TUE-AM-027

Power Bundle Adjustment for Large-Scale 3D Reconstruction

Simon Weber[1,2], Nikolaus Demmel[1,2], Tin Chon Chan[1,2], Daniel Cremers[1,2,3]

[1]Technical University of Munich, [2]Munich Center for Machine Learning, [3]University of Oxford

Context

- The BA problem refers to the **joint estimation** of camera parameters and 3D landmark positions.
- It is the core component in **many 3D reconstruction and Structure from Motion applications**.
- Building accurate city-scale maps for applications such as augmented reality or autonomous driving brings current BA approaches to their limits.

(a) *Ladybug-1197*

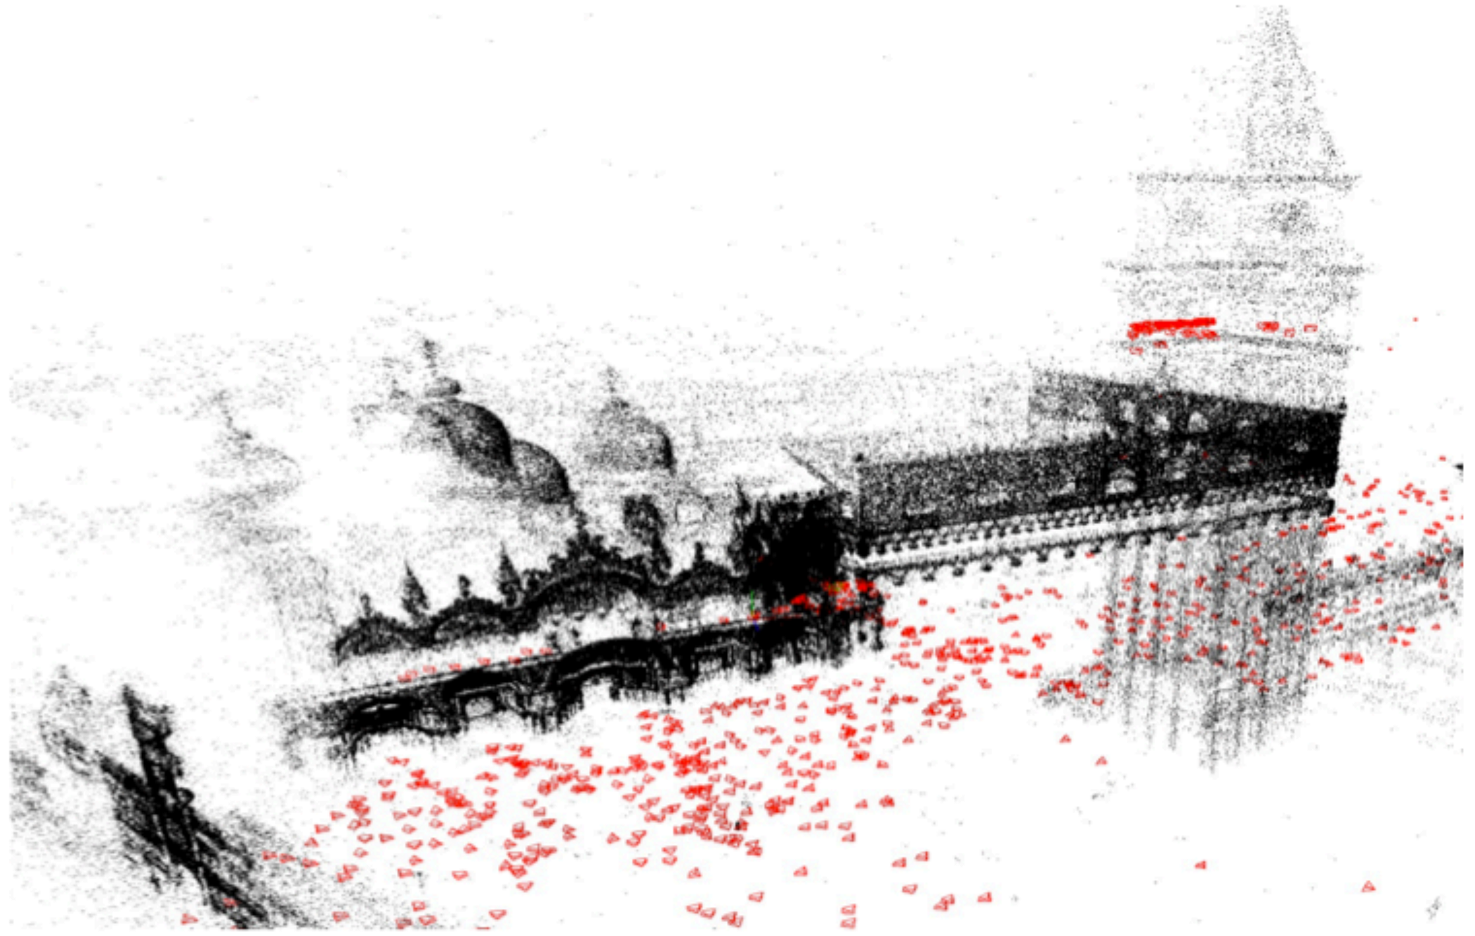

(b) *Venice-1102*

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Highlights

- We propose to solve **large-scale bundle adjustment** problem with a **power series** expansion of the Schur complement.

- We **theoretically prove** the convergence of our approach.

- We **experimentally** show that PoBA **significantly accelerates** the solution of the normal equation compared to state-of-the-art iterative methods, **even for very high accuracy**.

(a) *Ladybug-1197*

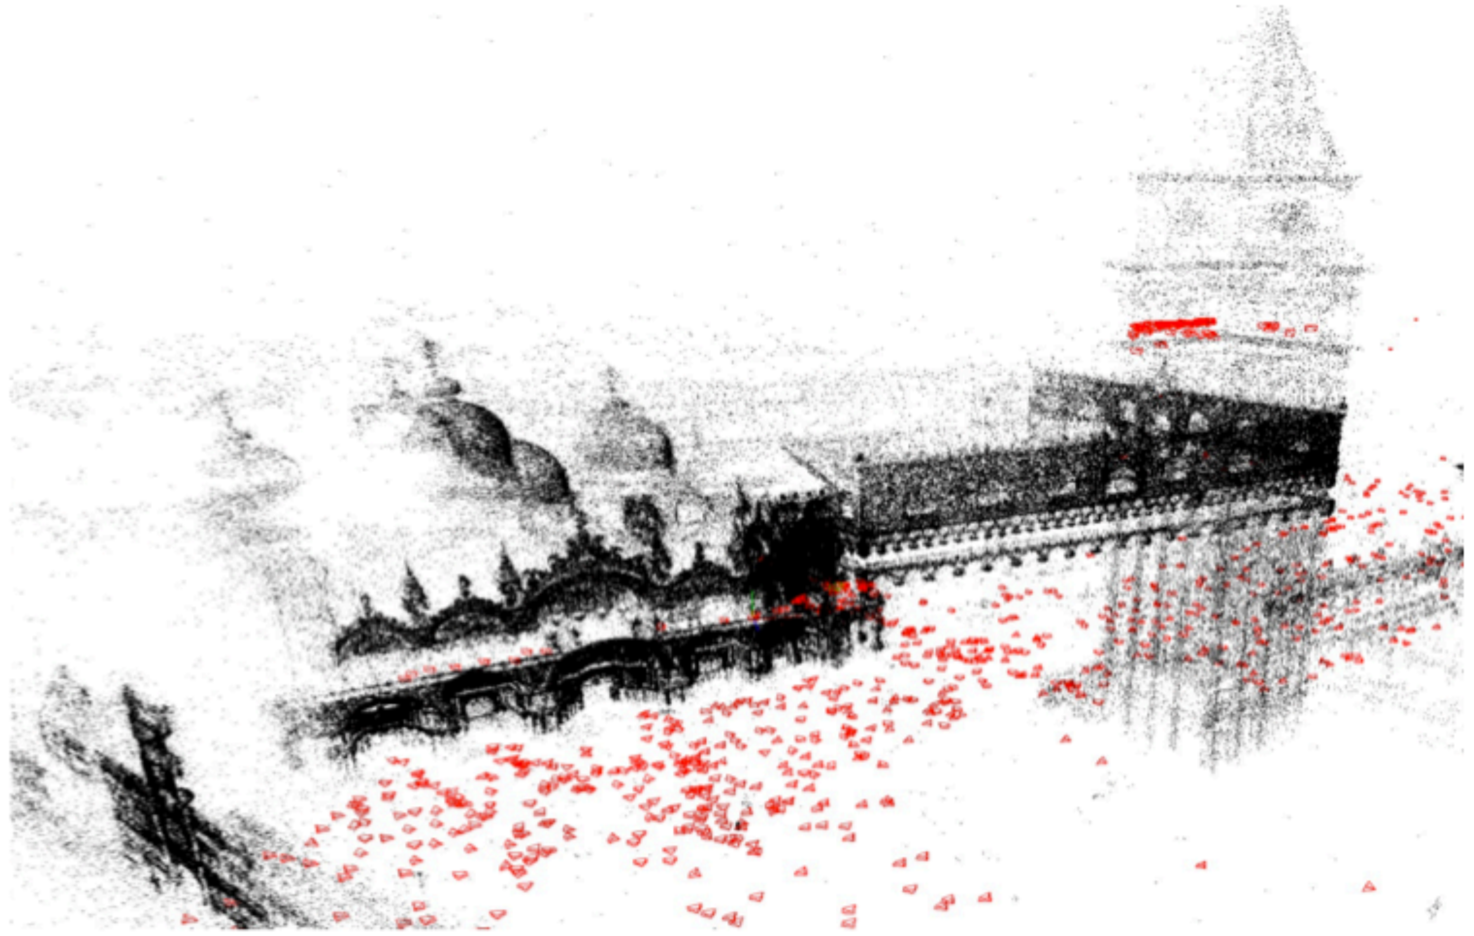

(b) *Venice-1102*

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Series

Proposition 1

Let M be a $n \times n$ matrix.

Power Series

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Proposition 1

Let M be a $n \times n$ matrix.

If $||M|| < 1$, then

Power Series

Proposition 1

Let M be a $n \times n$ matrix.

If $\|M\| < 1$, then

$$(I - M)^{-1} = \sum_{i=0}^{m} M^i + R \,.$$

$$R = \sum_{i=m+1}^{\infty} M^i$$

where $\|R\| \leq \dfrac{\|M\|^{m+1}}{1 - \|M\|}$.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$F(x) = \|r(x)\|^2 = \sum_i \|r_i(x)\|^2$$

residual

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$F(x) = \|r(x)\|^2 = \sum_i \|r_i(x)\|^2$$

$$r(x) \approx r^0 + J\Delta x$$

$$\min_{\Delta x_p, \Delta x_l} \left(\|r^0 + \begin{pmatrix} J_p & J_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix}\|^2 + \lambda \| \begin{pmatrix} D_p & D_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix}\|^2 \right)$$

pose/landmark delta

pose/landmark Jacobians

regularization

pose/landmark damping matrices

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$\min_{\Delta x_p, \Delta x_l} \left(\| r^0 + \begin{pmatrix} J_p & J_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 + \lambda \| \begin{pmatrix} D_p & D_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 \right)$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$\min_{\Delta x_p, \Delta x_l} \left(\| r^0 + \left(J_p \; J_l \right) \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 + \lambda \| \left(D_p \; D_l \right) \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 \right)$$

Lagrangian

$$\begin{pmatrix} U_\lambda & W \\ W^\top & V_\lambda \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} = - \begin{pmatrix} b_p \\ b_l \end{pmatrix}$$

$$U_\lambda = J_p^\top J_p + \lambda D_p^\top D_p \qquad b_p = J_p^\top r^0$$

$$V_\lambda = J_l^\top J_l + \lambda D_l^\top D_l \qquad b_l = J_l^\top r^0$$

$$W = J_p^\top J_l$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$\min_{\Delta x_p, \Delta x_l} \left(\| r^0 + \begin{pmatrix} J_p & J_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 + \lambda \| \begin{pmatrix} D_p & D_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 \right)$$

↓ Lagrangian

Normal equation

$$\begin{pmatrix} U_\lambda & W \\ W^\top & V_\lambda \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} = - \begin{pmatrix} b_p \\ b_l \end{pmatrix}$$

$$U_\lambda = J_p^\top J_p + \lambda D_p^\top D_p \qquad b_p = J_p^\top r^0$$

$$V_\lambda = J_l^\top J_l + \lambda D_l^\top D_l \qquad b_l = J_l^\top r^0$$

$$W = J_p^\top J_l$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$\min_{\Delta x_p, \Delta x_l} \left(\| r^0 + \begin{pmatrix} J_p & J_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 + \lambda \| \begin{pmatrix} D_p & D_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 \right)$$

↓ Lagrangian

$$\begin{pmatrix} U_\lambda & W \\ W^\top & V_\lambda \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} = - \begin{pmatrix} b_p \\ b_l \end{pmatrix}$$

$$U_\lambda = J_p^\top J_p + \lambda D_p^\top D_p \qquad b_p = J_p^\top r^0$$

$$V_\lambda = J_l^\top J_l + \lambda D_l^\top D_l \qquad b_l = J_l^\top r^0$$

$$W = J_p^\top J_l$$

↓ Schur complement trick

$$S \Delta x_p = - \tilde{b}$$

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

$$\tilde{b} = b_p - W V_\lambda^{-1} b_l$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$\min_{\Delta x_p, \Delta x_l} \left(\| r^0 + \begin{pmatrix} J_p & J_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 + \lambda \| \begin{pmatrix} D_p & D_l \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 \right)$$

Lagrangian

$$\begin{pmatrix} U_\lambda & W \\ W^\top & V_\lambda \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} = - \begin{pmatrix} b_p \\ b_l \end{pmatrix}$$

$$U_\lambda = J_p^\top J_p + \lambda D_p^\top D_p \qquad b_p = J_p^\top r^0$$

$$V_\lambda = J_l^\top J_l + \lambda D_l^\top D_l \qquad b_l = J_l^\top r^0$$

$$W = J_p^\top J_l$$

Schur complement trick

Reduced camera system

$$S \Delta x_p = - \tilde{b}$$

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

$$\tilde{b} = b_p - W V_\lambda^{-1} b_l$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Bundle Adjustment

$$\min_{\Delta x_p, \Delta x_l} \left(\| r^0 + \left(J_p \; J_l \right) \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 + \lambda \| \left(D_p \; D_l \right) \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} \|^2 \right)$$

Lagrangian

$$\begin{pmatrix} U_\lambda & W \\ W^\top & V_\lambda \end{pmatrix} \begin{pmatrix} \Delta x_p \\ \Delta x_l \end{pmatrix} = - \begin{pmatrix} b_p \\ b_l \end{pmatrix}$$

$$U_\lambda = J_p^\top J_p + \lambda D_p^\top D_p \qquad b_p = J_p^\top r^0$$

$$V_\lambda = J_l^\top J_l + \lambda D_l^\top D_l \qquad b_l = J_l^\top r^0$$

$$W = J_p^\top J_l$$

Schur complement trick

$$S \Delta x_p = - \widetilde{b}$$

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

$$\widetilde{b} = b_p - W V_\lambda^{-1} b_l$$

back-substitution

$$\Delta x_l = - V_\lambda^{-1} \left(-b_l + W^\top \Delta x_p \right)$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Bundle Adjustment (PoBA)

How to link **power series** theory to the **bundle adjustment problem**?

