
(1)

whose solution comes from the normal equation below

$$\begin{bmatrix} \mathbf{J}(\mathbf{x}) \\ \sqrt{\lambda} \mathbf{D} \end{bmatrix} \Delta \mathbf{x} = \begin{bmatrix} -\mathbf{f}(\mathbf{x}) \\ \mathbf{0} \end{bmatrix}, \qquad (2)$$

where  $\lambda > 0$  is the damping parameter and typically  $\mathbf{D} = \text{diag}(\mathbf{J}^T(\mathbf{x})\mathbf{J}(\mathbf{x}))^{\frac{1}{2}}$ . For notational simplicity, we write  $\mathbf{J} \triangleq \mathbf{J}(\mathbf{x}), \mathbf{J}_{\mathbf{c}} \triangleq \mathbf{J}(\mathbf{c}), \mathbf{J}_{\mathbf{p}} \triangleq \mathbf{J}(\mathbf{p})$  and  $\mathbf{f} \triangleq \mathbf{f}(\mathbf{x})$ . After multiplying  $[\mathbf{J}^T, \sqrt{\lambda}\mathbf{D}^T]$  at both sides of Eq. 2, we have

$$(\mathbf{J}^T \mathbf{J} + \lambda \mathbf{D}^T \mathbf{D}) \Delta \mathbf{x} = -\mathbf{J}^T \mathbf{f}, \qquad (3)$$

which can be re-written in the form

$$\begin{bmatrix} \mathbf{B} & \mathbf{E} \\ \mathbf{E}^T & \mathbf{C} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{c} \\ \Delta \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix}, \tag{4}$$

where  $\mathbf{B} = \mathbf{J}_{\mathbf{c}}^T \mathbf{J}_{\mathbf{c}} + \lambda \operatorname{diag}(\mathbf{J}_{\mathbf{c}}^T \mathbf{J}_{\mathbf{c}}), \ \mathbf{C} = \mathbf{J}_{\mathbf{p}}^T \mathbf{J}_{\mathbf{p}} + \lambda \operatorname{diag}(\mathbf{J}_{\mathbf{p}}^T \mathbf{J}_{\mathbf{p}}), \ \mathbf{E} = \mathbf{J}_{\mathbf{c}}^T \mathbf{J}_{\mathbf{p}}, \ \mathbf{v} = -\mathbf{J}_{\mathbf{c}}^T \mathbf{f}, \ \text{and} \ \mathbf{w} = -\mathbf{J}_{\mathbf{p}}^T \mathbf{f}. \ \text{Eq. 4 can be simplified by the Schur complement [25],}$ which leads to

$$\mathbf{S}\Delta\mathbf{c} = \mathbf{v} - \mathbf{E}\mathbf{C}^{-1}\mathbf{w},\tag{5}$$

$$\Delta \mathbf{p} = \mathbf{C}^{-1} (\mathbf{w} - \mathbf{E}^T \Delta \mathbf{c}), \tag{6}$$

where  $\mathbf{S} = \mathbf{B} - \mathbf{E}\mathbf{C}^{-1}\mathbf{E}^{T}$  is the Schur complement of **C**. Here, **S**, known as the *reduced camera matrix*, is a block structured symmetric positive definite matrix.



Fig. 2. Illustration of point splitting/binding over the visibility graph (top) and the corresponding structure of the reduced camera matrix (bottom). Point splitting helps to reshape the reduced camera matrix into the block-diagonal structure, while point binding does the inverse.

The block  $\mathbf{S}_{ij} \in \mathbb{R}^{d \times d}$  is nonzero iff cameras  $\mathbf{c}_i$  and  $\mathbf{c}_j$  observe at least one common point. Although a variety of sparse Cholesky factorization techniques [3, 12, 30] and preconditioned conjugate gradient methods [2, 20, 23] have been developed to solve the *reduced camera system (RCS)* of Eq. 5, it still can be prohibitive when the camera number m grows large.

# 4 Stochastic Bundle Adjustment

In this section, we present our stochastic bundle adjustment (STBA) method that decomposes the RCS into clusters inside the LM iterations. In Sec. 4.1, we first reformulate problem (1) based on the clustering of the visibility graph  $\mathcal{G}$ , yet subject to additional equality constraints. Next, in Sec. 4.2, we apply chance constrained relaxation to the reformulation and solve it by sampled convex program [7,8]. It manages to decompose the RCS into cluster-related linear sub-problems and yield an approximate STBA step efficiently. Third, a steepest correction step is proposed to remedy the approximation error of the STBA steps in Sec. 4.3. Finally, in Sec. 4.4, we present a practical implementation of the random constraint sampler required by the chance constrained relaxation.

### 4.1 Clustering Based Reformulation

In constrast to the previous methods [29, 37, 38, 17, 16, 36] that partition the problem in the pre-processing stage, we present a reformulation of problem (1) to decompose the RCS into clusters inside the LM iterations.

