

PRACTICAL TIME BUNDLE ADJUSTMENT FOR 3D RECONSTRUCTION ON THE GPU

Siddharth Choudhary (IIIT Hyderabad), Shubham Gupta (IIIT Hyderabad), P J Narayanan (IIIT Hyderabad)

Abstract



Takes around 9 seconds to perform one iteration on GPU for 488 cameras, giving it a speedup of around 10 times

Outline

- Motivation
- Related Work
- Problem Statement
- What is Bundle Adjustment ?
- Sparse Bundle Adjustment on the GPU
- Results and Analysis
- Preliminary Results on Fermi
- Summary
- Future Work

Motivation



Related Work

- Building Rome in a Day (ICCV 2009)
 - Uses 500 computer cores to maximize parallelization in the SFM pipeline
- Building Rome in a Cloudless Day (ECCV 2010)

Problem Statement

The goal is to develop a practical time implementation of Bundle Adjustment by exploiting all computing resources of the CPU and the GPU

What is Bundle Adjustment ?

Objective Function:

$$\min_{P,X} \sum_{i=1}^{n} \sum_{j=1}^{m} d(Q(P_j, X_i), x_{ij})^2$$

 $Q(P_j, X_i)$ is the predicted projection of point *i* on image *j* d(x, y) is the Euclidean distance between the image points represented by x and y

Minimizing this function is a sparse non linear least squares problem

What is Bundle Adjustment ?

The non-linear least squares objective function is solved using LM Algorithm which is an interpolation of Gauss Newton and Gradient descent iteration.

The normal equation to be solved during each linear LM iteration is given as:

$$\left(J^T \Sigma_X^{-1} J + \mu I\right) \delta = J^T \Sigma_X^{-1} \epsilon ,$$
Where $J = \frac{\partial X}{\partial P}$, $\epsilon = X - \hat{X}$ and μ is the damping factor

Data Structure for the SBA



Compressed Column Storage of Visibility Mask having 4 Cameras and 4 3D Points. Each CUDA Block processes one set of 3D points.

Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

Compute Hessian Matrix and augment it ($J^T \Sigma_x^{-1} J + \mu I$)

Compute Schur Complement to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

Compute $J^T \epsilon$

Compute $\epsilon_a - WV^{*-1}\epsilon_b$

Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

Compute Hessian Matrix and augment it ($J^T \Sigma_x^{-1} J + uI$)

Compute Schur Complement to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

Compute $\mathbf{J}^{\mathrm{T}} \boldsymbol{\epsilon}$

Compute $\epsilon_a - WV^{*-1}\epsilon_b$

- Computation of the Predicted Projection and Error Vector
 - For *m* cameras, m blocks are launched, with block *j* computing projections corresponding to camera *j*.
 - The computation is limited by the number of registers available per block and a maximum limit of number of threads per block.



Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

Compute Hessian Matrix and augment it ($J^T \Sigma_x^{-1} J + uI$)

Compute Schur Complement to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

Compute $\mathbf{J}^{\mathrm{T}} \boldsymbol{\epsilon}$

Compute $\epsilon_a - WV^{*-1}\epsilon_b$

Computation of the Jacobian Matrix and L2 Error

The Jacobian Matrix is calculated as $\frac{\partial X}{\partial P}$, with a grid structure similar to the computation of initial projections.



 $\mathbf{J} = (A_{10}, B_{10}, A_{20}, B_{20}, A_{01}, B_{01}, A_{31}, B_{31}, A_{12}, B_{12}, A_{32}, B_{32}, A_{03}, B_{03}, A_{13}, B_{13})$

Block i computes the $A_{ij} \& B_{ij}$ submatrices corresponding the jth camera

Computation of the Jacobian Matrix and L2 Error

- The Jacobian Matrix is calculated as $\frac{\partial X}{\partial P}$, with a grid structure similar to the computation of initial projections.
- While the GPU is computing Jacobian Matrix, CPU computes the L2 Error using the Error vector

Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

Compute Hessian Matrix and augment it $(J^T \Sigma_x^{-1} J + \mu I)$

Compute Schur Complement to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

Compute $J^T \epsilon$

Compute $\epsilon_a - WV^{*-1}\epsilon_b$

- $\hfill\square$ Computation of the Hessian Matrix and $J^T \varepsilon$
 - The Hessian Matrix is computed using three different and independent kernels, two of which computes the diagonal sub-matrices and one of them computes the non diagonal sub-matrix