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Bundle Adjustment (PoBA)

How to link **power series** theory to the **bundle adjustment problem**?

PoBA

Expand the inverse Schur complement into a power series.

Power Schur Complement

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

$$\downarrow$$

$$S = U_\lambda (I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

$$\downarrow$$

$$S = U_\lambda (I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)$$

$$\downarrow$$

$$S^{-1} = (I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)^{-1} U_\lambda^{-1}$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

$$S = U_\lambda - W V_\lambda^{-1} W^\top$$

↓

$$S = U_\lambda (I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)$$

↓

$$S^{-1} = (I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)^{-1} U_\lambda^{-1}$$

May we expand $(I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)^{-1}$ into a power series?

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

Lemma

Let μ be an eigenvalue of $U_\lambda^{-1} W V_\lambda^{-1} W^\top$.

Then $\mu \in [0,1[$.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

Lemma

Let μ be an eigenvalue of $U_\lambda^{-1} W V_\lambda^{-1} W^\top$.

Then $\mu \in [0,1[$.

Consequence:

Proposition 1 can be applied to the matrix $U_\lambda^{-1} W V_\lambda^{-1} W^\top$.

$$S^{-1} = (I - U_\lambda^{-1} W V_\lambda^{-1} W^\top)^{-1} U_\lambda^{-1}$$

$$\tilde{S}_{-1}(m) = \sum_{i=0}^{m} (U_\lambda^{-1} W V_\lambda^{-1} W^\top)^i U_\lambda^{-1}.$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

$$S\Delta x_p = -\tilde{b} \longrightarrow x(m) = -\tilde{S}_{-1}(m)\tilde{b}\,.$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

$$S\Delta x_p = -\tilde{b} \longrightarrow x(m) = -\tilde{S}_{-1}(m)\tilde{b}\,.$$

Proposition 2: Convergence result

$$\|x(m) - \Delta x_p\|_2 \underset{m\to+\infty}{\longrightarrow} 0\,.$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

$$S\Delta x_p = -\tilde{b} \longrightarrow x(m) = -\tilde{S}_{-1}(m)\tilde{b}\,.$$

Proposition 2: Convergence result

$$\|x(m) - \Delta x_p\|_2 \underset{m\to+\infty}{\longrightarrow} 0\,.$$

Consequences:

- an approximation of Δx_p is directly obtained by applying $\tilde{S}_{-1}(m)$ to the right hand side of $S\Delta x_p = -\tilde{b}$;
- the accuracy of this approximation only depends on the order m and can be as good as possible.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

How to chose the order m?

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Power Schur Complement

How to chose the order m?

We set a **stop criterion:**

$$(i+1)\frac{\|x(i) - x(i-1)\|_2}{\|x(i)\|_2} < \epsilon\,.$$

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Experiments

- BAL [1] dataset.
- Two kinds of experiments:

    1. Comparison with competitive iterative solvers $\sqrt{BA}$ [2] and Ceres.
    2. Extension of the distributed stochastic framework $STBA$ [3].

[1] Agarwal et al., *Bundle Adjustment in the Large*, ECCV10.

[2] Demmel et al., *Square Root Bundle Adjustment for Large-Scale Reconstruction*, CVPR21.

[3] Zhou et al., Stochastic Bundle Adjustment for Efficient and Scalable 3D Reconstruction, ECCV20.

Results

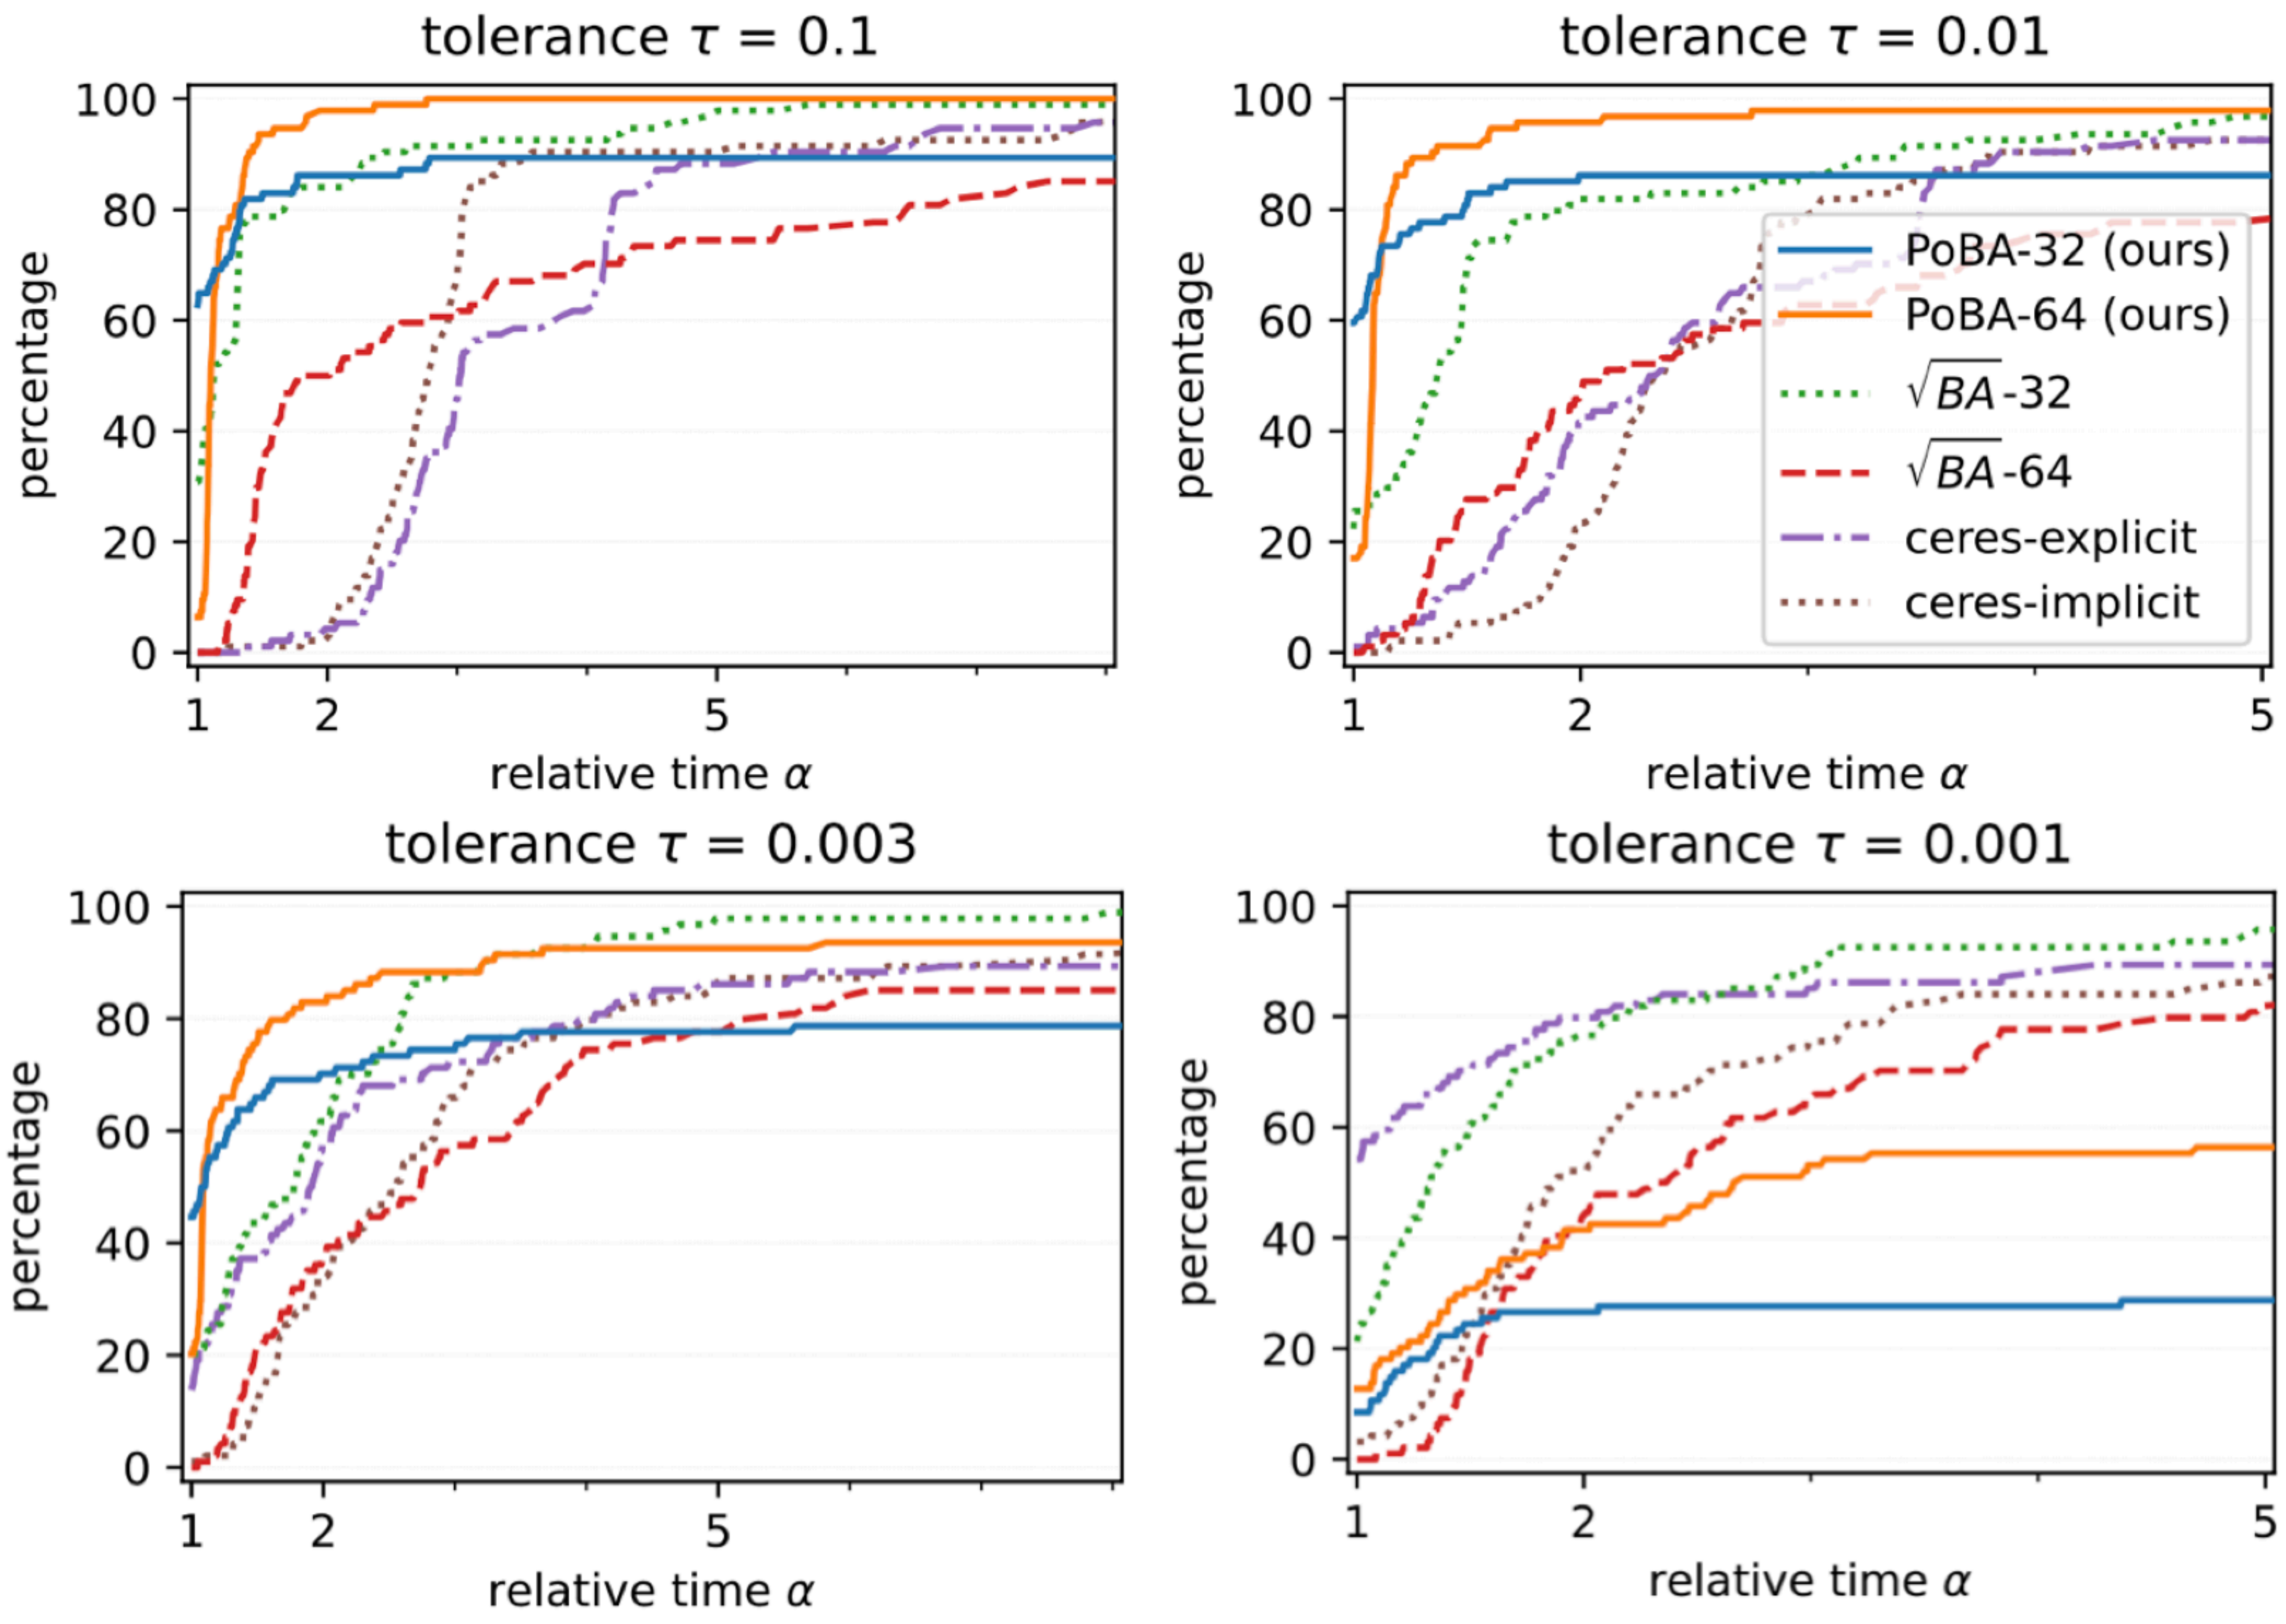

Performance profiles for all 97 BAL problems.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