Particularly, we consider the most general case that every single camera forms a cluster. In order to preserve all the projections  $\mathcal{E}$ , we apply **point splitting** to the physical points, as shown in Fig. 2. For a physical point  $\mathbf{p}_i$  viewed by  $v_i$  cameras, we split it into  $v_i$  virtual points  $\{\mathbf{p}_i^j\}_{j=1}^{v_i}$ , each assigned to one cluster. Such a clustering will reformulate problem (1) equivalently as a new constrained quadratic programming (QP) problem as below

$$\min_{\Delta \mathbf{x}'} \|\mathbf{J}' \Delta \mathbf{x}' + \mathbf{f}\|^2 + \lambda \|\mathbf{D}' \Delta \mathbf{x}'\|^2,$$
(7)

s.t. 
$$\mathbf{A} \Delta \mathbf{x}' = \mathbf{0}$$
. (8)

6

Here,  $\Delta \mathbf{x}' = [\Delta \mathbf{c}^T, \Delta \mathbf{p}'^T]^T$  is an expansion of  $\Delta \mathbf{x}$  which considers the update steps for all the virtual points, and so is the Jacobian  $\mathbf{J}' = [\mathbf{J_c}, \mathbf{J'_p}]$ . Accordingly,  $\mathbf{D}' = \operatorname{diag}(\mathbf{J}'^T \mathbf{J}')^{\frac{1}{2}}$ . The noteworthy distinction between problems (7) and (1) is that the new equality constraints of Eq. 8 are imposed to enforce that the steps of the same points in different clusters are identical. For example,  $\Delta \mathbf{p}_i^s = \Delta \mathbf{p}_i^t$  $(\forall s, t \in \{1, ..., v_i\})$  for point  $\mathbf{p}_i$  and the corresponding j-th row of  $\mathbf{A}$  appears in a form as  $\mathbf{a}_j = [0..., 1, ..., -1, ...0]$ . Since a point  $\mathbf{p}_i$  introduces  $v_i - 1$  equations,  $\mathbf{A}$ has a row number of  $r = \sum_{i=1}^n (v_i - 1)$ . Besides,  $\mathbf{A}$  is full row rank, because the rows each of which defines a unique equality constraint are linearly independent.

The constrained QP problem above can be easily solved by Lagrangian duality, which turns the problem into

$$\mathbf{H}_{\lambda} \Delta \mathbf{x}' = -(\mathbf{J}'^T \mathbf{f} + \mathbf{A}^T \boldsymbol{\nu}), \qquad (9)$$

where  $\mathbf{H}_{\lambda} = \mathbf{J}'^T \mathbf{J}' + \lambda \mathbf{D}'^T \mathbf{D}'$  and  $\boldsymbol{\nu} = -(\mathbf{A}\mathbf{H}_{\lambda}^{-1}\mathbf{A}^T)^{-1}\mathbf{A}\mathbf{H}_{\lambda}^{-1}\mathbf{J}'^T\mathbf{f}$  are the Lagrangian multipliers. Eq. 9 is in the similar format to Eq. 3, but has an additional term  $\mathbf{A}^T \boldsymbol{\nu}$  on the right hand side compared with Eq. 3. While  $\mathbf{J}'^T\mathbf{f}$  includes the gradients w.r.t. the independent virtual points,  $\mathbf{A}^T\boldsymbol{\nu}$  acts as a **correction term** to ensure that the solution complies with the constraints of Eq. 8.

The ultimate benefit of the clustering-based reformulation is revealed below. By likewise applying Schur complement to Eq. 9, we have

$$\mathbf{S}' \Delta \mathbf{c} = \mathbf{v}' - \mathbf{E}' \mathbf{C}'^{-1} \mathbf{w}', \tag{10}$$

where  $\mathbf{E}' = \mathbf{J}_{\mathbf{c}}^T \mathbf{J}_{\mathbf{p}}', \mathbf{C}' = \mathbf{J}_{\mathbf{p}}'^T \mathbf{J}_{\mathbf{p}}' + \lambda \operatorname{diag}(\mathbf{J}_{\mathbf{p}}'^T \mathbf{J}_{\mathbf{p}}'), \mathbf{S}' = \mathbf{B} - \mathbf{E}' \mathbf{C}'^{-1} \mathbf{E}'^T$ , and  $[\mathbf{v}'^T, \mathbf{w}'^T]^T = -(\mathbf{J}'^T \mathbf{f} + \mathbf{A}^T \boldsymbol{\nu})$ . Due to the fact that any two cameras do not share any common virtual points,  $\mathbf{S}'$  now becomes a **block-diagonal** matrix, *i.e.*,  $\mathbf{S}'_{ij} = \mathbf{0}, \forall i \neq j$ . Then Eq. 10 can be equivalently decomposed into *m* most elementary linear systems each corresponding to one camera.

#### 4.2 Chance Constrained Relaxation

The major problem with the clustering based reformulation above is the excessive cost of evaluating the Lagrangian multipliers  $\boldsymbol{\nu}$ , because it requires the evaluation of  $\mathbf{H}_{\lambda}^{-1}$ . In order to make the problem practically solvable, we would like to eliminate the need to evaluate the correction term  $\mathbf{A}^T \boldsymbol{\nu}$  of Eq. 9 by means of relaxation for problem (7).

