

Distributed Bundle Adjustment

Karthikeyan Natesan Ramamurthy
IBM Research
knatesa@us.ibm.com

Chung-Ching Lin
IBM Research
cclin@us.ibm.com

Aleksandr Aravkin
University of Washington
saravkin@uw.edu

Sharath Pankanti
IBM Research
sharat@us.ibm.com

Raphael Viguier
University of Missouri-Columbia
rvbb3@missouri.edu

Abstract

Most methods for Bundle Adjustment (BA) in computer vision are either centralized or operate incrementally. This leads to poor scaling and affects the quality of solution as the number of images grows in large scale structure from motion (SfM). Furthermore, they cannot be used in scenarios where image acquisition and processing must be distributed. We address this problem with a new distributed BA algorithm. Our distributed formulation uses alternating direction method of multipliers (ADMM), and, since each processor sees only a small portion of the data, we show that robust formulations improve performance. We analyze convergence of the proposed algorithm, and illustrate numerical performance, accuracy of the parameter estimates, and scalability of the distributed implementation in the context of synthetic 3D datasets with known camera position and orientation ground truth. The results are comparable to an alternate state-of-the-art centralized bundle adjustment algorithm on synthetic and real 3D reconstruction problems. The runtime of our implementation scales linearly with the number of observed points.

1. Introduction

Estimating accurate poses of cameras and locations of 3D scene points from a collection of images obtained by the cameras is a classic problem in computer vision, referred to as *structure from motion* (SfM). Optimizing for the camera parameters and scene points using the corresponding points in images, known as *Bundle Adjustment* (BA), is an important component of SfM [7, 8, 21].

Many recent approaches for BA can be divided into three categories: (a) those that pose BA as non-linear least squares [10, 13, 21], (b) those that decouple the problem in each camera using a triangulation-resection procedure for estimation [15, 18], and (c) those that pose and solve

BA in a linear algebraic formulation [6]. Some important considerations of these methods are reducing the computational complexity by exploiting the structure of the problem [1, 4, 13], incorporating robustness to outlier observations or correspondence mismatches [2, 26], distributing the computations or making the algorithm incremental [9, 11, 23, 24, 5] and making the algorithm insensitive to initial conditions [6]. In this paper, we develop robust distributed BA over camera and scene points.

Our approach is ideally suited for applications where image acquisition and processing must be distributed, such as in a network of unmanned aerial vehicles (UAVs). We assume that each UAV in the network has a camera and a processor; each camera acquires an image of the 3D scene, and the processors in the different UAVs cooperatively estimate the 3D point cloud from the images. Therefore, we use the terms camera, processor, and UAV in an equivalent sense throughout the paper. We also assume that corresponding points from the images are available (possibly estimated using a different distributed algorithm), and are only concerned about estimating the 3D scene points given the correspondences.

Robust approaches, such as [2, 26], are typically used to protect world point and camera parameter estimates from effects of *outliers*, which for BA are incorrect point correspondences that have gone undetected. In contrast, we use robust formulations to *accelerate consensus* in the distributed formulation. Depending on how distribution is achieved, every processor performing computation may see *only a small portion* of the total data, and attempt to use it to infer its local parameters. Small sample means can be extreme, even when the original sample is well-behaved (i.e. even when re-projection errors are truly Gaussian). In the limiting case, each processor may only base its computation on *one data point*, and therefore outliers are guaranteed to occur (from the point of view of individual processors) as an artifact of distributing the computation. Hence we hypothe-

size that using robust losses for penalizing re-projection errors, and quadratic losses for enforcing consensus improves performance.

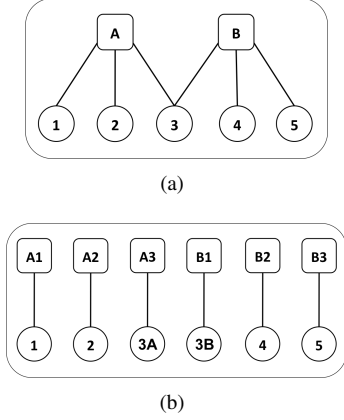


Figure 1: (a) Original configuration of cameras A, B and scene points 1, 2, 3, 4, 5, (b) distributing both the camera parameter and scene point estimation with the constraints $A1 = A2 = A3$, $B1 = B2 = B3$, and $3A = 3B$.