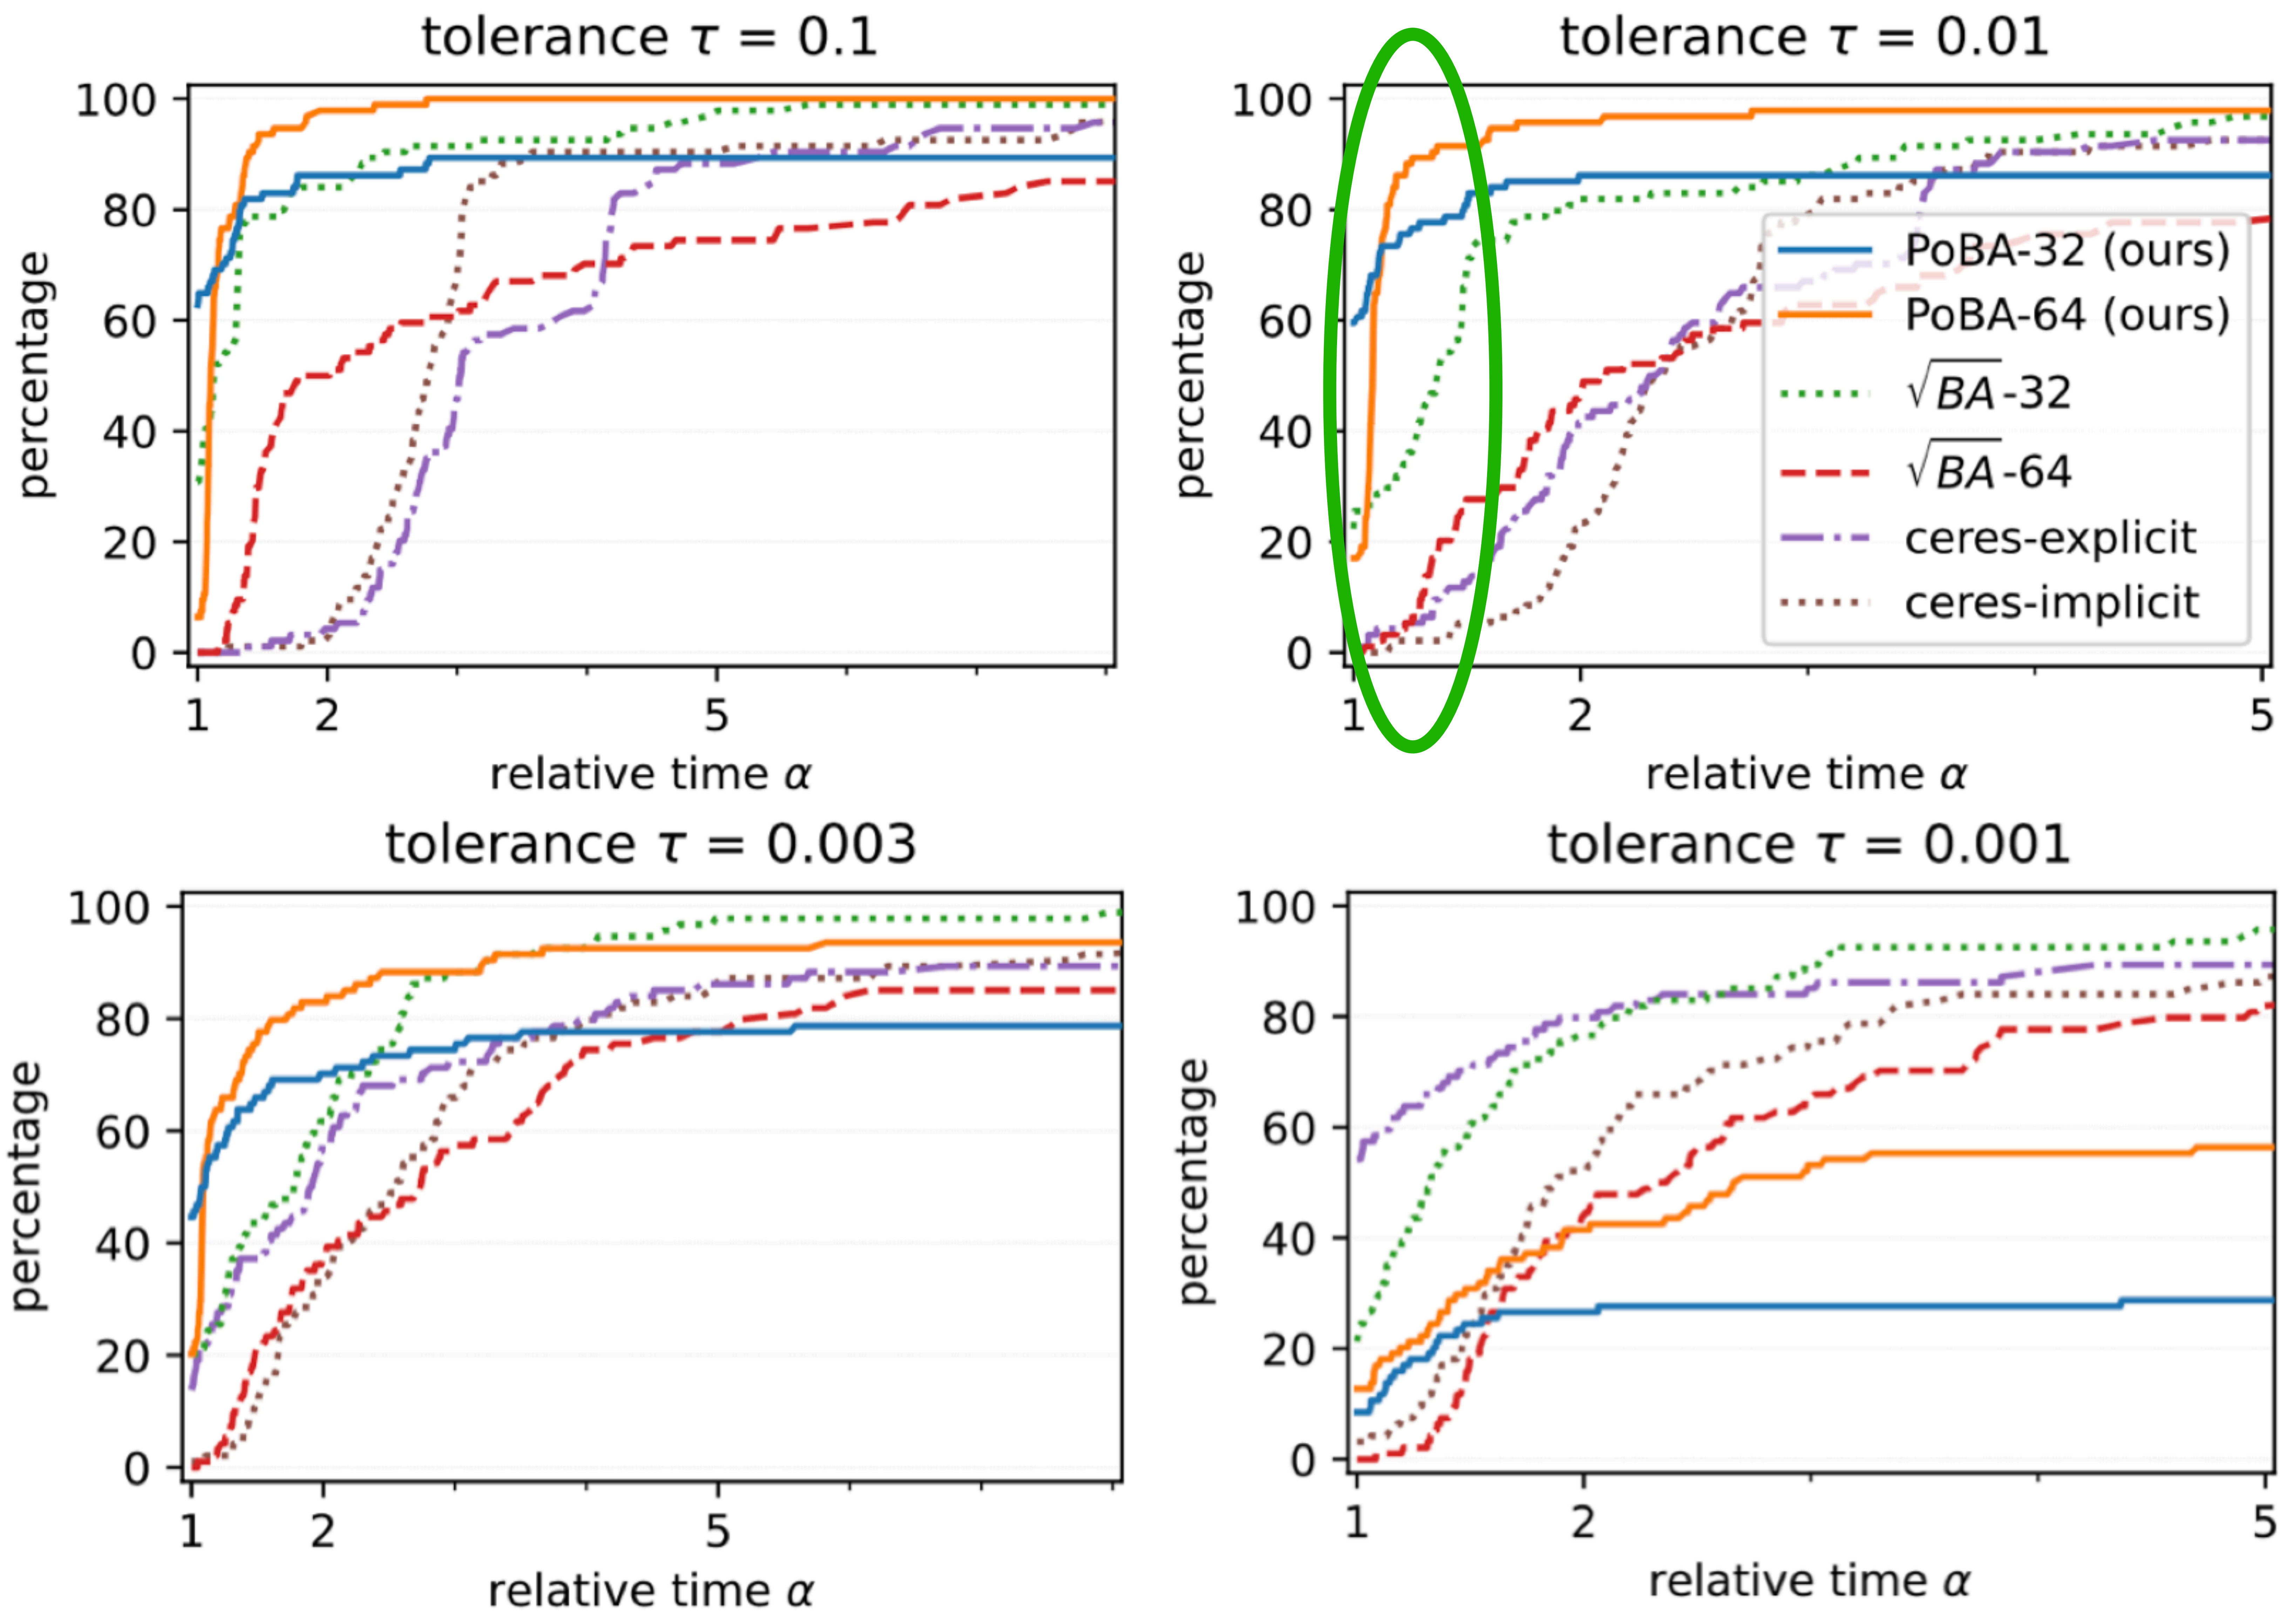

Performance profiles for all 97 BAL problems.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