We multiply a random binary variable  $\theta_i$  with the i-th equality constraint of Eq. 8, which results in  $\theta_i \mathbf{a}_i \Delta \mathbf{x}' = 0$ . It could be interpreted that, if  $\theta_i = 1$ , the constraint  $\mathbf{a}_i \Delta \mathbf{x}' = 0$  must be satisfied; otherwise, the constraint is allowed to be violated. In this way, Eq. 8 will be relaxed into chance constraints [24], which leads to

$$\min_{\Delta \mathbf{x}'} \|\mathbf{J}' \Delta \mathbf{x}' + \mathbf{f}\|^2 + \lambda \|\mathbf{D}' \Delta \mathbf{x}'\|^2,$$
(11)

s.t. 
$$\operatorname{Prob}(\mathbf{a}_i \Delta \mathbf{x}' = 0) = \operatorname{Prob}(\theta_i = 1) \ge \alpha, \ i = 1, ..., r,$$
 (12)

where  $\alpha \in (0, 1]$  is a predefined confidence level. It means that, instead of enforcing the hard constraints, we allow them to be satisfied with a probability above  $\alpha$ . The advantage of the chance constrained relaxation is that we can determine the reliability level of approximation by controlling  $\alpha$ . The larger  $\alpha$  is, the closer the chance constrained problem (11) will be to the original deterministic problem (7). One approach to problem (11) is called sampled convex program [7,8]. It extracts N independent samples  $\Theta(\alpha) = \{\theta_i^{(n)} | i = 1, ..., r, n = 1, ..., N\}$  with a minimum sampling probability of  $\alpha$  for each variable  $\theta_i^{(n)}$  and replaces the chance constraints (12) with the sampled ones. Below we will elaborate on how problem (11) can be solved given the samples  $\Theta(\alpha)$ .

For a sample  $\theta_i^{(n)} \in \Theta(\alpha)$ , if  $\theta_i^{(n)} = 0$ , the equality constraint  $\mathbf{a}_i \Delta \mathbf{x}' = 0$  is dropped; if  $\theta_i^{(n)} = 1$ , it enforces the equality of the steps of two virtual points, e.g.,  $\Delta \mathbf{p}_j^s = \Delta \mathbf{p}_j^t$ . Here the virtual point  $\mathbf{p}_j^s$  belongs to the single-camera cluster of camera  $\mathbf{c}_s$  and similarly  $\mathbf{p}_j^t$  to  $\mathbf{c}_t$ . Then we merge  $\mathbf{p}_j^s$  and  $\mathbf{p}_j^t$  into one point as shown in Fig. 2, and we call the operation **point binding** as opposed to **point splitting** introduced in Sec. 4.1. On the one hand, the point binding leads to the consequence that the Lagrangian multiplier  $\nu_i = 0$ , because the equality  $\Delta \mathbf{p}_j^s = \Delta \mathbf{p}_j^t$  always holds. On the other hand, the block  $\mathbf{S}'_{st}$  of  $\mathbf{S}'$  (*c.f.* Eq. 10) becomes nonzero, since the merged point is now shared by cameras  $\mathbf{c}_s$  and  $\mathbf{c}_t$ . After applying point binding to all the virtual points involved in the sampled constraints, *i.e.*,  $\{\theta_i^{(n)} \in \Theta(\alpha) | \theta_i^{(n)} = 1\}$ , all the constraints will be eliminated and there is no need to evaluate  $\boldsymbol{\nu}$  any more. Meanwhile, it will bring the cameras sharing common points into the same clusters.

Since the cameras in different clusters have no points in common after the point binding, the matrix  $\mathbf{S}'$  will appear in a block-diagonal structure which we call **cluster-diagonal**, as illustrated in Fig. 2. It means that each diagonal block of  $\mathbf{S}'$  corresponds to a camera cluster. In particular, we can intentionally design the sampler  $\Theta(\alpha)$  in order to shape  $\mathbf{S}'$  into the desired cluster-diagonal structures, as we will present in Sec. 4.4. As a result, this structure of  $\mathbf{S}'$  still enables the decomposition of Eq. 10 into smaller independent linear systems each relating to one camera cluster. Provided that there are l clusters, Eq. 10 can be equivalently re-written as

$$\begin{cases} \mathbf{S}_{1}^{\prime} \Delta \mathbf{c}_{1} = \mathbf{b}_{1}, \\ \dots \\ \mathbf{S}_{l}^{\prime} \Delta \mathbf{c}_{l} = \mathbf{b}_{l}, \end{cases}$$
(13)

where  $[\Delta \mathbf{c}_1^T, ..., \Delta \mathbf{c}_l^T]^T = \Delta \mathbf{c}$  and  $[\mathbf{b}_1^T, ..., \mathbf{b}_l^T]^T = \mathbf{v}' - \mathbf{E}' \mathbf{C}'^{-1} \mathbf{w}'$ . After Eqs. 13 are evaluated, we substitute  $\Delta \mathbf{c}$  into Eq. 9 to give