Our proposed robust BA approach supports a natural distributed parallel implementation. We distribute the world points and camera parameters as illustrated for a simple case of 2 cameras and 5 scene points in Figure 1. The algorithm is developed using distributed alternating direction method of multipliers (D-ADMM) [3]. Each processor updates its copy of a set of parameters, while the updated estimates and dual variables ensure consensus. Distributing *both* the world points and the camera parameters yields iterations with $O(l)$ required operations in a serial setting, where l is the total number of 2D observations. In a fully parallel setting, it is possible to bring the time complexities down to $O(1)$ per iteration, a vast improvement compared to traditional and sparse versions of BA, whose complexities are $O((m+n)^3)$ and $O(m^3 + mn)$ respectively [13] (with m and n the number of cameras and 3D scene points). We also exploit the sparsity of the camera network, since not all cameras observe all scene points.

Another optimization-based distributed approach for BA was recently proposed [5]¹. Authors of [5] distributed camera parameters, and performed synthetic experiments using an existing 3D point cloud reconstruction, perturbing it using moderate noise, and generating image points using known camera models. We go further, distributing both world points and camera parameters in a flexible manner, and we implement the entire BA pipeline for 3D reconstruction: performing feature detection, matching corresponding points, and applying the robust distributed D-ADMM BA

¹The initial version of our method was proposed at the same time as [5]

technique in real data settings.

The rest of the paper is organized as follows. We provide background in Section 2, and present the new formulation in Section 3. We show experiments and results on synthetic and real data in Section 4, and conclude with Section 5.

2. Background

2.1. The camera imaging process

We denote the m camera parameter vectors by $\{y_j\}_{j=1}^m$, the n 3D scene points as $\{x_i\}_{i=1}^n$, and the 2D image points as $\{z_{ij}\}$. Each 2D image point $z_{ij} \in \mathbb{R}^2$ is obtained by the transformation and projection of a 3D scene point $x_i \in \mathbb{R}^q$ by the camera $y_j \in \mathbb{R}^p$. BA is an inverse problem, where camera parameters and 3D world points are estimated from the observations $\{z_{ij}\}$. The forward model is a non-linear camera transformation function $f(x_i, y_j)$.

The number of image points is typically much smaller than mn , since not all cameras image all scene points. The camera parameter vector (y_j) usually includes position, Euler angles, and focal length. In this discussion, we assume focal length is known for simplicity, and $y_j \in \mathbb{R}^6$ comprises Euler angles α, β, γ and the translation vector $t \in \mathbb{R}^3$.

Denote the diagonal focal length matrix as $K \in \mathbb{R}^{3 \times 3}$, with the first two diagonal elements set to the focal length and the last element set to 1. The rotation matrix is represented as $R = R_3(\gamma)R_2(\beta)R_1(\alpha)$, where R_1, R_2, R_3 are rotations along the three axes of \mathbb{R}^3 . The camera transformation is now given as $\tilde{z} = Rx + t$. The final 2D image point z is obtained by a perspective projection, with co-ordinates given by

$$z_1 = \frac{\tilde{z}_1}{\tilde{z}_3}, z_2 = \frac{\tilde{z}_2}{\tilde{z}_3}. \quad (1)$$

2.2. Bundle adjustment

Given the 2D points in multiple images that represent the same scene point, BA is typically formulated as a nonlinear least squares problem:

$$\min_{\{x_i\}, \{y_j\}} \sum_{j=1}^m \sum_{i \in S(j)} \|z_{i,j} - f(x_i, y_j)\|_2^2. \quad (2)$$

The set $S(j)$ contains i if the scene point i is imaged by the camera j . The number of unknowns in this objective is $3n + 6m$, and hence it is necessary to have at least this many observations to obtain a good solution; in practice the number of observations is much larger. Problem (2) is solved iteratively, with descent direction (δ_x, δ_y) found by replacing f in (2) by its linearization

$$f(x + \delta_x, y + \delta_y) \approx f(x, y) + J(x)\delta_x + J(y)\delta_y,$$

where $J(x) = \partial_x f, J(y) = \partial_y f$. The Levenberg-Marquardt (LM) algorithm [16] is often used for BA.

The naive LM algorithm requires $O((m+n)^3)$ operations for each iteration, and memory on the order of $O(mn(m+n))$, since we must invert of an $O(m+n) \times O(m+n)$ matrix at each iteration. However, exploiting matrix structure and using the Schur complement approach proposed in [13], the number of arithmetic operations can be reduced to $O(m^3 + mn)$, and memory use to $O(mn)$. Further reduction can be achieved by exploiting secondary sparse structure [10]. The conjugate gradient approaches in [1, 4] can reduce the time complexity to $O(m)$ per iteration, making it essentially linear in the number of cameras.