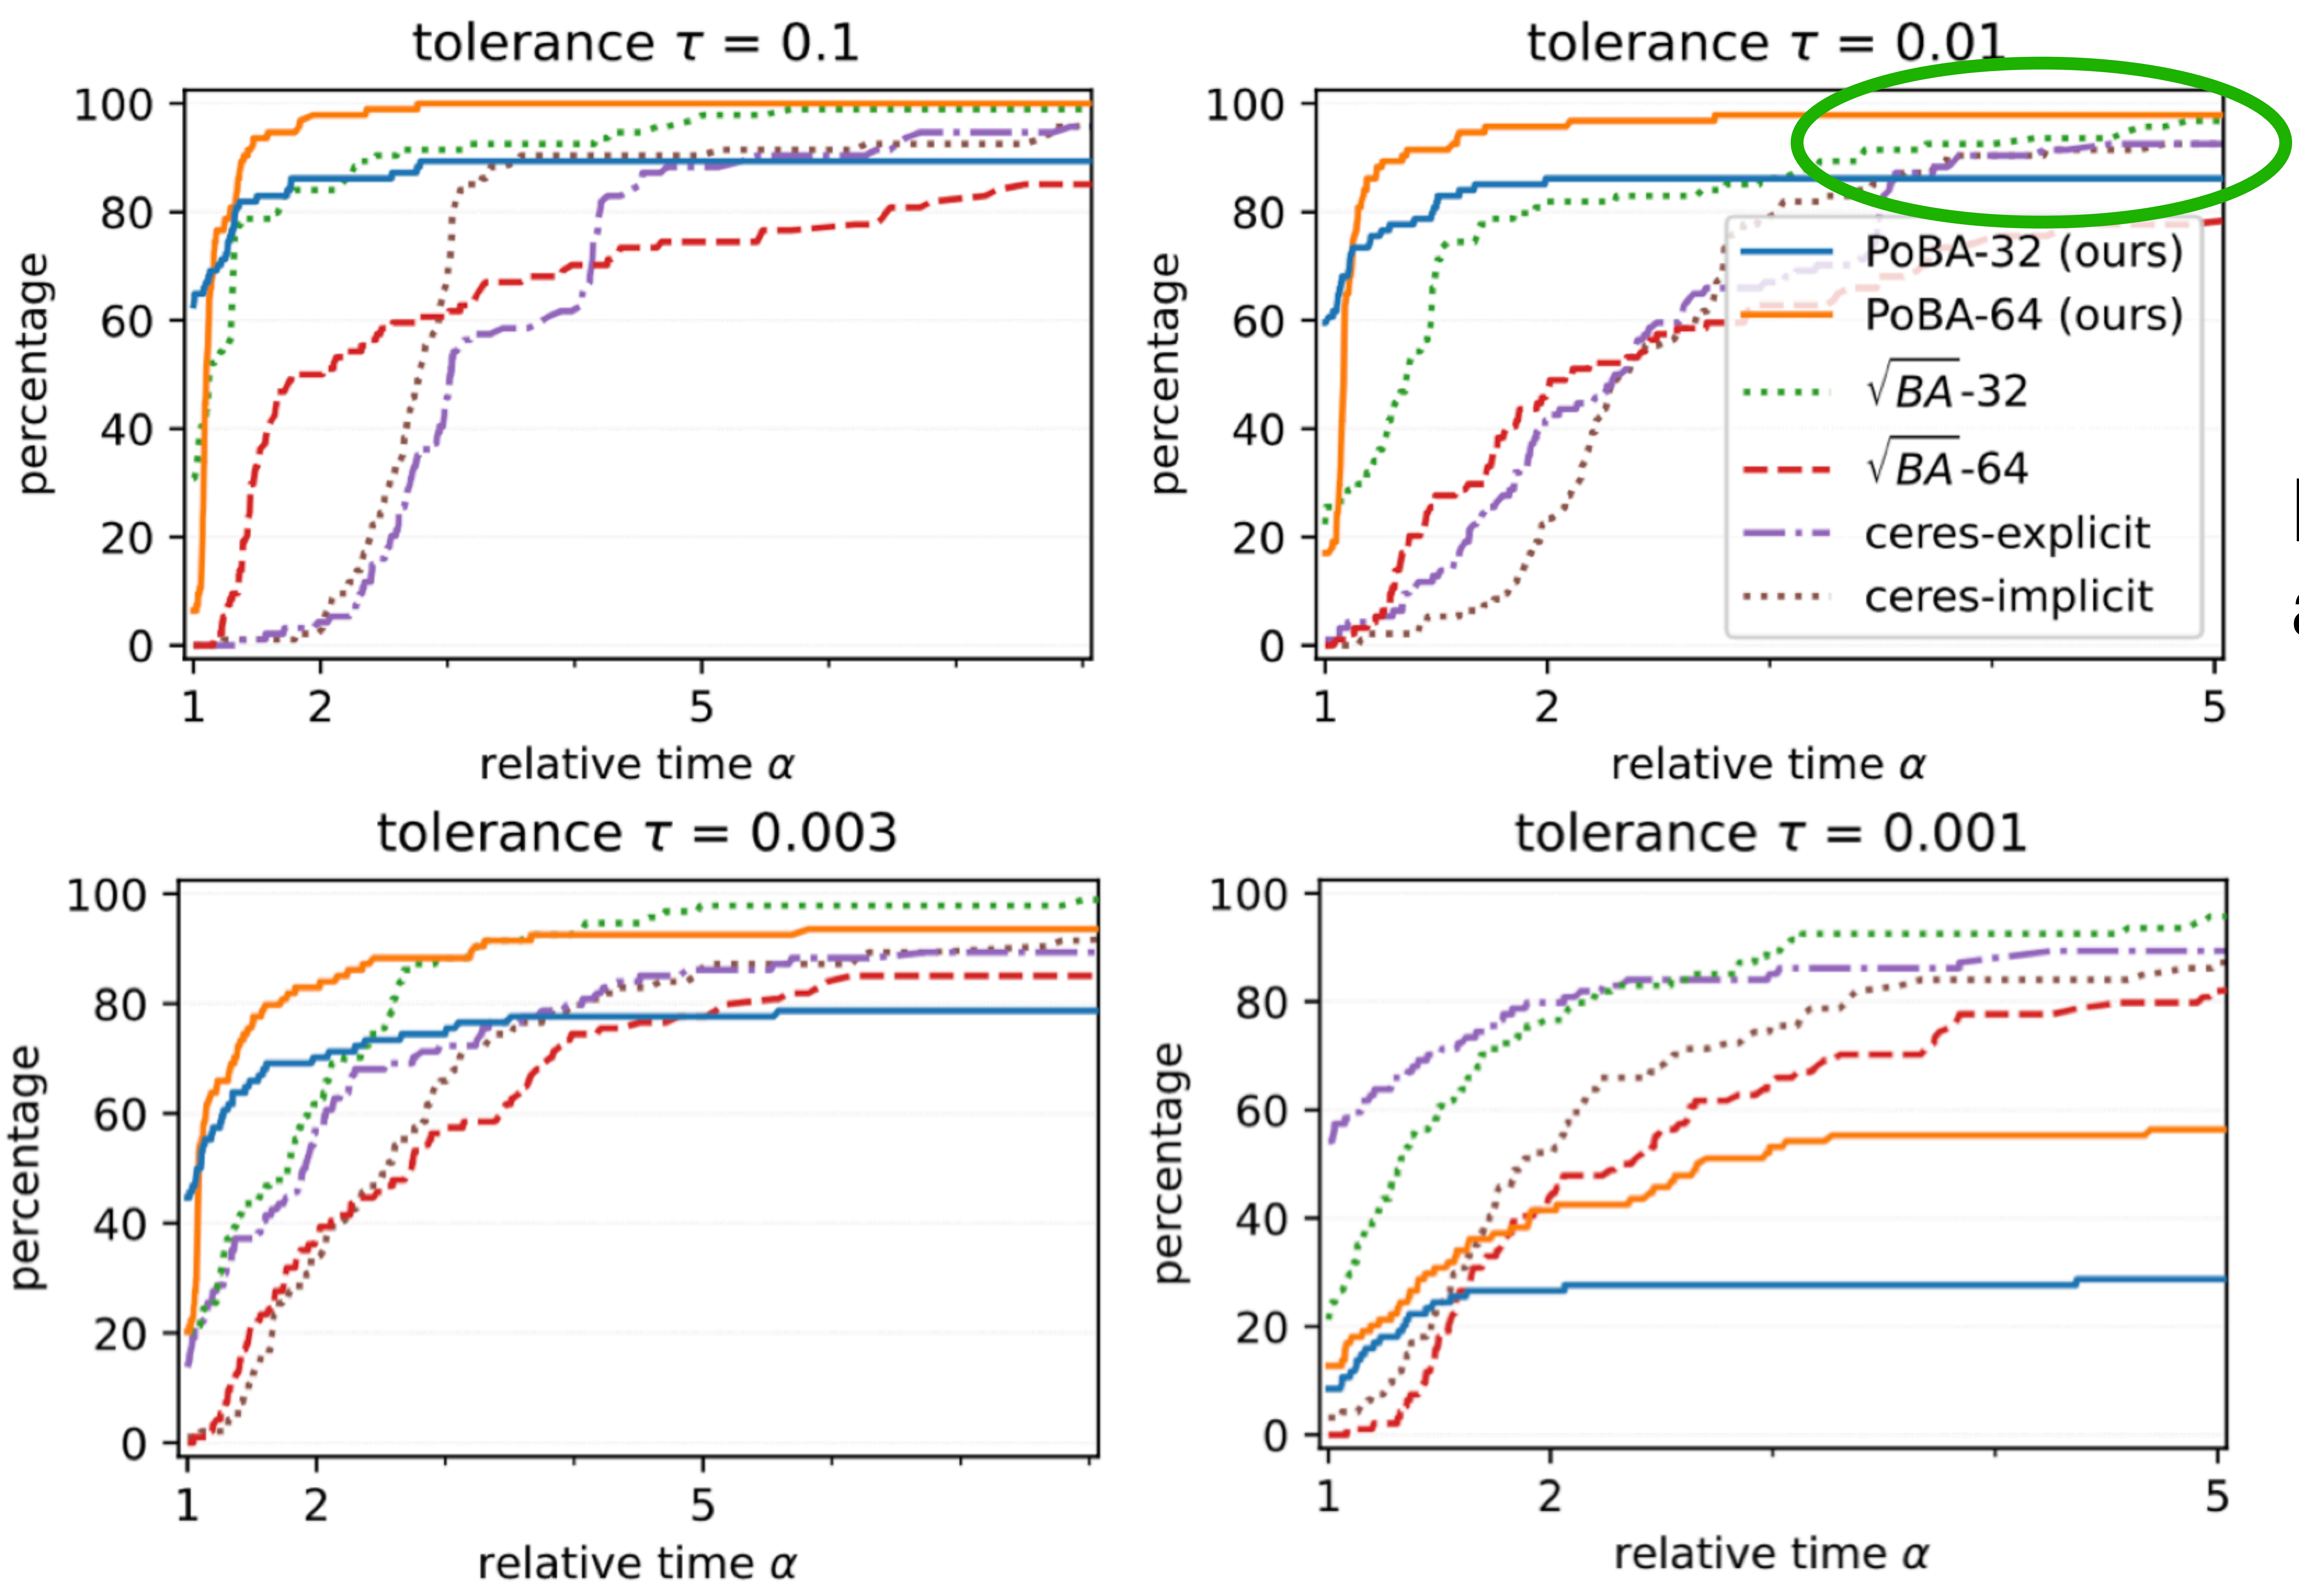

Performance profiles for all 97 BAL problems.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