$$\mathbf{J}_{\mathbf{p}}^{\prime} \Delta \mathbf{p}^{\prime} = \sum_{i=1}^{i} \mathbf{J}_{\mathbf{p}^{i}} \Delta \mathbf{p}^{i} = -\mathbf{f} - \mathbf{J}_{\mathbf{c}} \Delta \mathbf{c}, \qquad (14)$$

where  $\mathbf{J}'_{\mathbf{p}} = [\mathbf{J}_{\mathbf{p}^1}, ..., \mathbf{J}_{\mathbf{p}^l}]$  and  $\Delta \mathbf{p}' = [\Delta \mathbf{p}^1, ..., \Delta \mathbf{p}^l]$  include the Jacobians and virtual point steps w.r.t. the *l* clusters respectively and we omit  $\mathbf{D}'$  for ease of notation. To give a uniform step for a physical point, we equalize the steps of its

8

virtual points in different clusters by solving the linear system below in place of Eq. 14:

$$\sum_{i=1}^{l} \mathbf{J}_{\mathbf{p}^{i}} \Delta \mathbf{p} = -\mathbf{f} - \mathbf{J}_{\mathbf{c}} \Delta \mathbf{c}.$$
 (15)

Since  $\sum_{i=1}^{l} \mathbf{J}_{\mathbf{p}^{i}} = \mathbf{J}_{\mathbf{p}}$ , Eq. 15 gives the same solution as the point steps in Eq. 6:  $\Delta \mathbf{p} = \mathbf{C}^{-1}(\mathbf{w} - \mathbf{E}^{T}\Delta \mathbf{c}).$ 

So far we have presented how an update step of the camera and point parameters is determined approximately by STBA. Besides this, we keep the other components of the LM algorithm unchanged [27]. For reference, we detail the full algorithm in the supplementary material.

#### 4.3 Steepest Descent Correction

The chance constrained relaxation in the last section effectively decomposes the RCS, but leads to approximate solutions with decreased feasibility due to the random constraint sampling. Below, we provide an empirical analysis on the effect of the approximation and present a conditional correction step to remedy the approximation error.

The LM algorithm is known to be the interpolation of the Gauss-Newton and gradient descent methods, depending on the trust region radius controlled by the damping parameter  $\lambda$ . When  $\lambda$  is small and the LM algorithm behaves more like the Gauss-Newton method, the approximation induced by STBA in Sec. 4.2 is admissible, in that problem (1) itself is derived from the first order approximation of the error function  $\mathbf{f}(\mathbf{x})$ . And the LM algorithm can automatically contract the trust region when the approximation leads to the increase of the objective.

When  $\lambda$  is large, *i.e.*, the trust region is small, the LM algorithm is closer to the gradient descent method, which gives a step towards the steepest descent direction defined by the right hand side of Eq. 9, *i.e.*,  $-(\mathbf{J}'^T \mathbf{f} + \mathbf{A}^T \boldsymbol{\nu})$ . However, a problem with STBA is that the correction term  $\mathbf{A}^T \boldsymbol{\nu}$  is eliminated approximately by the chance constrained relaxation. As a consequence, the derived step would deviate from the steepest descent direction and thus hamper the convergence. Therefore, we propose to recover  $\mathbf{A}^T \boldsymbol{\nu}$  to remedy the deviation in such a case. Especially, when  $\lambda$  is large enough, the matrix  $\mathbf{H}_{\lambda}$  in Eq. 9 will be dominated by the diagonal terms, so that we can approximate  $\mathbf{H}_{\lambda}$  by diag $(\mathbf{H}_{\lambda})$ . After that,  $\mathbf{A}^T \boldsymbol{\nu} = -\mathbf{A}^T (\mathbf{A} \mathbf{H}_{\lambda}^{-1} \mathbf{A}^T)^{-1} \mathbf{A} \mathbf{H}_{\lambda}^{-1} \mathbf{J}'^T \mathbf{f}$  can be evaluated efficiently because of the sparsity of  $\mathbf{A}$ . Since the approximation  $\mathbf{H}_{\lambda} \approx \text{diag}(\mathbf{H}_{\lambda})$  is not accurate unless  $\lambda$ is large, in practice, we enable the steepest descent correction particularly when  $\lambda \geq 0.1$ . Fig. 3 visualizes the effect of the correction.

### 4.4 Stochastic Graph Clustering

The chance constrained relaxation in Sec. 4.2 necessitates an effective random constraint sampler  $\Theta(\alpha)$ . Among many of the possible designs, we propose a viable implementation named stochastic graph clustering in this section.



Fig. 3. Visualization of the steepest descent correction steps of sample camera parameters when  $\lambda = 0.1$ . It shows that the deviations of approximate STBA steps from the LM steps are effectively corrected.

The design of the clustering method considers the following requirements. First, the sampler should be randomized with respect to the chance constraints (12). Since (12) indicates that the expectation  $E(\theta_i)$  should have  $E(\theta_i) \ge \alpha$ (i = 1, ..., r), the upper bound of the confidence level  $\alpha$  is defined as  $\min_{i=1}^{r} E(\theta_i)$ . Therefore, the sampler should sample as many constraints as possible on average to increase the upper bound of  $\alpha$ . Second, the random sampler is intended to partition the cameras into small independent clusters so that Eqs. 13 can be solved efficiently.