Another popular approach to reduce the computational complexity involves decoupling of the optimization by explicitly estimating the scene point using back-projection in the *intersection* step and estimating the camera parameters in the *resection* step [18]. The resection step decouples into m independent problems, and hence the overall procedure has a cost of $O(m)$ per iteration. A similar approach, but with the minimization of ℓ_∞ norm of the re-projection error was proposed in [15]. It was shown to be more reliable and degraded gracefully with noise compared to ℓ_2 based BA algorithms. Recently Wu proposed an incremental approach for bundle adjustment [23], where a partial BA or a full BA is performed after adding each camera and associated scene points to the set of unknown parameters, again with a complexity of $O(m)$. We use the ADMM framework to develop our approach.

2.3. Alternating Direction Method of Multipliers

ADMM is a simple and powerful procedure well-suited for distributed optimization [12], see also [3]. In order to understand D-ADMM, consider the objective $h(x) := \sum_{i=1}^n h_i(x)$. We introduce local variables with a *consensus* equality constraint:

$$\min_{\{x_i\}, u} \sum_{i=1}^n h_i(x_i) \quad (3)$$

subject to $x_i - u = 0, i \in \{1, \dots, n\}$.

To solve this problem, we first write down an *augmented Lagrangian* [19]:

$$\bar{l}_\phi(x, u, r, \rho) := \sum_{i=1}^n h_i(x_i) + r_i^T(x_i - u) + \frac{\rho}{2}\phi(x_i, u), \quad (4)$$

where $\rho > 0$ is the penalty parameter, r_i is the Lagrangian multiplier for the constraint, and $\phi(x_i, u)$ is the augmentation term that measures the distance individual variables x_i and the consensus variable u . We then find a saddle point using three steps to update $\{x_i\}$, u , and $\{r_i\}$. Typically $\phi(x_i, u)$ is chosen to the squared Euclidean distance in which case (4) becomes the proximal Lagrangian [19], but other distance or divergence measures can also be used.

3. Algorithmic formulation

3.1. Distributed estimation of scene points and camera parameters

We distribute the estimation among both the scene points and the camera parameters as illustrated in Figure 1. We estimate the camera parameter and the scene point corresponding to each image point independently, and then impose appropriate equality constraints. Eqn. (2) can be written as

$$\min_{\{x_i^j\}, \{y_j^i\}, \{x_i\}, \{y_j\}} \sum_{j=1}^m \sum_{i \in S(j)} \phi_m(z_{i,j} - f(x_i^j, y_j^i)), \quad (5)$$

$$\text{such that } x_i^j = x_i, \forall i, \text{ and } \{j : i \in S(j)\}, \quad (6)$$

$$y_j^i = y_j, \forall j, \text{ and } \{i \in S(j)\}. \quad (7)$$

The augmented Lagrangian, with dual variables r_i^j and s_j^i , is given by

$$\begin{aligned} & \sum_{j=1}^m \sum_{i \in S(j)} \phi_m(z_{i,j} - f(x_i^j, y_j^i)) + r_i^{jT}(x_i^j - x_i) + s_j^{iT}(y_j^i - y_j) \\ & + (\rho_x/2)\phi_a(x_i^j - x_i) + (\rho_y/2)\phi_a(y_j^i - y_j) \end{aligned} \quad (8)$$

Here ϕ_a measures the distance between the distributed world points and their consensus estimates, and distributed camera parameters and their consensus estimates. For ϕ_m we compare squared Euclidean and Huber losses, and ϕ_a is always the squared Euclidean loss.

The ADMM iteration is given by

$$\begin{aligned} (x_i^{j(k+1)}, y_j^{i(k+1)}) &:= \underset{\{x_i^j\}, \{y_j^i\}}{\operatorname{argmin}} \phi_m(z_{i,j} - f(x_i^j, y_j^i)) \\ &+ r_i^{j(k)T}(x_i^j - x_i^{(k)}) + s_j^{i(k)T}(y_j^i - y_j^{(k)}) \\ &+ (\rho_x/2)\phi_a(x_i^j - x_i^{(k)}) + (\rho_y/2)\phi_a(y_j^i - y_j^{(k)}), \end{aligned} \quad (9)$$

$$x_i^{(k+1)} := \frac{1}{|j : i \in S(j)|} \sum_{j: i \in S(j)} \left(x_i^{j(k+1)} + (1/\rho_x)r_i^{j(k)} \right), \quad (10)$$

$$y_j^{(k+1)} := \frac{1}{|i \in S(j)|} \sum_{i \in S(j)} \left(y_j^{i(k+1)} + (1/\rho_y)s_j^{i(k)} \right), \quad (11)$$

$$r_i^{j(k+1)} := r_i^{j(k)} + \rho_x \left(x_i^{j(k+1)} - x_i^{(k+1)} \right), \quad (12)$$

$$s_j^{i(k+1)} := s_j^{i(k)} + \rho_y \left(y_j^{i(k+1)} - y_j^{(k+1)} \right). \quad (13)$$