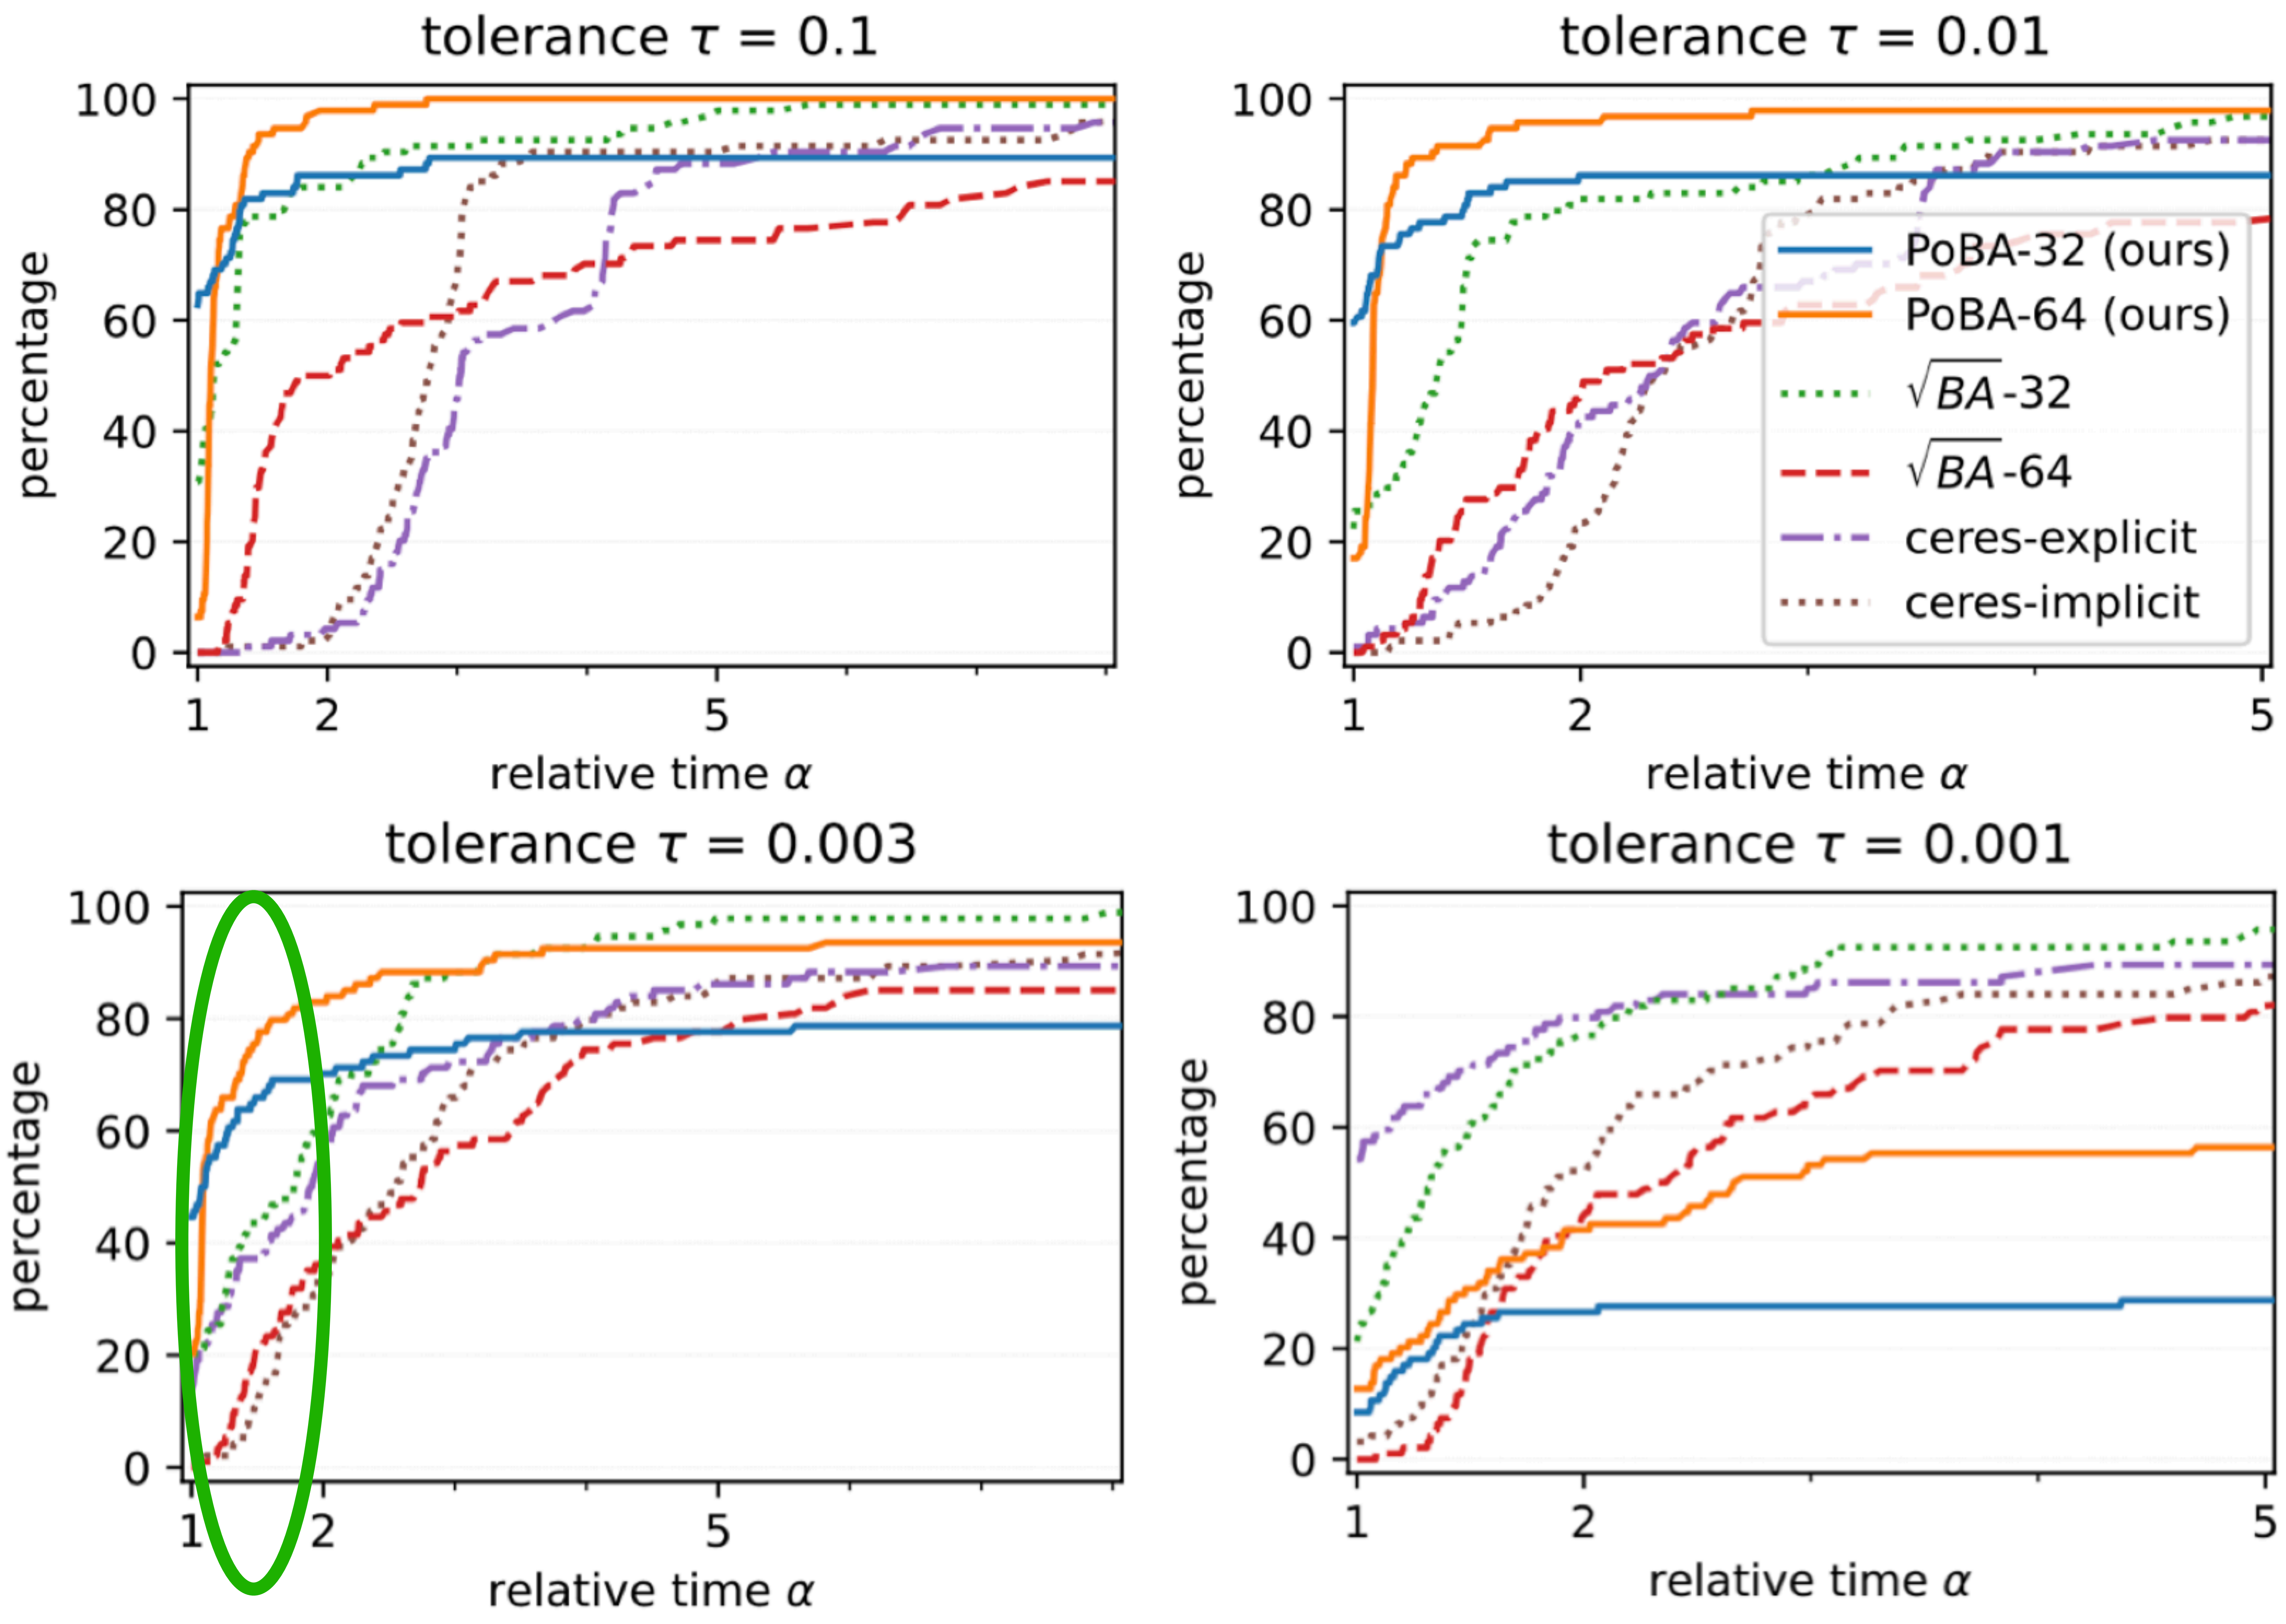

Performance profiles for all 97 BAL problems.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

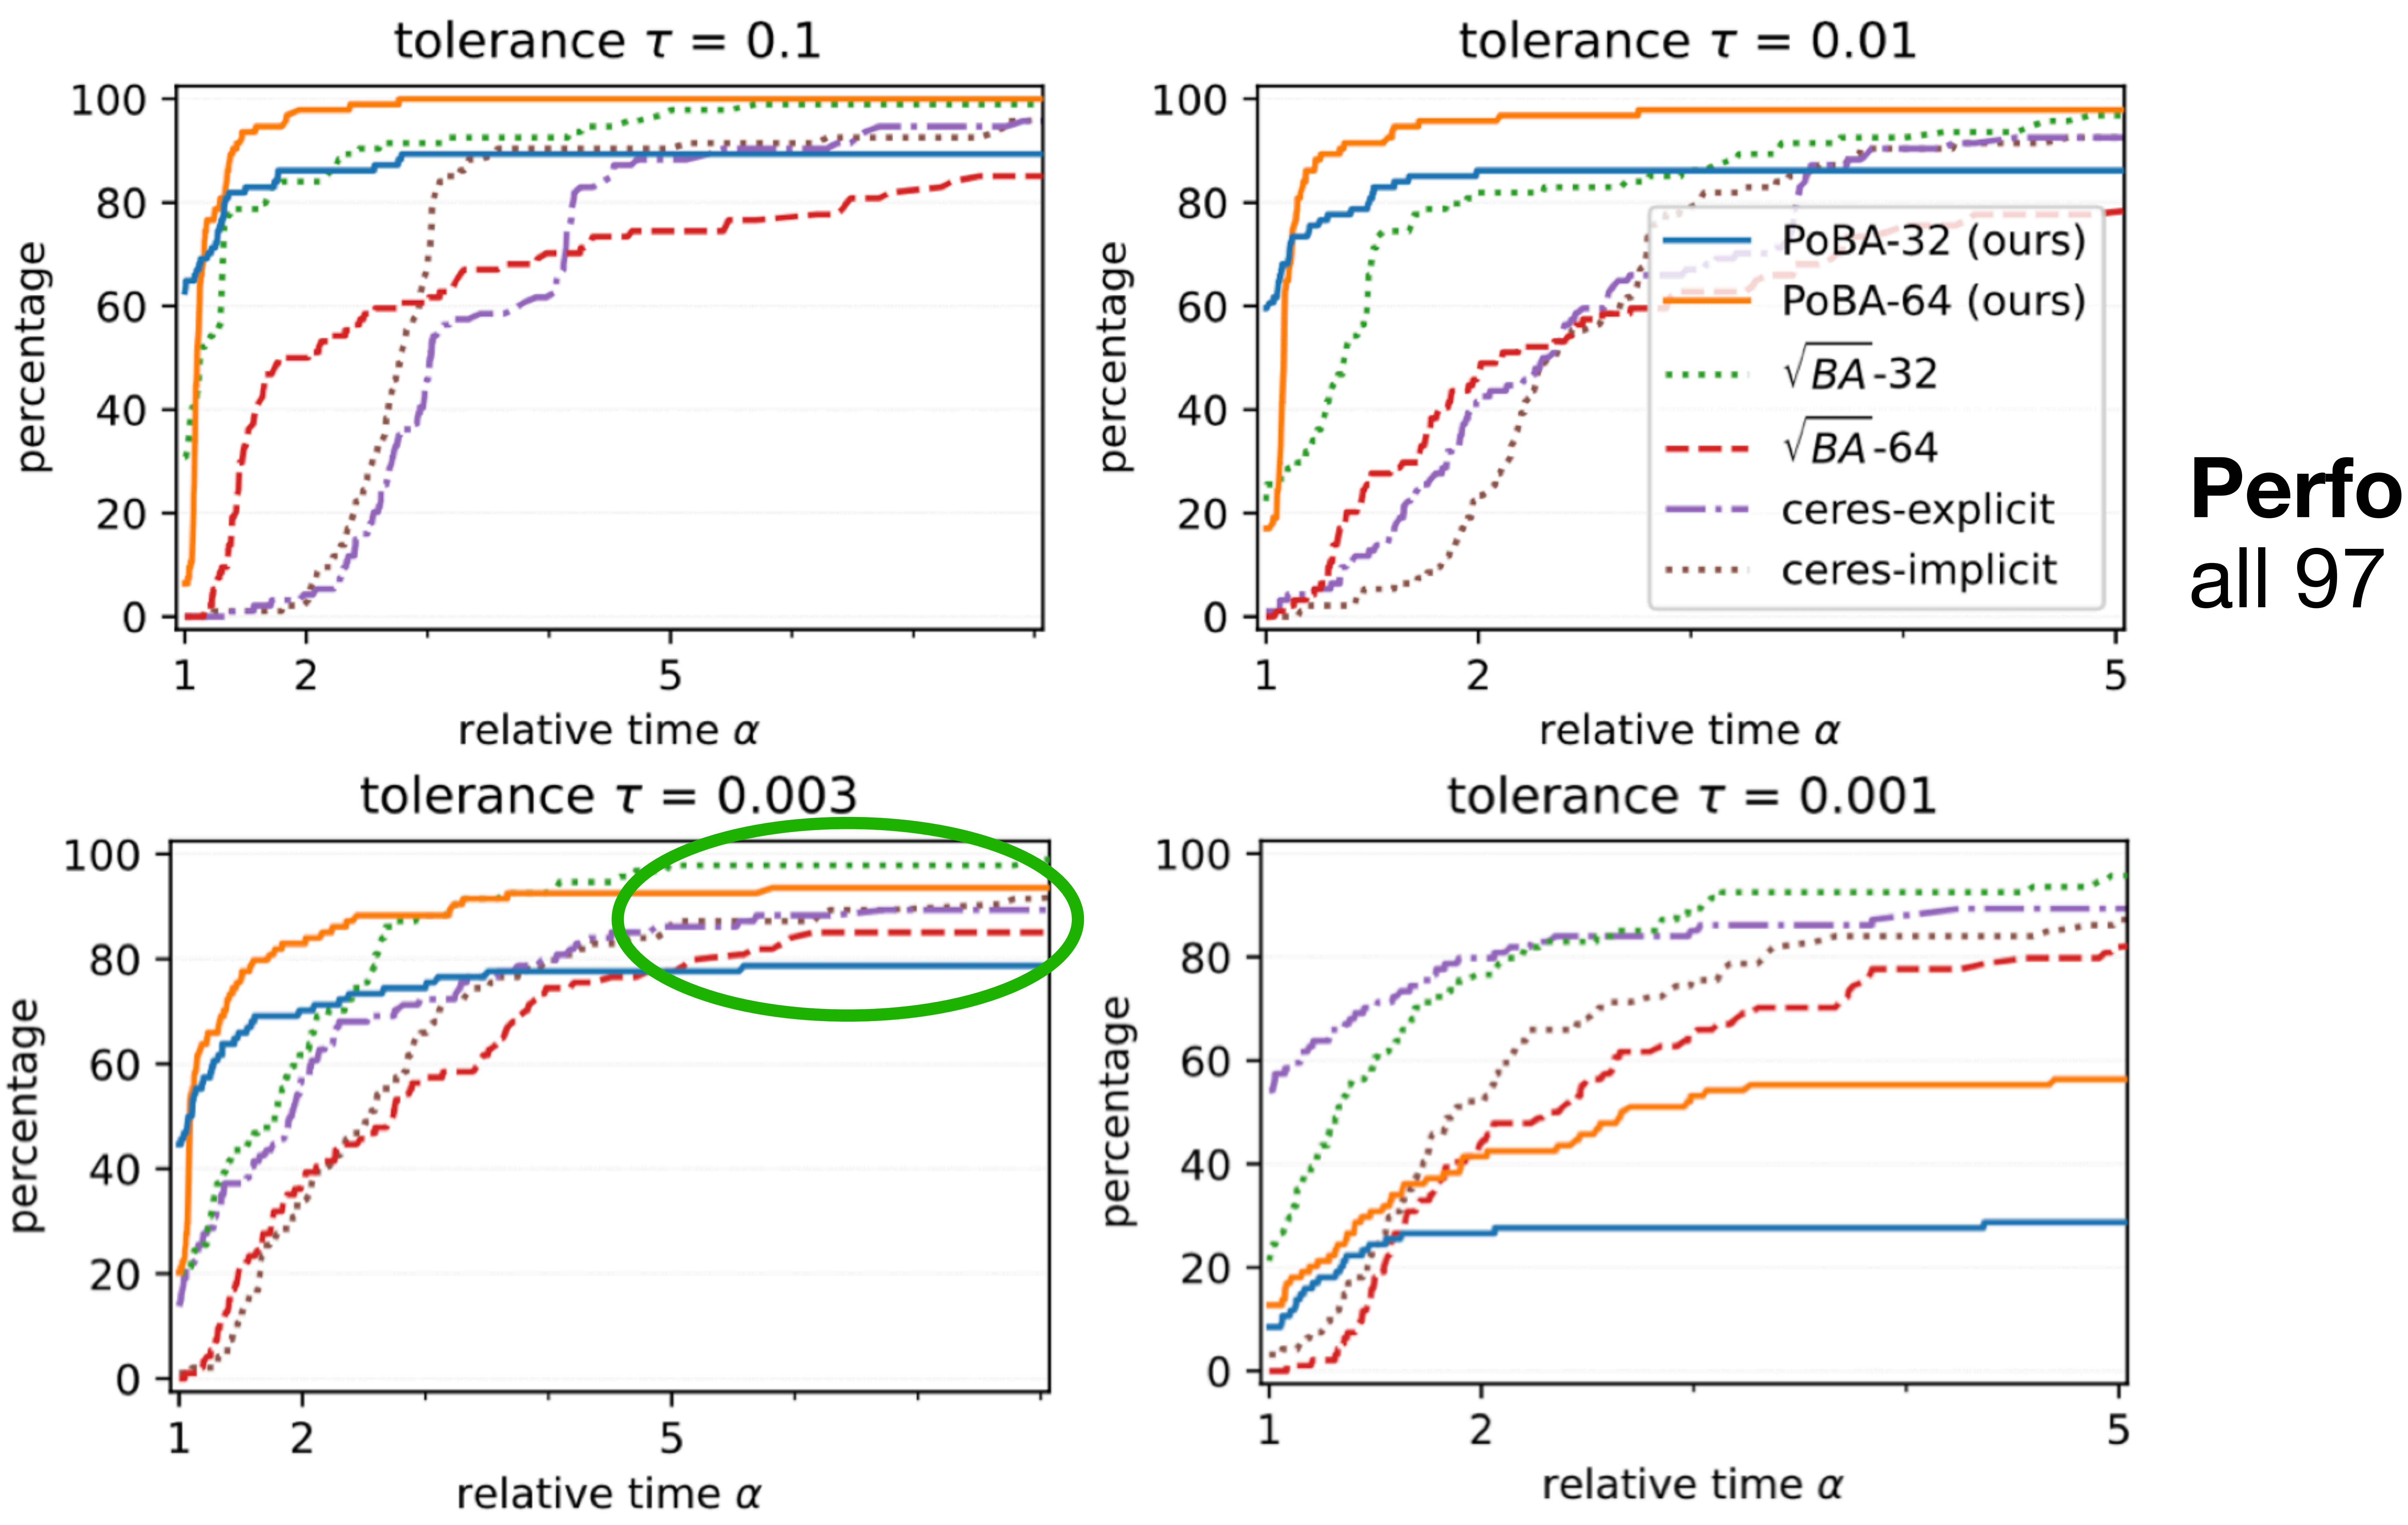

Performance profiles for all 97 BAL problems.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