Concretely, the stochastic graph clustering operates over a camera graph  $\mathcal{G}_c = (\mathcal{C}, \mathcal{E}_c)$ , where the weight  $w_{ij}$  of an edge  $e_{ij} \in \mathcal{E}_c$  between cameras  $\mathbf{c}_i$  and  $\mathbf{c}_j$ is equal to the number of points covisible by the two cameras. At the beginning, each camera forms an individual cluster as formulated in Sec. 4.1. Next, if  $\mathbf{c}_i$  and  $\mathbf{c}_i$  are joined, a number of  $w_{ij}$  pairs of virtual points viewed by  $\mathbf{c}_i$  and  $\mathbf{c}_j$  will be merged. Therefore,  $w_{ij}$  equality constraints will be satisfied. In order to join as many virtual points as possible while yielding a cluster structure of  $\mathcal{G}_c$ , we aim at finding a clustering that maximizes the modularity below inspired by [6]: Q = $\frac{1}{2s} \sum_{e_{ij} \in \mathcal{E}_c} \delta(\nu_i, \nu_j) \left( w_{ij} - \frac{k_i k_j}{2s} \right), \text{ where } s = \sum_{e_{ij} \in \mathcal{E}_c} w_{ij} \text{ is the total sum of edge}$ weights,  $k_i = \sum_j w_{ij}$  is the sum of weights of edges incident to camera  $\mathbf{c}_i$ , and  $\nu_i$ denotes the cluster of  $\mathbf{c}_i$ .  $\delta(\nu_i, \nu_j) = 1$  if  $\nu_i = \nu_j$  and 0 otherwise. The modularity  $Q \in [-1,1]$  measures the density of connections inside clusters as opposed to those across clusters [6]. Therefore, a larger modularity generally indicates that more virtual points are merged inside clusters. Maximizing the modularity is NPhard [31], but Louvain's algorithm [6] provides a greedy strategy which greedily joins the two clusters giving the maximum increase in modularity in a bottomup manner. It can be efficiently applied to large graphs since its complexity is shown to be linear in the node number on sparse data [6].

However, Louvain's algorithm [6] is deterministic due to its greedy nature. To ensure that every pair of virtual points is likely to be merged, we instead join clusters randomly according to a probability distribution defined based on the modularity increments [11], which is

$$\operatorname{Prob}(N_x \cup N_y) = \frac{\exp(\beta \Delta Q(N_x, N_y))}{\sum_i \sum_j \exp(\beta \Delta Q(N_i, N_j))},$$
(16)



Fig. 4. Visualization of the random clustering results produced by stochastic graph clustering. Cameras of different clusters are in different colors.

where  $\beta > 0$  is a scaling parameter. Two neighboring clusters  $N_x$  and  $N_y$  are more likely to join together if it leads to a larger modularity increment  $\Delta Q(N_x, N_y)$ . In order to limit the sizes of the sub-problems of STBA, we stop joining clusters if their sizes exceed  $\Gamma$ . In Fig. 4, we visualize the stochastic clustering results.

# 5 Experiments

#### 5.1 Experiment Settings

**Datasets.** We run experiments on three different types of datasets: 1) 1DSfM dataset [33] which is composed of 14 sets of unordered Internet images; 2) KITTI dataset [18] containing 11 street-view image sequences; and 3) Large-Scale dataset which is collected by ourselves due to the absence of publicly available large-scale 3D datasets. It includes 4 image sets each comprising more than 30,000 images. The problem sizes all exceed the memory of a single compute node and thus we use them particularly for distributed experiments.

**Comparisons.** On 1DSfM and KITTI datasets, we compare our method with two standard trust region algorithms, Levenberg-Marquardt (LM) [25] and Dogleg (DL) [26]. For the LM algorithm, we use two variants: LM-sparse and LMiterative, which exploit the exact sparse method and inexact iterative method [23] to solve the RCS (Eq. 5), respectively. For the distributed experiments on the Large-Scale dataset, we compare our distributed implementation of STBA against the state-of-the-art distributed solver DBACC [36]. The ablation studies on steepest descent correction and stochastic graph clustering are presented in Sec. 5.4 and the supplementary material, respectively.

**Implementations.** We implement LM, DL and STBA in C++, using Eigen for linear algebra computations. All the algorithms are implemented from the same code base, which means that they share the same elementary operations so that they can be compared equitably. For robustness, we use the Huber loss with a scale factor of 0.5 for the errors [35]. LM-sparse exploits the supernodal  $LL^T$  Cholesky factorization with COLAMD ordering [12] based on CHOLMOD, which is well suited for handling sparse data like KITTI [18]. LM-iterative uses the conjugate gradient method with the advanced cluster-jacobi preconditioner [23]. DL uses the same exact sparse solver as LM-sparse since it requires a reasonably good estimation of the Gauss-Newton step [26]. Dense  $LL^T$  factorization is used to solve the decomposed RCS (Eqs. 13) for STBA due to the dense connectivity inside camera clusters. Multi-threading is applied to the operations including the reprojection error and Jacobian computation, the preconditioner construction and the matrix-vector multiplications for all the methods as in [34].