Computation of U

- The grid structure consists of m blocks with each block processing one Uj where j is the block id
- Summation is done using segmented scan



Computation of V

Computation of V is done in a way similar to U, with each block computing Vi for the ith point

$$\mathbf{V}_i = \sum_j \mathbf{B}_{ij}^T \mathbf{\Sigma}_{x_{ij}}^{-1} \mathbf{B}_{ij}$$



Computation of W

- Computation of Ws are independent of each other
- nnz/10 blocks are launched with each block computing 10 Ws

$$\mathbf{W}_{ij} = \mathbf{A}_{ij}^T \boldsymbol{\Sigma}_{x_{ij}}^{-1} \mathbf{B}_{ij}$$



$\hfill\square$ Computation of the Hessian Matrix and $J^T \varepsilon$

- The Hessian Matrix is computed using three different and independent kernels, two of which computes the diagonal sub-matrices and one of them computes the non diagonal sub-matrix
- After U,V,W is calculated, the augmentation is done using another kernel, which calculates µ and add to the diagonal terms.
- While the H matrix is computation is done on GPU, J^T ∈ computation is done on CPU.

Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

Compute Hessian Matrix and augment it ($J^T \Sigma_x^{-1} J + \mu I$)

Compute Schur Complement to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

Compute $\mathbf{J}^{\mathrm{T}} \boldsymbol{\epsilon}$

Compute $\epsilon_a - WV^{*-1}\epsilon_b$

- Computation of the Schur Complement to from the Reduced Camera System
 - It is the most computationally expensive step of all the modules.
 - S is a symmetric matrix. So, we calculate only the upper diagonal of S.
 - The CUDA grid structure consists of $m \times m$ blocks with each block computing a 9×9 sub-matrix in the upper diagonal
 - The computation is limited by the amount of shared memory available and number of registers available per block.

Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

Compute Hessian Matrix and augment it ($J^T \Sigma_x^{-1} J + uI$)

Compute Schur Complement to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

Compute $\mathbf{J}^{\mathrm{T}} \boldsymbol{\epsilon}$

Compute $\epsilon_a - WV^{*-1}\epsilon_b$

- \square Computation of the Inverse of the Reduced Camera System and computation of δ_a
 - We use Cholesky Decomposition implemented in MAGMA library in order to compute the inverse.
 - It benefits from the hybrid computation by using both CPUs and GPUs and achieves a speedup of over 10 times than CPU.

Compute the Predicted Projection \hat{x}_{ij} and Error Vectors $\epsilon_{ij} = x_{ij} - \hat{x}_{ij}$

Compute Jacobian Matrix (J)

GPU

Compute Schur Complement

 μ I)

to form Reduced Camera System

Compute Inverse of Reduced Camera System and calculate δ

Compute L2 Error using Error Vectors

CPU

Compute $e_j = \epsilon_{a_j} - \sum_i Y_{ij} \epsilon_{b_i}$

Hybrid Computation



Computation blocks are efficiently scheduled either on CPU or GPU. Arrows show the data dependency between various modules on CPU and GPU. Modules connected through a vertical line are computed in parallel on CPU and GPU

Results and Analysis



Timings shown are the time taken by each component either on CPU or GPU including the memory transfer time in one iteration for 488 cameras.

Results and Analysis



Time(sec) taken for each step in one iteration of Bundle Adjustment on GPU and CPU for various number of cameras. Total time is the time taken by hybrid implementation of BA using CPU and GPU in parallel

Results and Analysis



Time Comparison (one iteration) of Bundle Adjustment Computation on CPU and GPU

Preliminary Results on Fermi





We introduced a hybrid algorithm using the GPU and the CPU to perform practical time bundle adjustment.

We achieve a speedup of around 8 – 10 times over the CPU implementation on one quarter of Nvidia Tesla S1070 GPU

Future Work

We are adapting our approach to the Fermi and expecting significant speedups on it.

A multi GPU implementation is also being explored for faster overall processing

Thank You