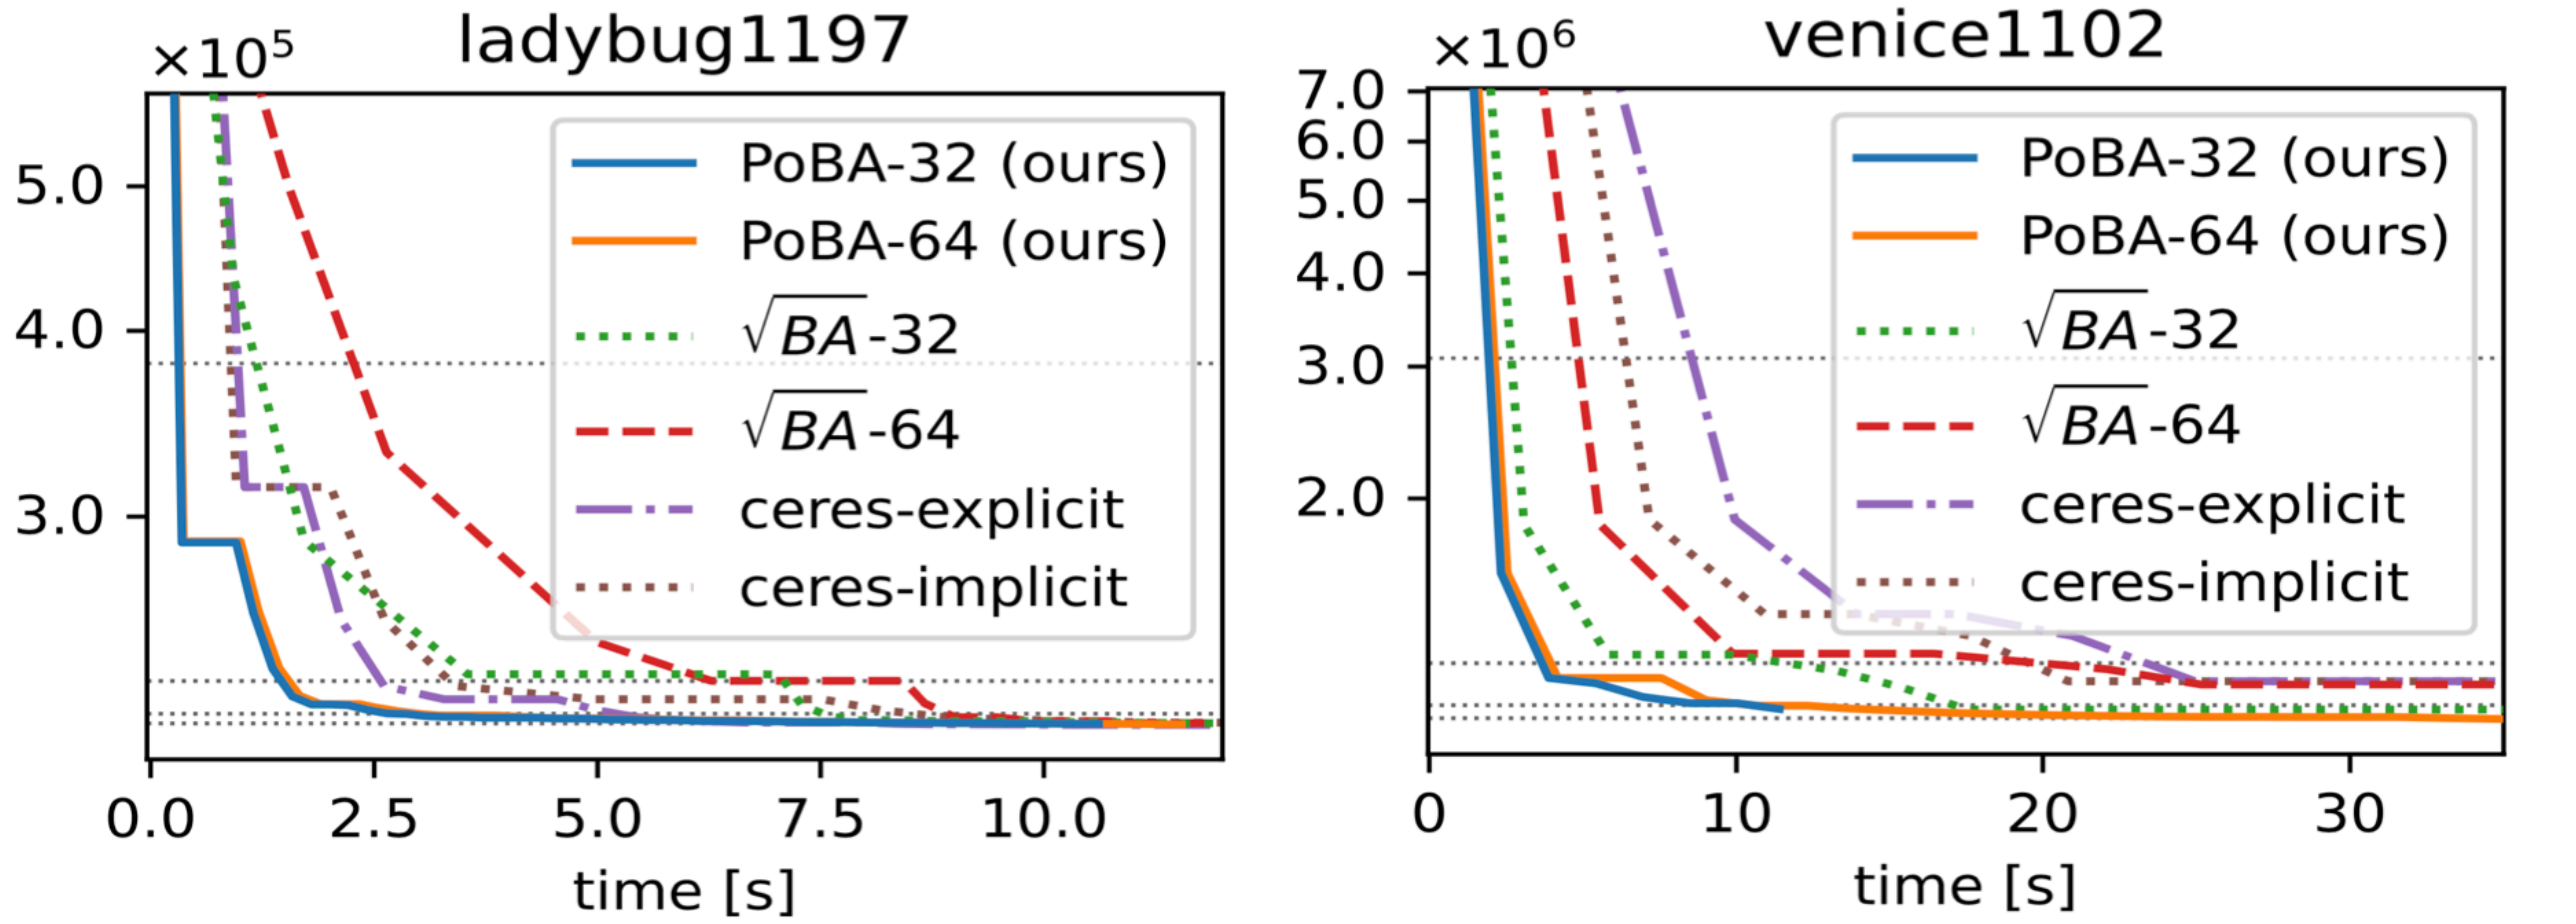

Convergence plots for BAL problems with 1197 poses (left) and 1102 poses (right).

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

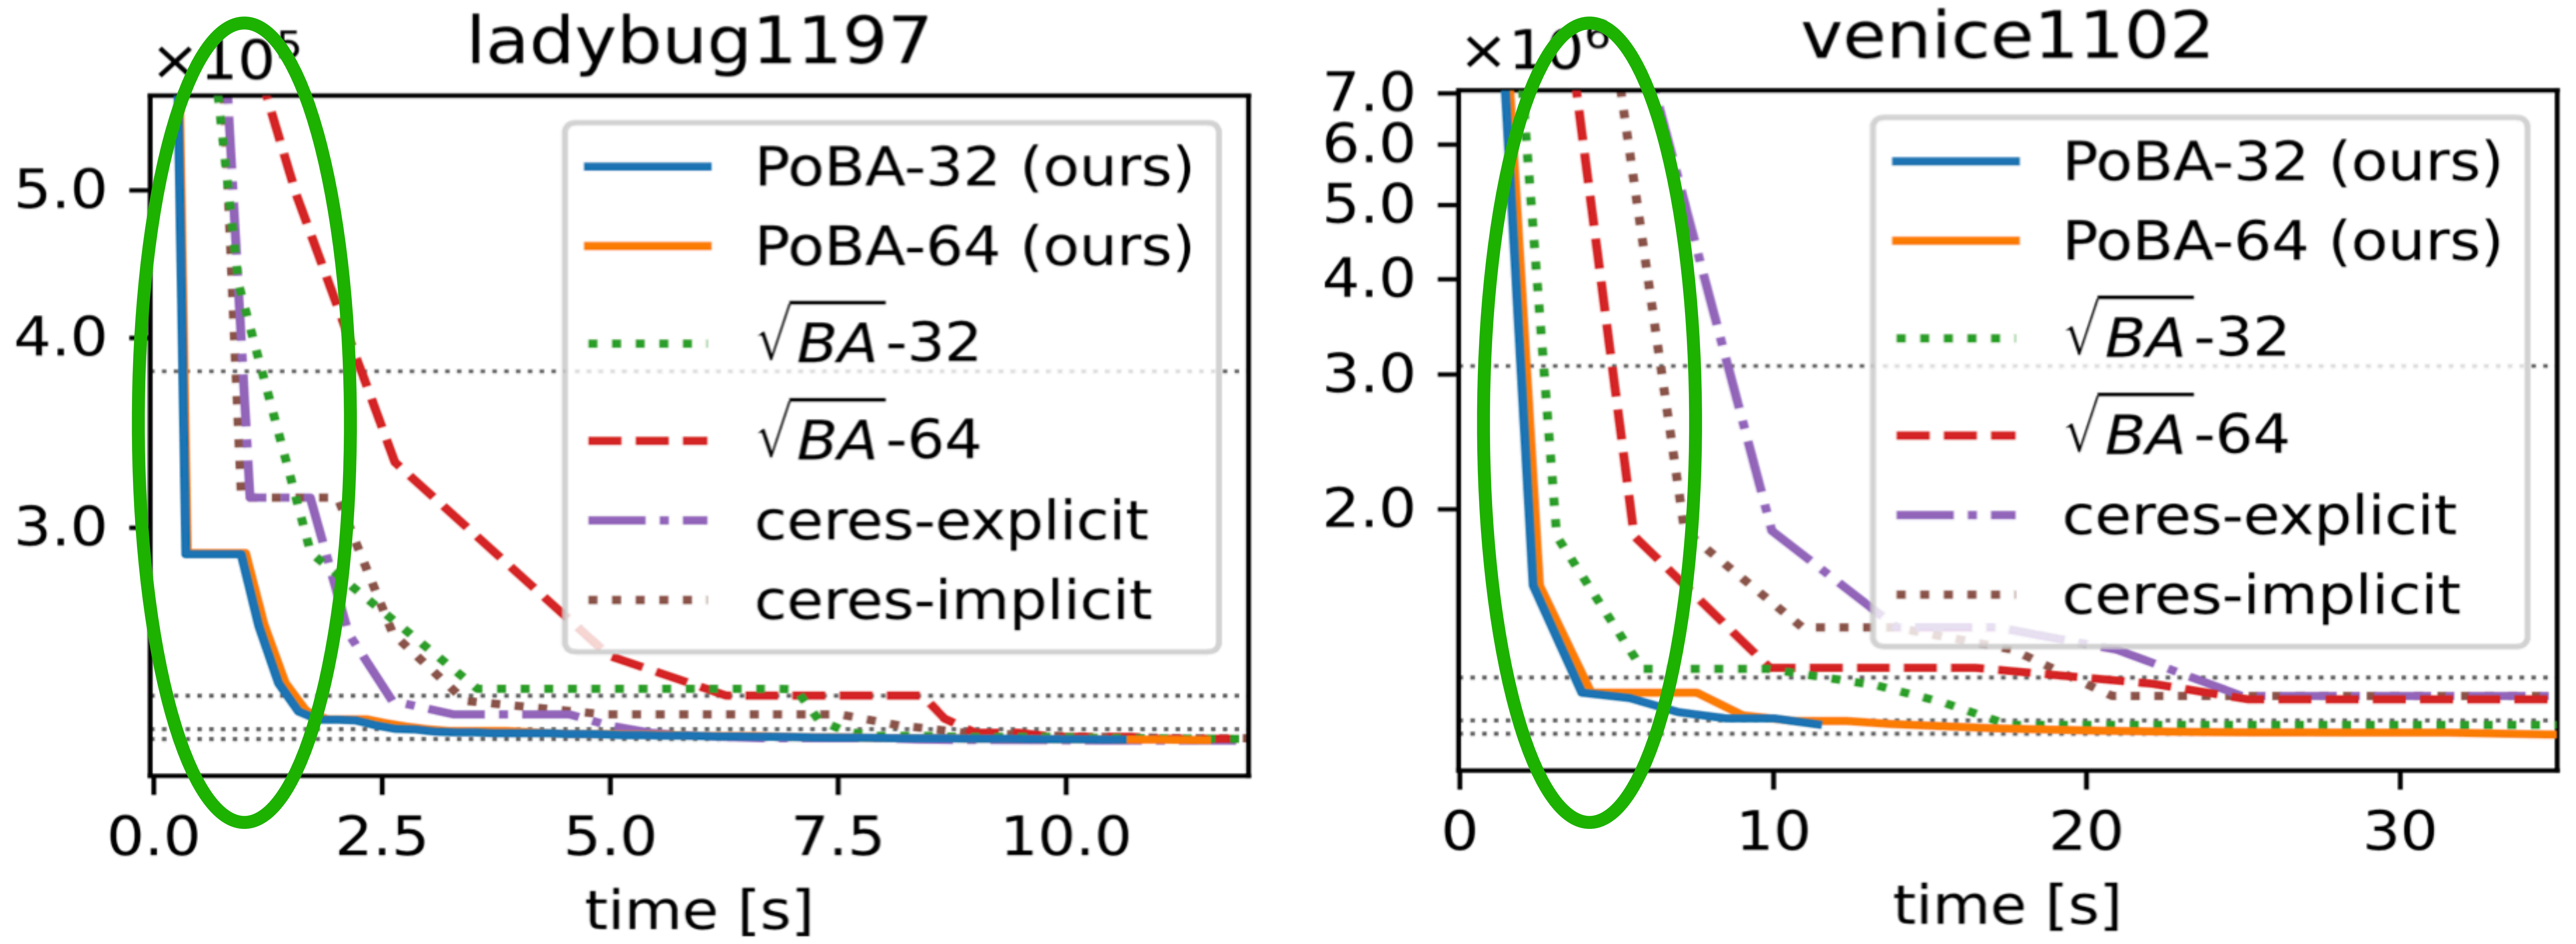

Convergence plots for BAL problems with 1197 poses (left) and 1102 poses (right).