**Parameters.** In the experiments, we assume that the camera intrinsics have been calibrated as in [26, 25]. Camera extrinsics are parameterized with 6-d vectors, using axis-angle representations for rotations. We set the initial damping parameter  $\lambda$  to 1e-4 and the max iteration number to 100 for all the methods. The iterations could terminate early if the cost, gradient or parameter tolerance [1] drops below 1e-6. For STBA, we empirically set the scaling parameter  $\beta$  to 10 (Eq. 16) and the max cluster size  $\Gamma$  to 100.

**Hardware.** We use a compute node with an 8-core Intel i7-4790K CPU and a 32G RAM. The distributed experiments are deployed on a cluster with 6 compute nodes.

#### 5.2 Performance Profiles

Following previous works [14, 23], we evaluate the solvers with *Performance Profiles* over the total of 25 problems of 1DSfM [33] and KITTI [18]. We obtain the SfM results for 1DSfM by COLMAP [32]<sup>3</sup> and the SLAM results for KITTI by stereo ORB-SLAM2 [28], while disabling the final bundle adjustment (BA). Since the SfM/SLAM results are generally accurate because the pipeline uses repeated BA for robust reconstruction, we make the problems more challenging by adding Gaussian noise to the points and camera centers following [16, 20, 29].

First of all, we give a brief introduction of performance profiles [14]. Given a problem  $p \in \mathcal{P}$  and a solver  $s \in \mathcal{S}$ , let F(p, s) denote the final objective the solver s attained when solving problem p. Then, for a number of solvers in  $\mathcal{S}$ , let  $F^*(p) = \min_{s \in \mathcal{S}} F(p, s)$  denote the minimum objective the solvers  $\mathcal{S}$  attained when solving problem p. Next, we define an objective threshold for problem p which is  $F_{\tau}(p) = F^*(p) + \tau(F_0(p) - F^*(p))$ , where  $F_0(p)$  is the initial objective and  $\tau \in (0, 1)$  is the pre-defined tolerance determining how close the threshold is to the minimum objective. After this, we measure the efficiency of a solver s by computing the time it takes to reduce the objective to  $F_{\tau}(p)$ , which is denoted by  $T_{\tau}(p, s)$ . And the most efficient solver is the one who takes the minimum time, *i.e.*,  $\min_{s \in \mathcal{S}} T_{\tau}(p, s)$ .

The method Performance Profiles regards that the solver s solves the problem p if  $T_{\tau}(p,s) \leq \alpha \min_{s \in S} T_{\tau}(p,s)$ , where  $\alpha \in [1,\infty)$ . Therefore, if  $\alpha = 1$ , only the most efficient solver is thought to solve the problem, while if  $\alpha \to \infty$ , all the solvers can be seen to solve the problem. Finally, the performance profile of the solver s is defined w.r.t.  $\alpha$  over the whole problem set  $\mathcal{P}$  as  $\rho(s,\alpha) =$ 

 $<sup>^{3}</sup>$  Since one of the image sets *Union Square* has only 10 reconstructed images, we replace it with another public image set *ArtsQuad*.

12 L. Zhou, Z. Luo, M. Zhen, T. Shen, et al.



Fig. 5. Performance profiles [14] of different solvers when solving the total of 25 problems of 1DSfM [33] and KITTI [18].



Fig. 6. The convergence curves of 4 scenes from 1DSfM and KITTI.

100 \*  $\frac{|\{p \in \mathcal{P} | T_{\tau}(p,s) \leq \alpha \min_{s \in S} T_{\tau}(p,s)\}|}{|\mathcal{P}|}$ . It is basically the percentage of problems solved by s and is non-decreasing w.r.t.  $\alpha$ .

We plot the performance profiles of the solvers in Fig. 5. To verify the benefits of using stochastic clustering, herein we also compare STBA with its variant STBA-fixed which uses a fixed clustering as previous methods [23, 16, 36]. When  $\tau$  is equal to 0.1 and 0.01, our STBA is able to solve nearly 100% of the problems for any  $\alpha$ , because it always reaches the objective threshold  $F_{\tau}(p)$  with less time than LM and DL methods by a factor of more than 5. This is mainly attributed to the reduced per-iteration cost as we can see from the convergence curves in Fig. 6. When  $\tau$  becomes 0.001 and the threshold  $F_{\tau}(p)$  is harder to achieve. STBA is less efficient but still performs on par with LM-iterative and LM-sparse when  $\alpha < 3$  and better than DL for any  $\alpha$ . On the contrary, the performance of STBA-fixed drops drastically when  $\tau$  decreases to 0.001. The performance change for STBA and STBA-fixed when  $\tau$  decreases is mainly caused by the fact that the clustering methods come with the price of slower convergence near the stationary points [16, 36]. Compared with the full second-order solvers such as LM and DL, clustering methods only utilize the second order information within clusters. However, as opposed to STBA-fixed which uses fixed clustering, STBA has mitigated the negative effect of clustering by introducing stochasticity so that different second-order information can be utilized to boost convergence as the clustering changes. Despite the slower convergence rate, the benefit of STBA that it can reduce most of the loss with the lowest time cost (e.q., 99%loss reduction with only 1/5 time of the counterparts when  $\tau = 0.01$  in Fig. 5(b)) is still supposed to be highlighted, especially for the real-time SLAM applications where bundle adjustment is called repeatedly to correct the drift.