The equation (9) has to be solved for all $j \in S(i), i \in \{1, \dots, m\}$, and it can be trivially distributed across multiple processes. When ϕ_m is squared ℓ_2 distance, (9) can be

solved using the Gauss-Newton method [17], where we repeatedly linearize f around the current solution and update (x, y) . When ϕ_m is the Huber loss, we use limited memory BFGS (L-BFGS) [17] to update the distributed scene points. Upon convergence, we will obtain the consensus estimates x_i and y_j for all scene points and cameras.

3.1.1 Convergence Analysis

We show that under certain assumptions the proposed D-ADMM algorithm in Section 3.1 converges, using the non-convex and non-smooth framework developed by [22].

Theorem 1 The D-ADMM algorithm proposed in Section 3.1 to the stationary point of the augmented Lagrangian in 8 when:

1. $f(\cdot, \cdot)$ is the perspective camera projection model,
2. ϕ_m is any convex, smooth loss function, and ϕ_a is the squared Euclidean loss.
3. ρ_x and ρ_y are sufficiently large.

Proof Let d_{ij} be the stack of $\{x_i^j, y_j^i\}$, and $\hat{d} = [\hat{d}_{ij}]_{\forall i, \forall j}$. Similarly each pair of consensus variables are stacked as the vector $\hat{c}_{ij} = [x_i^T y_j^T]^T$, and $\hat{c} = [\hat{c}_{ij}]_{\forall i, \forall j}$. \hat{d} and \hat{c} are respectively equivalent to x and y in [22]. We show that the five assumptions (A1-A5) of [22, Thm. 1] are satisfied.

1. Given our assumptions, the objective function in (6) is coercive, i.e., it tends to ∞ as $\hat{d} \rightarrow \infty$ (A1).
2. The feasibility and sub-minimization path conditions are also satisfied since the constraint matrices are easily seen to be full rank (A2-A3).
3. Each additive part of the objective $\phi_m(z_{i,j} - f(x_i^j, y_j^i))$ is restricted prox-regular if ϕ_m is a smooth convex function and f is the perspective camera model. The gradient will be steep when \tilde{z}_3 in (1) is less than some $\epsilon > 0$ and $\phi_m(z_{i,j} - f(x_i^j, y_j^i))$ is prox-regular for $\epsilon > 0$; hence A4 in [22, Thm. 1] holds.
4. Our objective with respect to the consensus variable is identically 0, which is trivially regular (A5).

Since all the assumptions hold, the iterative algorithm in eqns. (9)-(13) converges to a stationary point of the augmented Lagrangian for sufficiently large ρ_x and ρ_y .

3.1.2 Time Complexity

Optimizing (9) takes $O(l)$ time for each round of updates, since (9) must be solved l times, with each solve requiring constant time. The time complexity of the consensus steps for camera parameters and world points given by (10) and (11) are $O(m)$ and $O(n)$ respectively. For the Lagrangian

parameter updates given by (12) and (13), the time complexity is $O(l)$. Hence the dominant time complexity of the proposed algorithm is $O(l)$ for each round. Since the algorithm can be trivially parallelized, the complexity can be brought down to $O(1)$ for each round, if we distribute all the observations to individual processors.

3.1.3 Communication Overhead

Considering a sparse UAV network, assume that each world point is imaged by d cameras. Each camera needs to maintain a copy of the consensus world points x_i . Therefore to update x_i using (10), each camera needs to obtain $d - 1$ individual estimates of x_i^j and send its version of x_i^j to $d - 1$ other cameras. Values r_i^j can be updated locally in each camera, given x_i^j, x_i and previous versions of r_i^j using (12). Hence, for each world point we have a communication overhead of $3(d - 1)d$ floating points per iteration (each world point is a 3D vector). Hence for n world points, the communication overhead is $3(d - 1)dn$ floating points per iteration, where d depends on the distance of the camera from the scene.