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

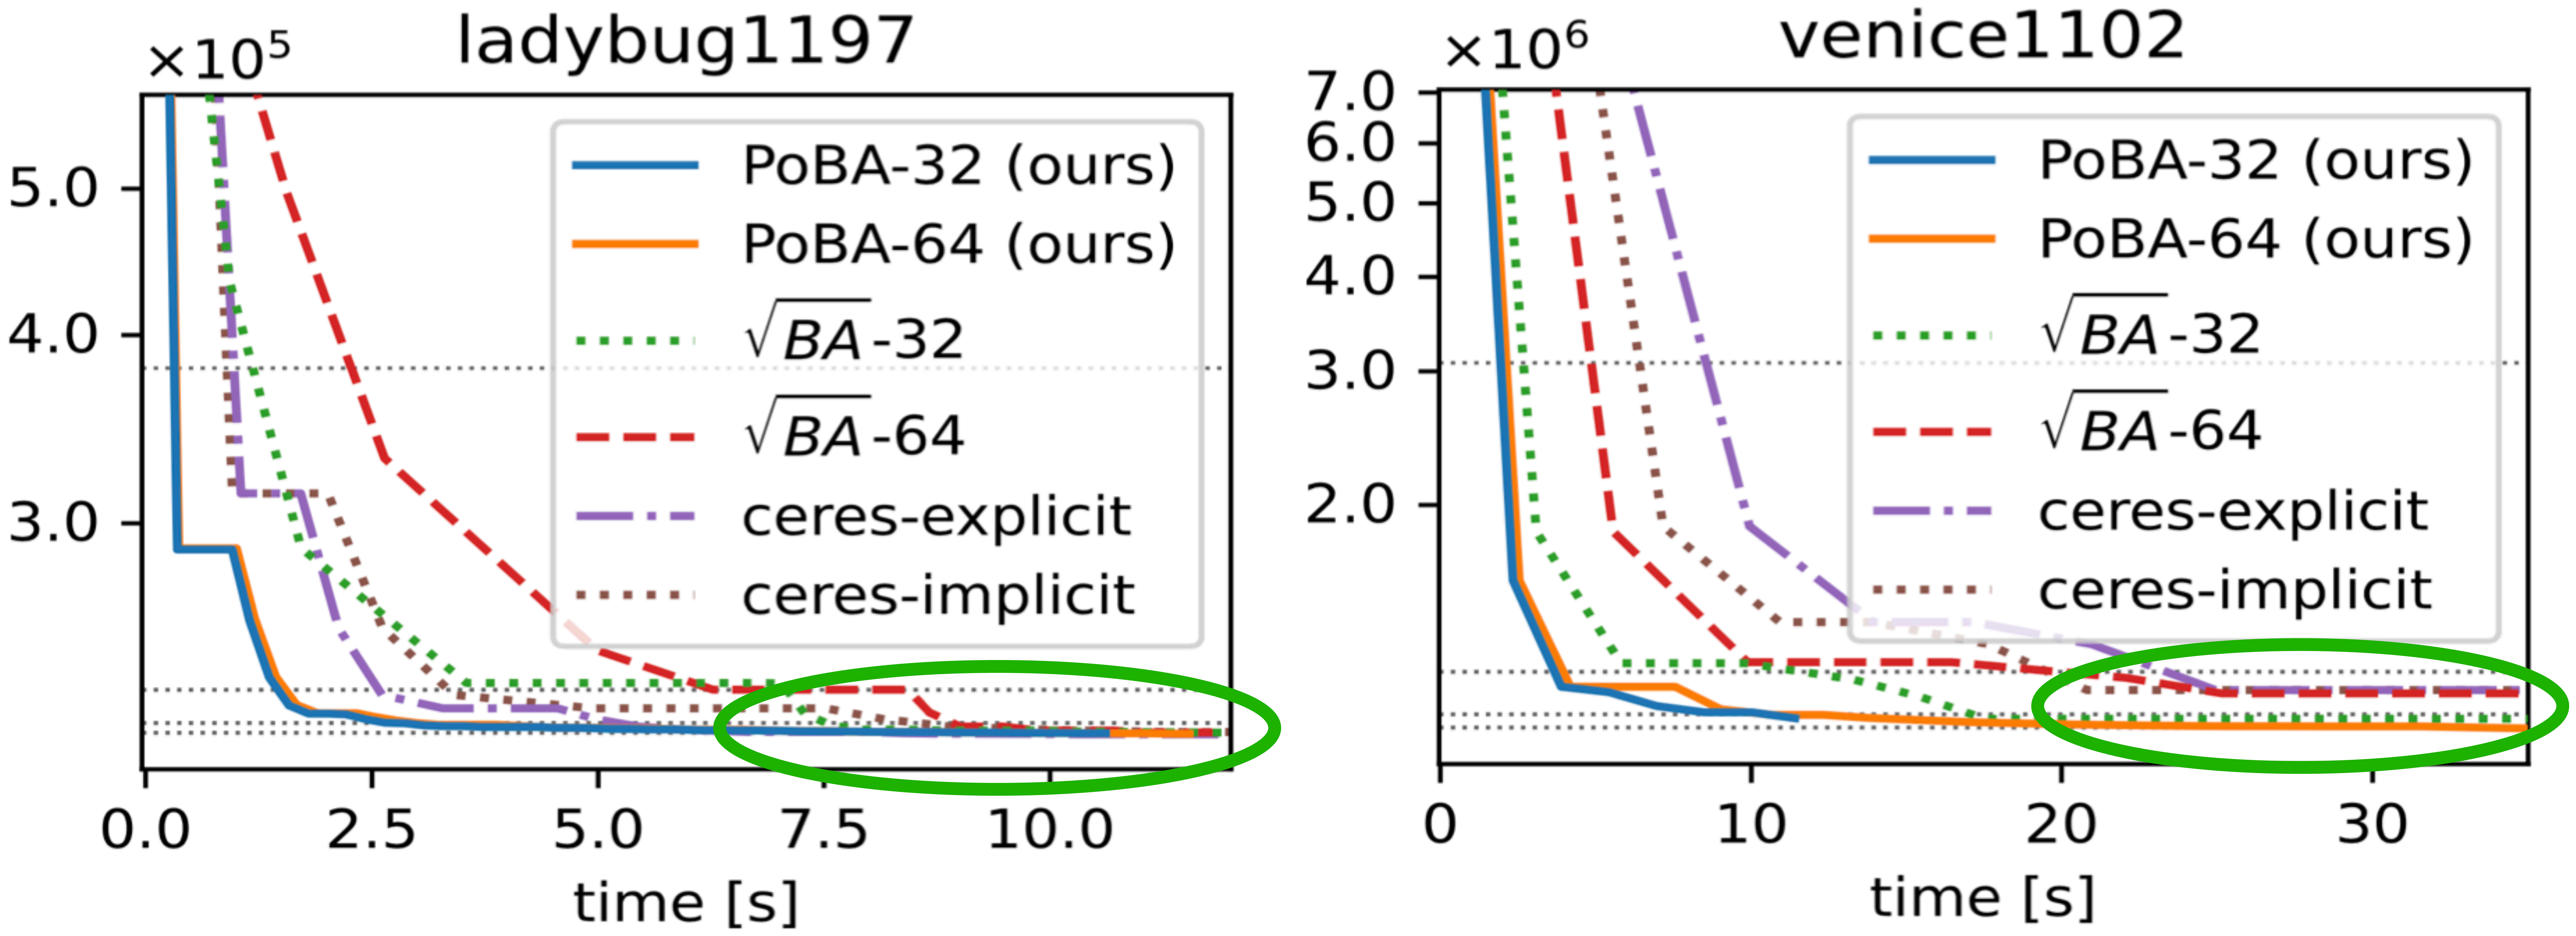

Convergence plots for BAL problems with 1197 poses (left) and 1102 poses (right).