Stochastic Bundle Adjustment for Efficient and Scalable 3D Reconstruction 13



Fig. 7. Visualizations of SfM results (top) and convergence curves (bottom) of the Larse-Scale dataset. Cameras are drawn as blue pyramids.

Table 1. Statistics of the distributed bundle adjustment solvers on the Large-Scale dataset. DBACC [36] consume many more Jacobian and RCS evaluations than STBA.

Data	#images	#clusters		RPE (pixel)		#Jacobian	/RCS evaluations	Mean iteration time (s)	
		DBACC	STBA	DBACC	STBA	DBACC	STBA	DBACC	STBA
LS-1	29975	300	340	0.823	0.818	1011/1080	49/100	912.5	71.0
LS-2	33634	336	386	0.766	0.783	854/860	48/100	934.8	79.3
LS-3	33809	339	391	1.083	1.056	1025/1100	49/100	1107.0	89.9
LS-4	44276	444	505	0.909	0.882	877/900	49/100	988.1	71.2

#### 5.3 Results on Large-Scale Dataset

To evaluate the scalability, we conduct distributed experiments on the Large-Scale (LS) dataset, which includes the urban scenes of four cities named LS-1, LS-2, LS-3 and LS-4. We run a distributed SfM program of our own to produce initial sparse reconstructions of the four scenes and add Gaussian noise with a standard deviation of 3 meters to the camera centers and points. Then we compare our distributed STBA against the state-of-the-art distributed bundle adjustment framework DBACC [36].

We visualize the sparse reconstructions and the convergence curves of the four scenes in Fig. 7 and report the statistics in Table 1. As we can see from Fig. 7, STBA achieves faster convergence rates than DBACC by an order of magnitude. The main cause of the gap is that DBACC, which is based on the ADMM formulation [4], has to take inner iterations to solve a new minimization problem in every ADMM iteration. Although we have set the maximum inner iteration number to merely 10, DBACC still takes many more Jacobian and RCS evaluations and thus has much longer iterations than STBA by an order of magnitude, as shown in Table 1. Besides, as opposed to DBACC, our STBA is free of too many hyper-parameters.

#### 5.4 Ablation Study on Steepest Descent Correction

Here we perform an ablation study on steepest descent correction proposed in Sec. 4.3 to validate its efficacy. We compare our STBA with its variant called



Fig. 8. Convergence curves of LM-sparse, STBA and STBA\* which does not use steepest descent correction (Sec. 4.3). Steepest descent correction helps to correct the deviations between the STBA and LM steps.

STBA\* which does not use the correction. We run STBA, STBA\* and LM-sparse on 1DSfM and KITTI. Since steepest descent correction is designed particularly for a small trust region, we set the lower bound of the damping parameter  $\lambda$  to 0.1 in the experiments. We observe that by using the correction, STBA consistently achieves a faster convergence than STBA\* and performs on par with LM-sparse on all the scenes. Visualizations of the sample convergence curves w.r.t. the iterations are shown in Fig. 8, where STBA and LM-sparse have very close convergence curves. It manifests that steep descent correction indeed facilitates the correction of the approximation errors of the STBA steps and hence boosts the convergence.

## 6 Conclusion

In this paper, we rethink the proper way of integrating the clustering scheme into solving bundle adjustment by proposing STBA. First, STBA reformulates an LM iteration based on the clustering of the visibility graph, but meanwhile introduces additional equality constraints across the clusters. Second, we approximately relax the constraints as chance constraints and solve the problem by sampled convex program which randomly samples the chance constraints with the intention of splitting the large reduced camera system into small clusters. Not only does it reduce the per-iteration cost, but also allows parallel and distributed computing to accommodate the increase of the problem size. Moreover, we present a steepest descent correction technique to remedy the approximation errors of the STBA steps for a small trust region, and provide a practical implementation of stochastic graph clustering for constraint sampling. Extensive experiments on Internet SfM data, SLAM data and large-scale data demonstrate the efficiency and scalability of our approach.

Acknowledgement. This work is supported by Hong Kong RGC GRF16206819 & 16203518 and T22-603/15N.