3.1.4 Generalized Distributed Estimation

The problem (9) requires each processor to estimate $p + q > 2$ parameters from a single 2D observation. To control the variability of individual estimates as the algorithm proceeds, we generalize the approach to use more than one observation and hence more than one scene point and camera vector during each update step. This generalized step provides flexibility to adjust the number of 3D scene points and cameras based on computational capability of each thread in a CPU or a GPU. We solve

$$\begin{aligned} (X_i^{j(k+1)}, Y_j^{i(k+1)}) &:= \underset{\{X_i^j\}, \{Y_j^i\}}{\operatorname{argmin}} \phi_m(Z_{i,j} - f(X_i^j, Y_j^i)) \\ &+ r_i^{j(k)T} (X_i^j - X_i^{(k)}) + s_j^{i(k)T} (Y_j^i - Y_j^{(k)}) \\ &+ (\rho_x/2)\phi_a(X_i^j - X_i^{(k)}) + (\rho_y/2)\phi_a(Y_j^i - Y_j^{(k)}), \end{aligned} \quad (14)$$

where

$$\begin{aligned} X_i^{j(k+1)} &:= \begin{bmatrix} x_{i_1}^{j(k)} & x_{i_2}^{j(k)} & \dots & x_{i_\pi}^{j(k)} \end{bmatrix}^T, \\ Y_j^{i(k+1)} &:= \begin{bmatrix} y_{j_1}^{i(k)} & y_{j_2}^{i(k)} & \dots & y_{j_\kappa}^{i(k)} \end{bmatrix}^T \end{aligned} \quad (15)$$

4. Experiments

We perform several experiments with synthetic data and real data to show the convergence of the re-projection error and the parameter estimates. We also compare the performance of the proposed approach with the centralized BA algorithm that we implemented using LM. The LM stops

when the re-projection error drops below 10^{-14} , or when the regularization parameter becomes greater than 10^{16} . We implement our distributed approach in a single multi-core computer and not in a sparse UAV network, but our architecture is well-suited for a networked UAV application.

4.1. Synthetic Data

We simulate a realistic scenario, with smooth camera pose transition, and noise parameters consistent with real-world sensor errors. Using the simulation, we evaluate the error in the estimated 3D scene point cloud and the camera parameters, and investigate how estimation error of camera pose affects the final tie points triangulation.

The camera positions are sampled around an orbit, with an average radius 1000m and altitude 1500m, with the camera directed towards a specific area. To each camera pose, a random translation and rotation is added as any real observer cannot move in a perfect circle while steadily aiming always in the same exact direction. The camera path and the 3D scene points for an example scenario are shown in Figure 2. In practice, tie points are usually visible only within a small subset of the available views, and it is generally not practical to try to match all key points within each possible pair of frames. Instead, points are matched within adjacent frames. In our synthetic data, we create artificial occlusions or mis-detection so that each point is only visible on a few consecutive frames.




Figure 2: Camera flight path (blue) and 3D scene points (red) for an example synthetic data set.

4.2. Convergence and Runtime

We investigate convergence of the re-projection error and parameters for D-ADMM BA, comparing the convergence when ϕ_m is squared ℓ_2 vs. Huber in (5), and ϕ_a always the squared ℓ_2 . The number of cameras is 5, the number of scene points is 10, and the number of 2D image points (observations) is 50. We fix the standard deviation for the additive Gaussian noise during the initialization of the camera angles and positions to be 0.1. We vary the standard deviation of noise for the scene points from 0.2 to 1.7. In-

roducing robust losses for misfit penalty helps the convergence of the re-projection error significantly, see Figure 3, (a) vs. (c). This behavior is observed with the convergence of the scene points, see Figures 3, (b) vs. (d), and camera parameters. The Huber penalty is used to guard against outliers; here, outliers come from processors working with limited information. The performance degrades gracefully with noise, see Figure 3, (c) and (d).

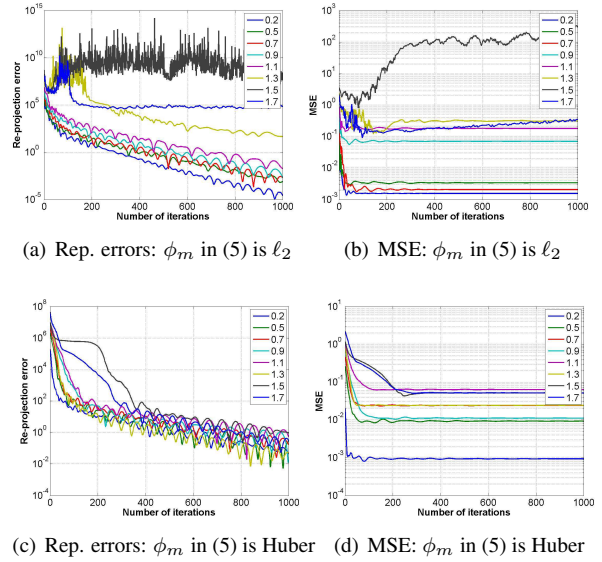


Figure 3: Choosing ϕ_m loss to be Huber penalty leads to better performance in distributed BA, even when there are no outliers in the original data. Panels (a) and (c) compare reprojection errors, while (b) and (d) compare MSE of scene points. In all figures, curves correspond to values σ of scene variance, as shown in the legend. Consensus penalty ϕ_a is always ℓ_2 .

We also compare D-ADMM BA with the centralized LM BA and present the results in Figure 4 (a) and (b). The number of camera parameters and 3D scene points are (10, 40), (15, 100), (25, 100), (30, 200), (100, 200), and (100, 250); with the number of observations increasing as shown on the x-axis of Figure 4. In most settings, D-ADMM BA has a better parameter MSE than centralized LM BA. The runtime of the proposed approach with respect to the number of observations and parallel workers is shown in Figure 4 (c). The parallel workers are configured in MATLAB, and the runtime is linear with respect to the observations and reduces with increasing workers. Our implementation is a simple demonstration of the capability of the algorithm — a fully parallel implementation in a fast language such as C can realize its full potential.

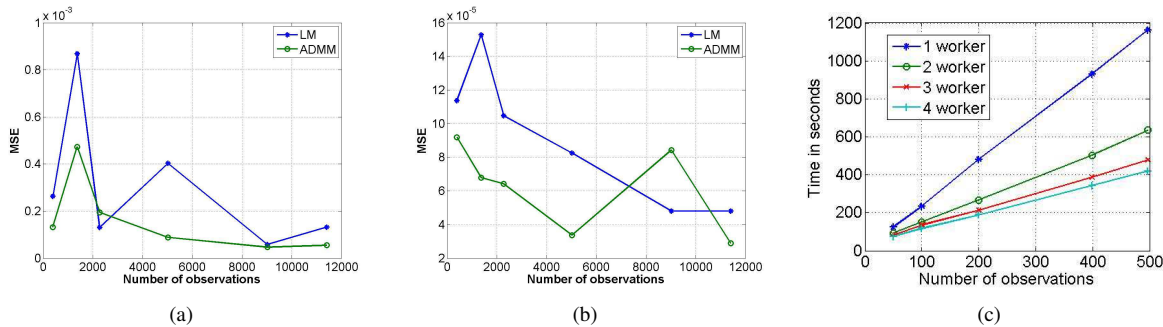


Figure 4: (a) MSE between the actual and estimated camera parameters, (b) MSE between the actual and estimated scene points, (c) runtime of the proposed D-ADMM algorithm with increasing number of processor cores.

4.3. Real Data

To demonstrate the performance of D-ADMM BA, we conducted experiments on real datasets with different settings. All experiments are done with MATLAB on a PC with a 2.7 GHz CPU and 16 GB RAM.

In our SFM pipeline, SIFT feature points [14] are used for detection and matching. The relative fundamental matrices are estimated for each pair of images with sufficient corresponding points, which are used to estimate relative camera pose and 3D structure. Next, the relative parameters are used to generate the global initial values for BA. The datasets were downloaded from the Princeton Vision Group and the EPFL Computer Vision Lab [20].

Since there are no ground truth 3D structures available for the real datasets, we compare the dense reconstruction results obtained using the method of [25]. The first dataset has five images and a sample image is shown in Figure 5 (a). After keypoint detection and matching, centralized LM BA and D-ADMM BA are given the same input. There are a total of 104 world points and 252 observations. The final re-projection error of LM and D-ADMM are 0.93 and 0.67 respectively. Figure 5 (c) and (d) shows that the dense reconstruction quality of LM and the D-ADMM are similar. Figure 5 (b) shows the convergence of re-projection error for the D-ADMM algorithm. Figure 6 (a) shows the convergence of re-projection error for different values of $\rho = \rho_x = \rho_y$. Setting ρ to a high value accelerates convergence.

We also estimate camera parameters and scene points, applying the approach of Section 3.1.4 to the same data set. Figure 6 (b) shows that as the number of scene points per iteration increase, the runtime decreases, with 32 scene points per iteration giving the fastest convergence, see figure 6 (c). Figure 6 (d) compares re-projection errors with different number of cameras in each iteration. Initial values are the same as in the castle-P30 experiment (Figure 11),

and the number of scene points in each iteration is 64. Re-projection errors decrease faster as the number of cameras in each iteration increases.