#### References

- 1. Agarwal, S., Mierle, K., Others: Ceres solver. http://ceres-solver.org
- Agarwal, S., Snavely, N., Seitz, S.M., Szeliski, R.: Bundle adjustment in the large. In: ECCV (2010)
- Amestoy, P.R., Davis, T.A., Duff, I.S.: An approximate minimum degree ordering algorithm. SIAM Journal on Matrix Analysis and Applications 17(4), 886–905 (1996)
- 4. Bertsekas, D.P.: Parallel and distributed computation: numerical methods, vol. 23 (1989)
- 5. Bertsekas, D.P.: Constrained optimization and Lagrange multiplier methods. Academic press (2014)
- Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. Journal of statistical mechanics: theory and experiment 2008(10) (2008)
- Calafiore, G., Campi, M.C.: Uncertain convex programs: randomized solutions and confidence levels. Mathematical Programming 102(1), 25–46 (2005)
- Campi, M.C., Garatti, S.: A sampling-and-discarding approach to chanceconstrained optimization: feasibility and optimality. Journal of Optimization Theory and Applications 148(2) (2011)
- Chatterjee, A., Madhav Govindu, V.: Efficient and robust large-scale rotation averaging. In: ICCV (2013)
- Combettes, P.L., Pesquet, J.C.: Proximal splitting methods in signal processing. In: Fixed-point algorithms for inverse problems in science and engineering (2011)
- Darmaillac, Y., Loustau, S.: Mcmc louvain for online community detection. arXiv preprint arXiv:1612.01489 (2016)
- Davis, T.A., Gilbert, J.R., Larimore, S.I., Ng, E.G.: Algorithm 836: Colamd, a column approximate minimum degree ordering algorithm. TOMS 30(3), 377–380 (2004)
- 13. Dellaert, F., Carlson, J., Ila, V., Ni, K., Thorpe, C.E.: Subgraph-preconditioned conjugate gradients for large scale slam. In: IROS (2010)
- Dolan, E.D., Moré, J.J.: Benchmarking optimization software with performance profiles. Mathematical programming 91(2), 201–213 (2002)
- 15. Engels, C., Stewénius, H., Nistér, D.: Bundle adjustment rules
- 16. Eriksson, A., Bastian, J., Chin, T.J., Isaksson, M.: A consensus-based framework for distributed bundle adjustment. In: CVPR (2016)
- Fang, M., Pollok, T., Qu, C.: Merge-sfm: Merging partial reconstructions. In: BMVC (2019)
- 18. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: CVPR (2012)
- Hestenes, M.R., et al.: Methods of conjugate gradients for solving linear systems. Journal of research of the National Bureau of Standards 49(6), 409–436 (1952)
- Jeong, Y., Nister, D., Steedly, D., Szeliski, R., Kweon, I.S.: Pushing the envelope of modern methods for bundle adjustment. PAMI 34(8), 1605–1617 (2011)
- Jian, Y.D., Balcan, D.C., Dellaert, F.: Generalized subgraph preconditioners for large-scale bundle adjustment. In: Outdoor and Large-Scale Real-World Scene Analysis (2012)
- 22. Konolige, K., Garage, W.: Sparse sparse bundle adjustment. In: BMVC (2010)
- Kushal, A., Agarwal, S.: Visibility based preconditioning for bundle adjustment. In: CVPR (2012)

- 16 L. Zhou, Z. Luo, M. Zhen, T. Shen, et al.
- Li, P., Arellano-Garcia, H., Wozny, G.: Chance constrained programming approach to process optimization under uncertainty. Computers & chemical engineering 32(1-2), 25–45 (2008)
- Lourakis, M.I., Argyros, A.A.: Sba: A software package for generic sparse bundle adjustment. TOMS 36(1), 2 (2009)
- 26. Lourakis, M., Argyros, A.A.: Is levenberg-marquardt the most efficient optimization algorithm for implementing bundle adjustment? In: ICCV (2005)
- Marquardt, D.W.: An algorithm for least-squares estimation of nonlinear parameters. Journal of the society for Industrial and Applied Mathematics 11(2), 431–441 (1963)
- Mur-Artal, R., Tardós, J.D.: ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. IEEE Transactions on Robotics 33(5), 1255–1262 (2017)
- Ni, K., Steedly, D., Dellaert, F.: Out-of-core bundle adjustment for large-scale 3d reconstruction. In: ICCV (2007)
- Rotkin, V., Toledo, S.: The design and implementation of a new out-of-core sparse cholesky factorization method. TOMS 30(1), 19–46 (2004)
- Schaeffer, S.E.: Survey: Graph clustering. Computer Science Review 1(1), 27–64 (2007)
- 32. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016)
- 33. Wilson, K., Snavely, N.: Robust global translations with 1dsfm. In: ECCV (2014)
- Wu, C., Agarwal, S., Curless, B., Seitz, S.M.: Multicore bundle adjustment. In: CVPR (2011)
- 35. Zach, C.: Robust bundle adjustment revisited. In: ECCV (2014)
- Zhang, R., Zhu, S., Fang, T., Quan, L.: Distributed very large scale bundle adjustment by global camera consensus. In: ICCV (2017)
- 37. Zhu, S., Shen, T., Zhou, L., Zhang, R., Wang, J., Fang, T., Quan, L.: Parallel structure from motion from local increment to global averaging. arXiv preprint arXiv:1702.08601 (2017)
- Zhu, S., Zhang, R., Zhou, L., Shen, T., Fang, T., Tan, P., Quan, L.: Very large-scale global sfm by distributed motion averaging. In: CVPR (2018)