We perform distributed BA on the Herz-Jesu dataset data set provided in [20] using the approach in Section 3.1.4. This data set has seven images, 1140 world points, and 2993 observations. In this experiment, the LM BA algorithm using the same setting as in previous experiments does not converge and has the final re-projection error about 2500. Therefore, the dense reconstruction result is not presented. D-ADMM BA with eight scene points in each update step has a final re-projection error of 0.76. Figure 7(b) shows the dense 3D point cloud estimated with D-ADMM BA.

Additional results on other datasets (fountain-P11, entry-P10, Herz-Jesu-P25, and castle-P30) are presented in Table 1, Figure 8, 9, 10 and 11. σ is mean re-projection error. Figure 8, 9, 10 and 11 present different perspectives of the dense reconstruction results to show the robustness of 3D parameter estimations.

Table 1: The dataset information and experiment results.

Dataset	Images	Scene pts	Obs	σ
fountain-P11	11	1346	3859	0.5
entry-P10	10	1382	3687	0.7
Herz-Jesu-P25	25	2161	5571	0.87
castle-P30	30	2383	6453	0.84

Settings are fixed across experiments, and the maximum iteration counter is set to 1600. The experiments on fountain-P11 and Herz-Jesu-P25 dataset (Figure 8 and 10) have better dense reconstruction results since there are more images covering the same regions. The real data experiments show D-ADMM BA achieves similar objective values (mean re-projection error < 1) as the number of observations increases; it is not necessary to increase the number

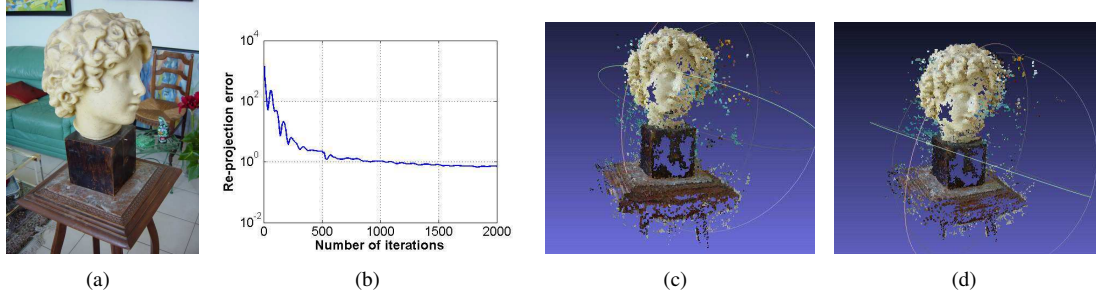


Figure 5: (a) Original 2D image, (b) re-projection error for D-ADMM BA, (c) dense 3D point cloud estimated with LM BA (mean re-projection error = 0.93), (d) dense 3D point cloud estimated using D-ADMM BA (mean re-projection error = 0.67).




Figure 6: (a) The re-projection error for different values of ρ , using generalized distribution approach, (b) runtime of D-ADMM BA (c) re-projection errors with increasing number of scene points, (d) re-projection errors for multiple cameras per estimation vector.

of iterations as the size of the data increases. D-ADMM BA scales linearly with the number of observations and can be parallelized on GPU clusters.

5. Conclusions

We presented a new distribution algorithm for bundle adjustment, D-ADMM BA, which compares well to centralized approaches in terms of performance and scales well for SfM. Experimental results demonstrated the importance of robust formulations for improved convergence in the distributed setting. Even when there are no outliers in the



Figure 7: (a) Original 2D image, (b) dense 3D point cloud estimated with D-ADMM BA (mean re-proj. error = 0.76).

initial data, robust losses are helpful because estimates of processors working with limited information can stray far from the aggregate estimates, see Figure 3. Formulation design for distributed optimization may yield further improvements; this is an interesting direction for future work.

Results obtained with D-ADMM BA are comparable to those obtained with state-of-the-art centralized LM BA, and D-ADMM BA scales linearly in runtime with respect to the number of observations. Our approach is well-suited for use in a networked UAV system, where distributed computation is an essential requirement.




Figure 8: Reconstructed Fountain-P11 views (11 images, 1346 world points, 3859 obs., mean re-proj. error = 0.5).




Figure 9: Reconstructed Entry-P10 views (10 images, 1382 world points, 3687 obs., mean re-proj. error = 0.7).




Figure 10: Reconstructed Herz-Jesu-P25 views (25 images, 2161 world points, 5571 obs., mean re-proj. error = 0.87).




Figure 11: Castle-P30 (30 images, 2383 world points, 6453 obs., mean re-proj. error = 0.84).

References

- [1] S. Agarwal, N. Snavely, S. M. Seitz, and R. Szeliski. Bundle adjustment in the large. In *Computer Vision—ECCV 2010*, pages 29–42. Springer, 2010.
- [2] A. Aravkin, M. Styer, Z. Moratto, A. Nefian, and M. Broxton. Student’s t robust bundle adjustment algorithm. In *Image Processing (ICIP), 2012 19th IEEE International Conference on*, pages 1757–1760. IEEE, 2012.
- [3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [4] M. Byröd and K. Åström. Conjugate gradient bundle adjustment. In *Computer Vision—ECCV 2010*, pages 114–127. Springer, 2010.
- [5] A. Eriksson, J. Bastian, T.-J. Chin, and M. Isaksson. A consensus-based framework for distributed bundle adjustment. In *Computer Vision and Pattern Recognition, 2015. CVPR 2015. IEEE Conference on*. Ieee, 2015.
- [6] A. Fusiello and F. Crosilla. Solving bundle block adjustment by generalized anisotropic procrustes analysis. *ISPRS Journal of Photogrammetry and Remote Sensing*, 102:209–221, 2015.
- [7] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [8] J. Heinly, J. L. Schonberger, E. Dunn, and J.-M. Frahm. Reconstructing the world* in six days*(as captured by the yahoo 100 million image dataset). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3287–3295, 2015.
- [9] V. Indelman, R. Roberts, C. Beall, and F. Dellaert. Incremental light bundle adjustment. In *Proceedings of the British Machine Vision Conference (BMVC 2012)*, pages 3–7, 2012.
- [10] K. Konolige and W. Garage. Sparse sparse bundle adjustment. In *BMVC*, pages 1–11. Citeseer, 2010.
- [11] J. Kopf, M. F. Cohen, and R. Szeliski. First-person hyperlapse videos. *ACM Transactions on Graphics (TOG)*, 33(4):78, 2014.
- [12] P.-L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
- [13] M. I. Lourakis and A. A. Argyros. Sba: A software package for generic sparse bundle adjustment. *ACM Transactions on Mathematical Software (TOMS)*, 36(1):2, 2009.
- [14] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999.
- [15] K. Mitra and R. Chellappa. A scalable projective bundle adjustment algorithm using the l infinity norm. In *Computer Vision, Graphics & Image Processing, 2008. ICVGIP’08. Sixth Indian Conference on*, pages 79–86. IEEE, 2008.
- [16] J. J. Moré. The levenberg-marquardt algorithm: implementation and theory. In *Numerical analysis*, pages 105–116. Springer, 1978.
- [17] J. Nocedal and S. Wright. *Numerical optimization*. Springer Series in Operations Research. Springer, 1999.
- [18] M. D. Pritt. Fast orthorectified mosaics of thousands of aerial photographs from small uavs. In *Applied Imagery Pattern Recognition Workshop (AIPR), 2014 IEEE*, pages 1–8. IEEE, 2014.
- [19] R. T. Rockafellar and R. J.-B. Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- [20] C. Strecha, W. von Hansen, L. V. Gool, P. Fua, and U. Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. Ieee, 2008.
- [21] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment - a modern synthesis. In *Vision algorithms: theory and practice*, pages 298–372. Springer, 2000.
- [22] Y. Wang, W. Yin, and J. Zeng. Global convergence of admm in nonconvex nonsmooth optimization. *arXiv preprint arXiv:1511.06324*, 2015.
- [23] C. Wu. Towards linear-time incremental structure from motion. In *3D Vision-3DV 2013, 2013 International Conference on*, pages 127–134. IEEE, 2013.
- [24] C. Wu, S. Agarwal, B. Curless, and S. M. Seitz. Multicore bundle adjustment. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3057–3064. IEEE, 2011.
- [25] J. Xiao, J. Chen, D.-Y. Yeung, and L. Quan. Learning two-view stereo matching. In *Computer Vision—ECCV 2008*, pages 15–27. Springer, 2008.
- [26] J. Zhang, M. Boutin, and D. G. Aliaga. Robust bundle adjustment for structure from motion. In *Image Processing, 2006 IEEE International Conference on*, pages 2185–2188. IEEE, 2006.