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

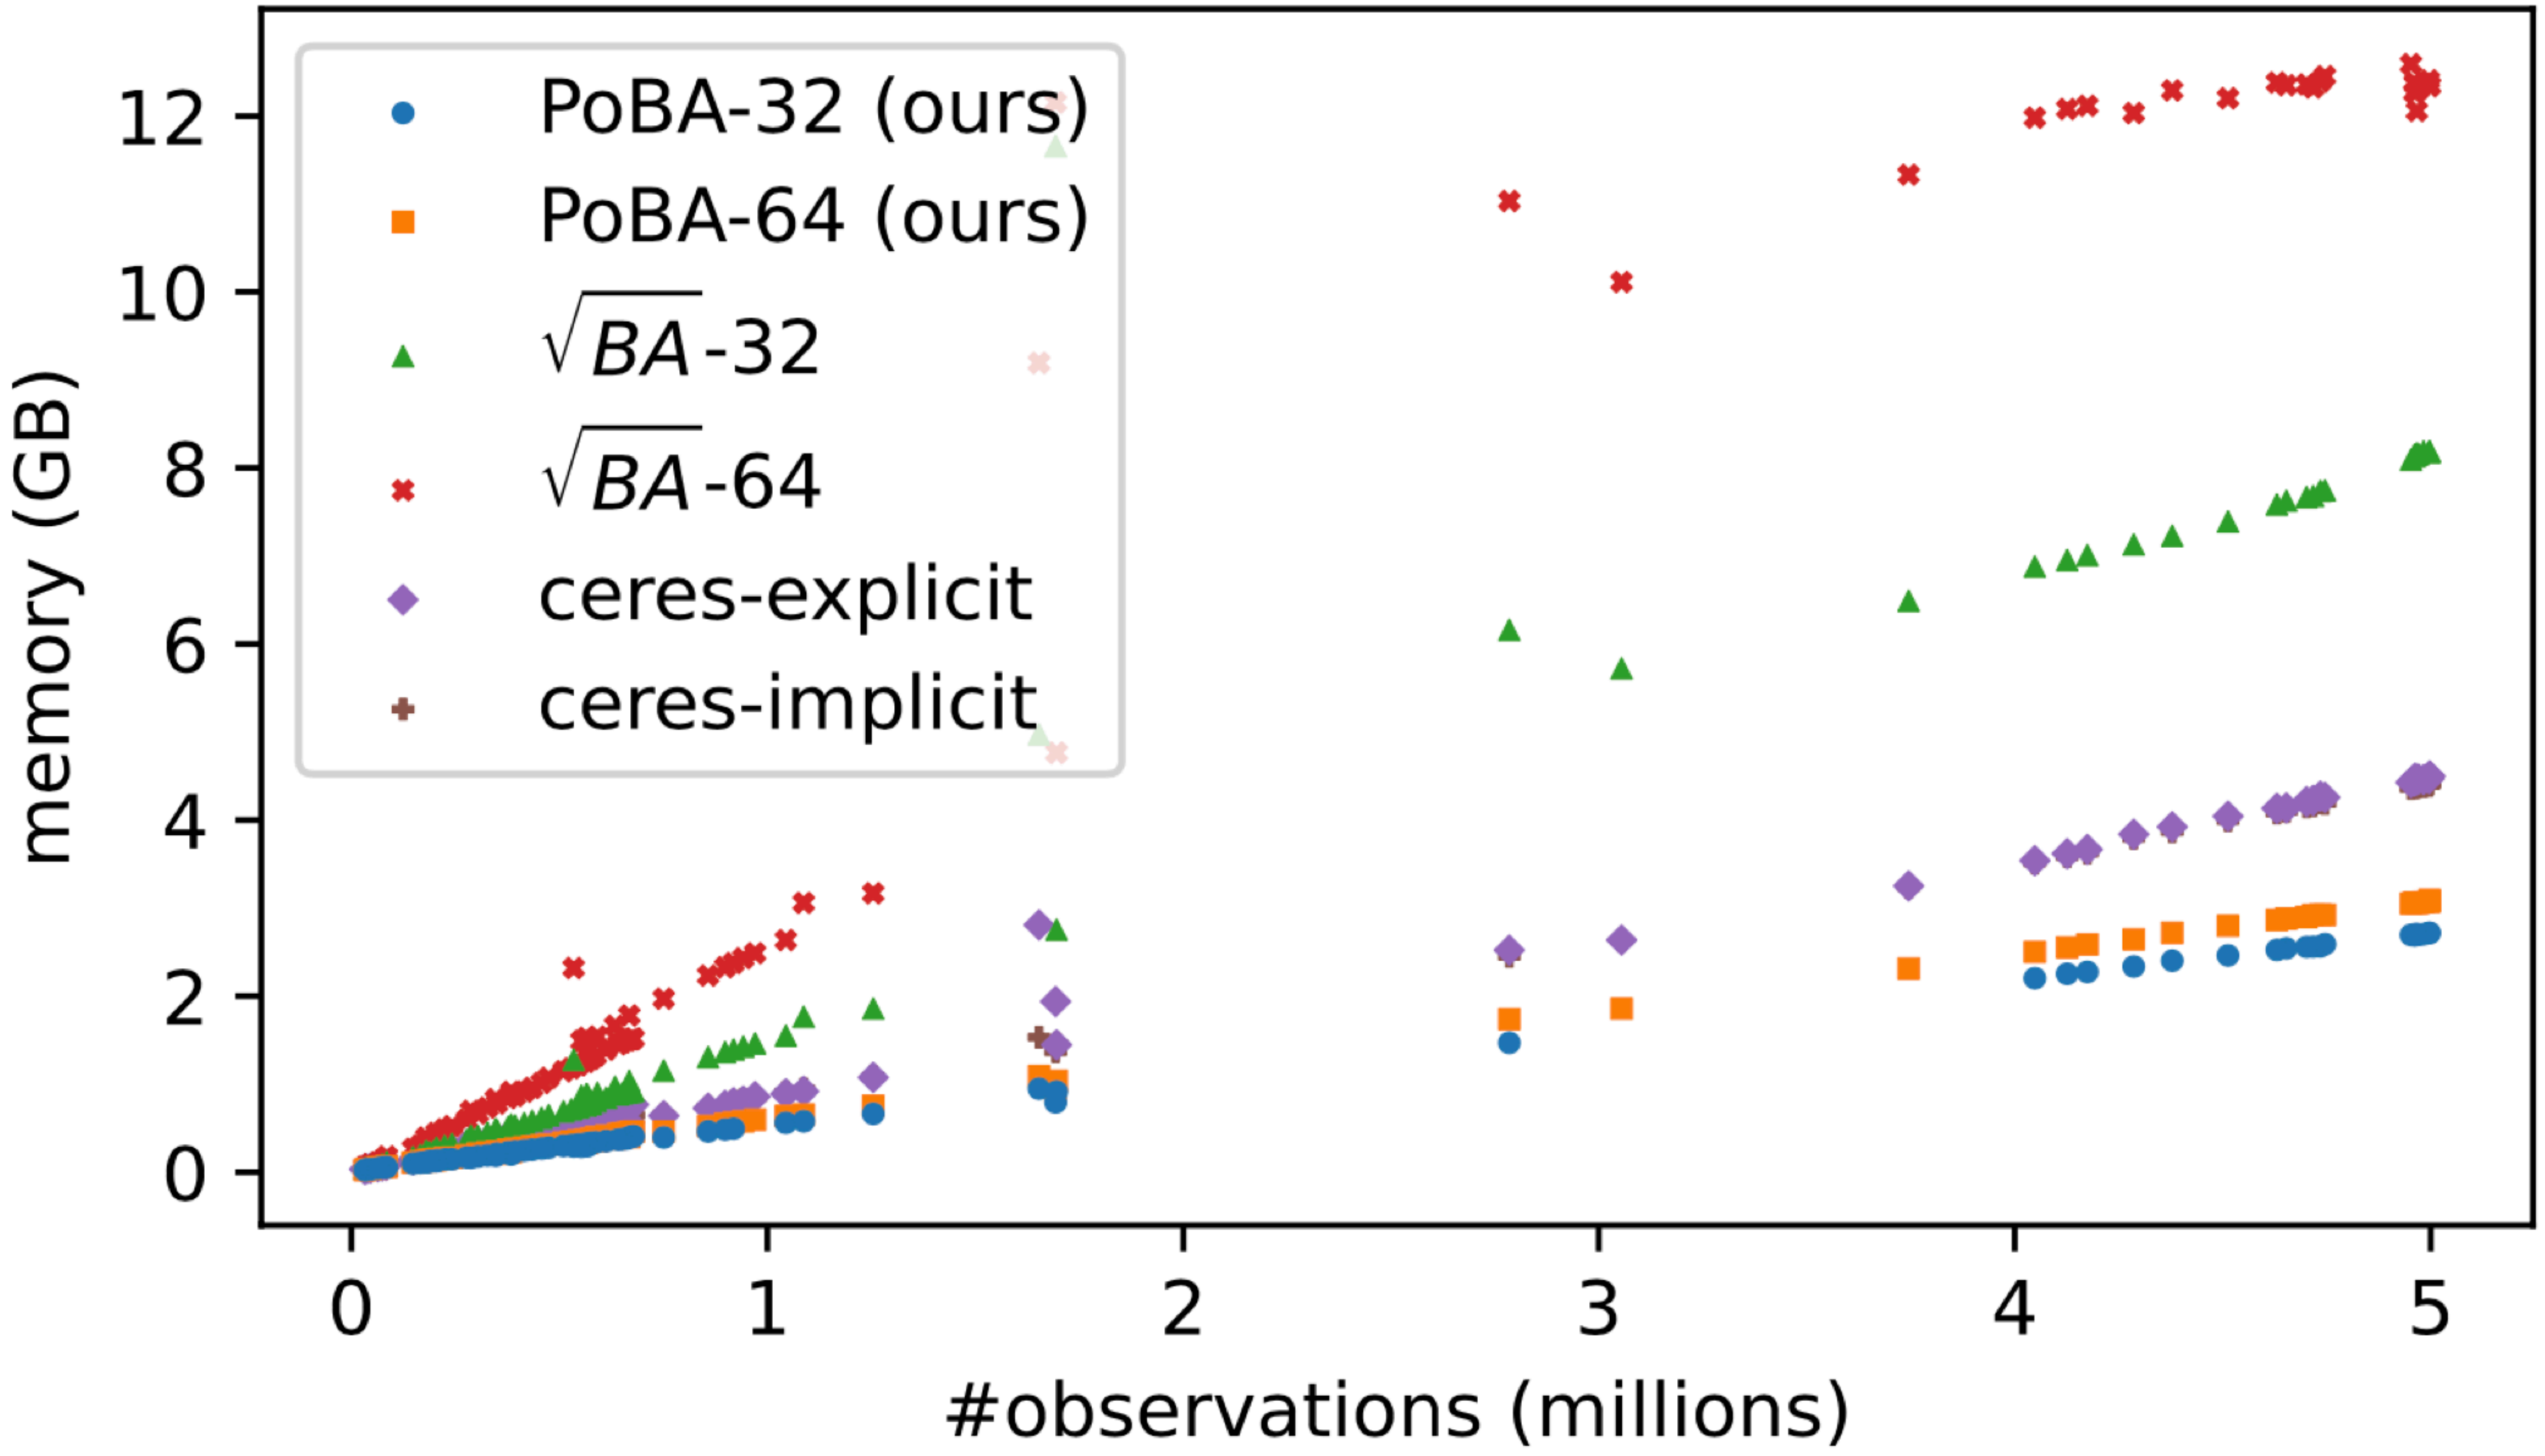

Memory consumption for all BAL problems.

- PoBA requires almost five times less memory than $\sqrt{BA}$ and twice less memory than Ceres.
- By far the **less memory-consuming** solver wrt its challengers.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

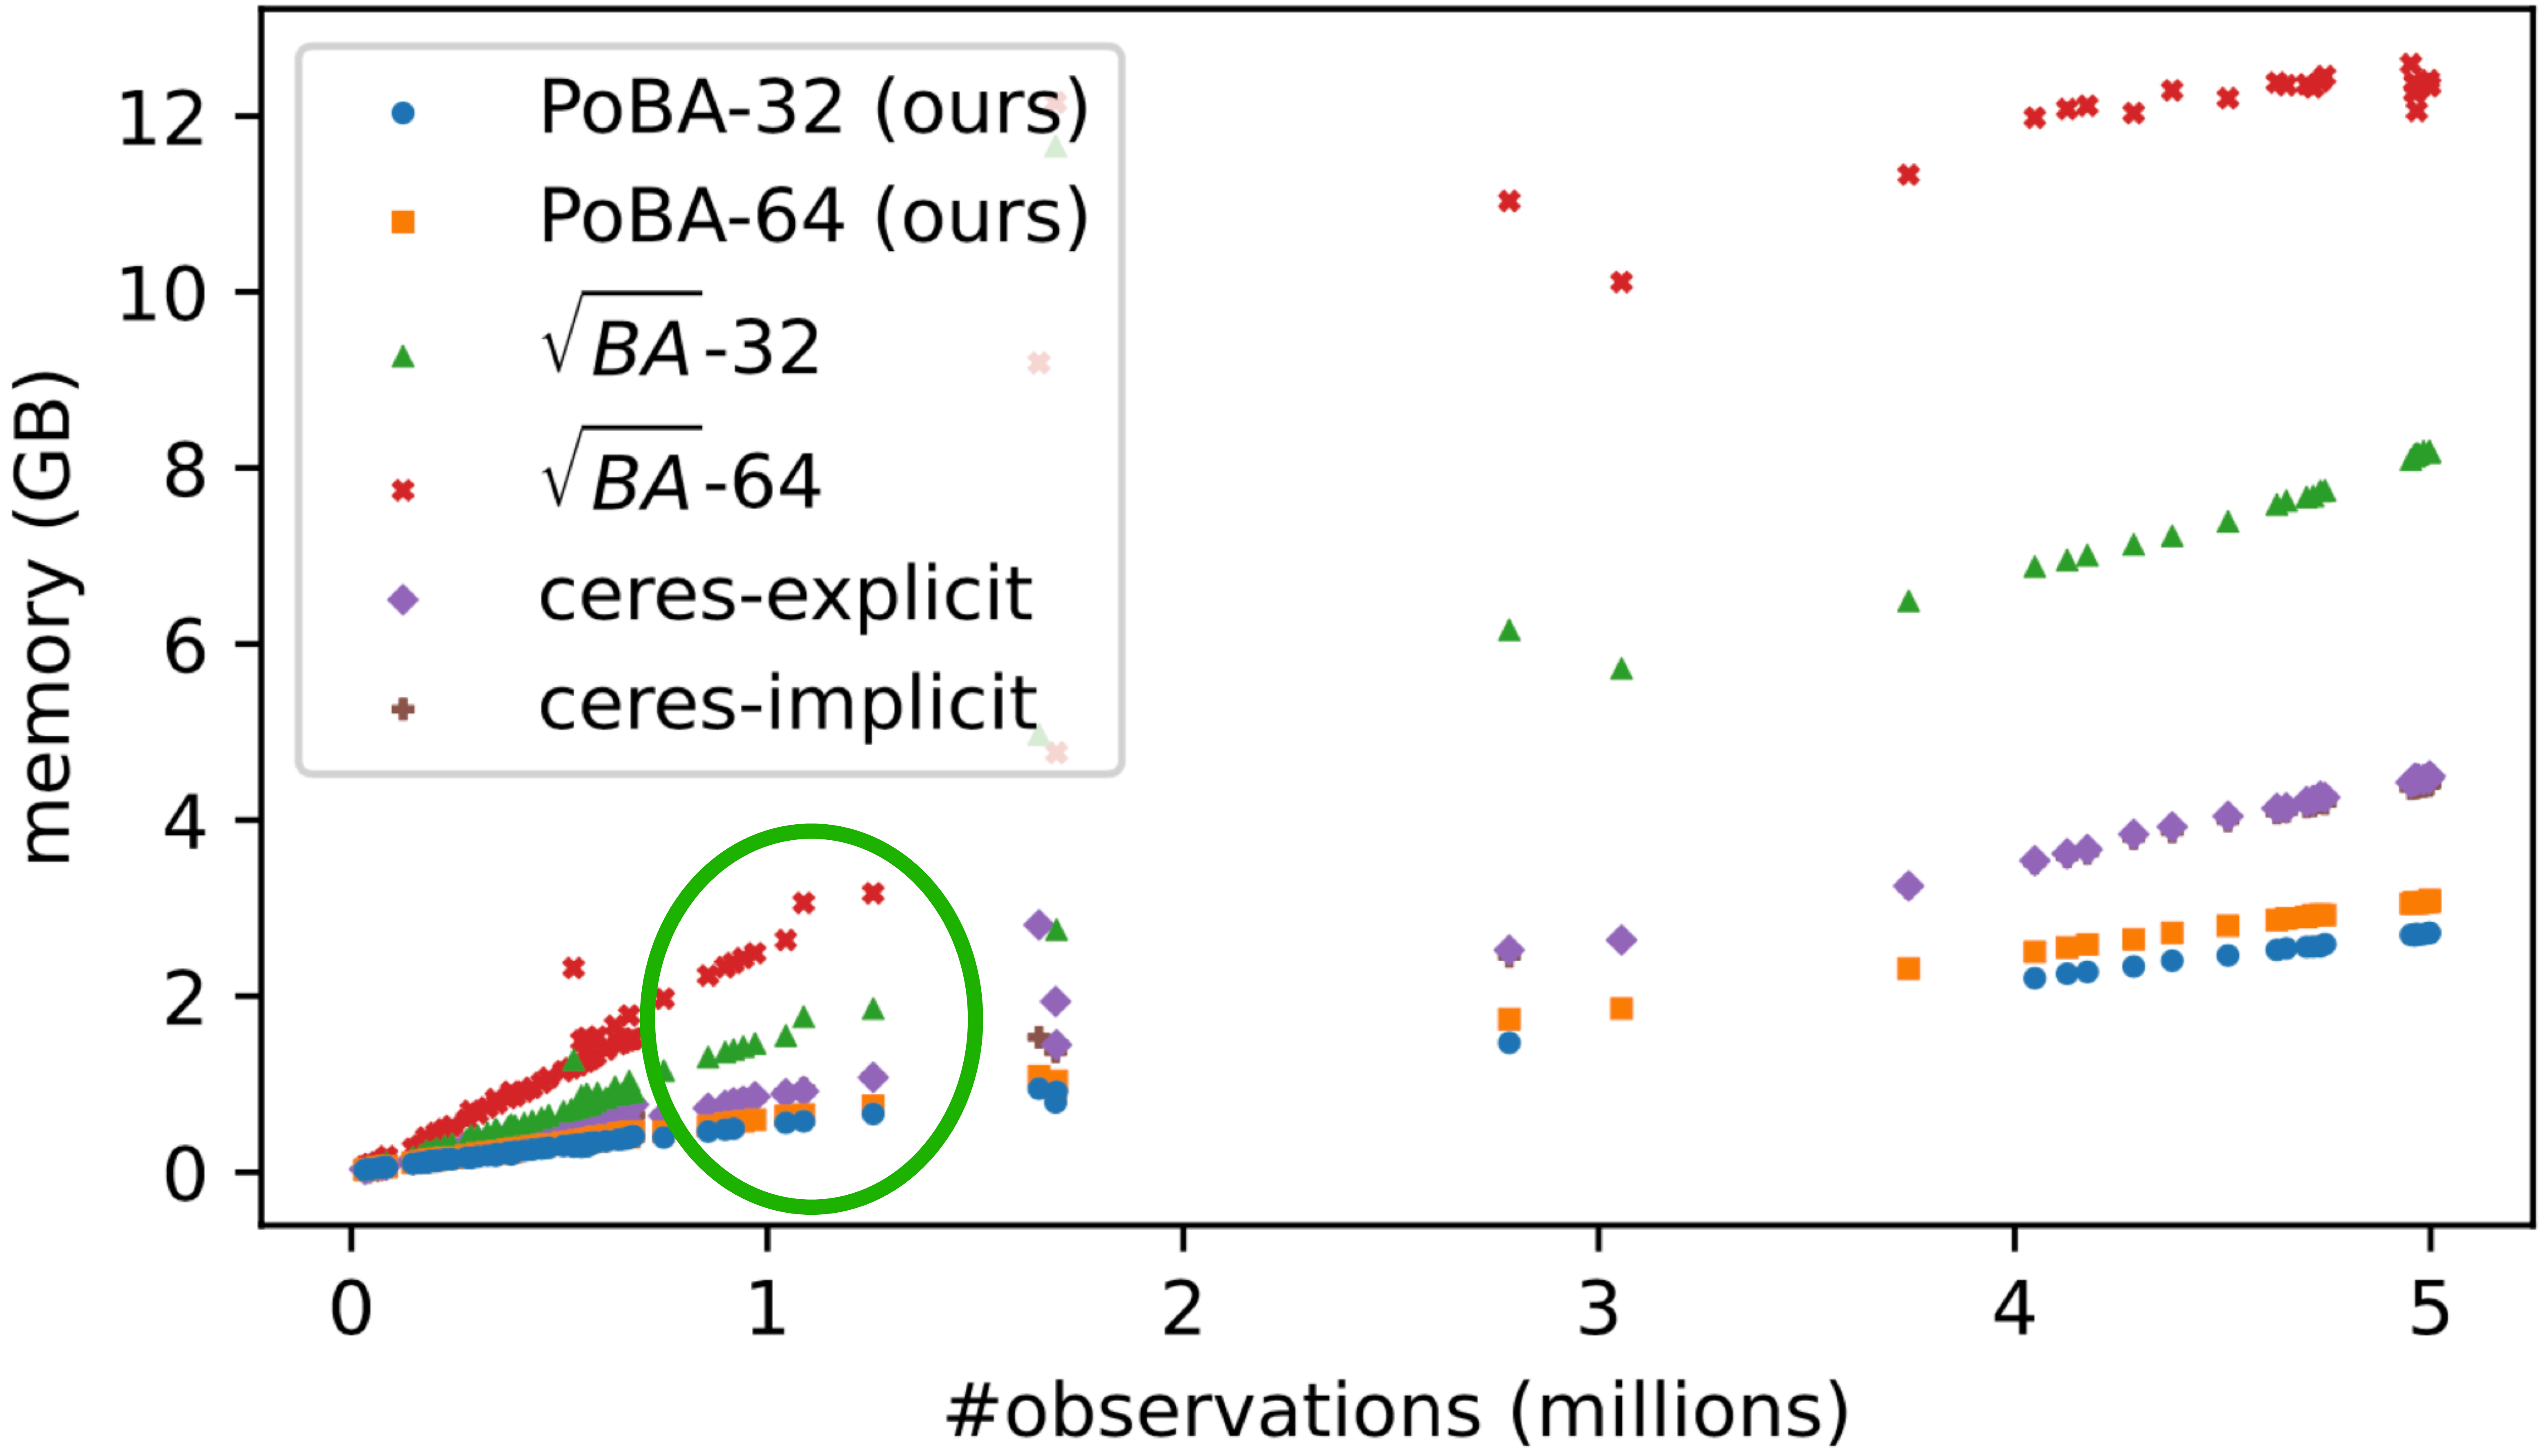

Memory consumption for all BAL problems.

- PoBA requires almost five times less memory than $\sqrt{BA}$ and twice less memory than Ceres.
- By far the **less memory-consuming** solver wrt its challengers.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

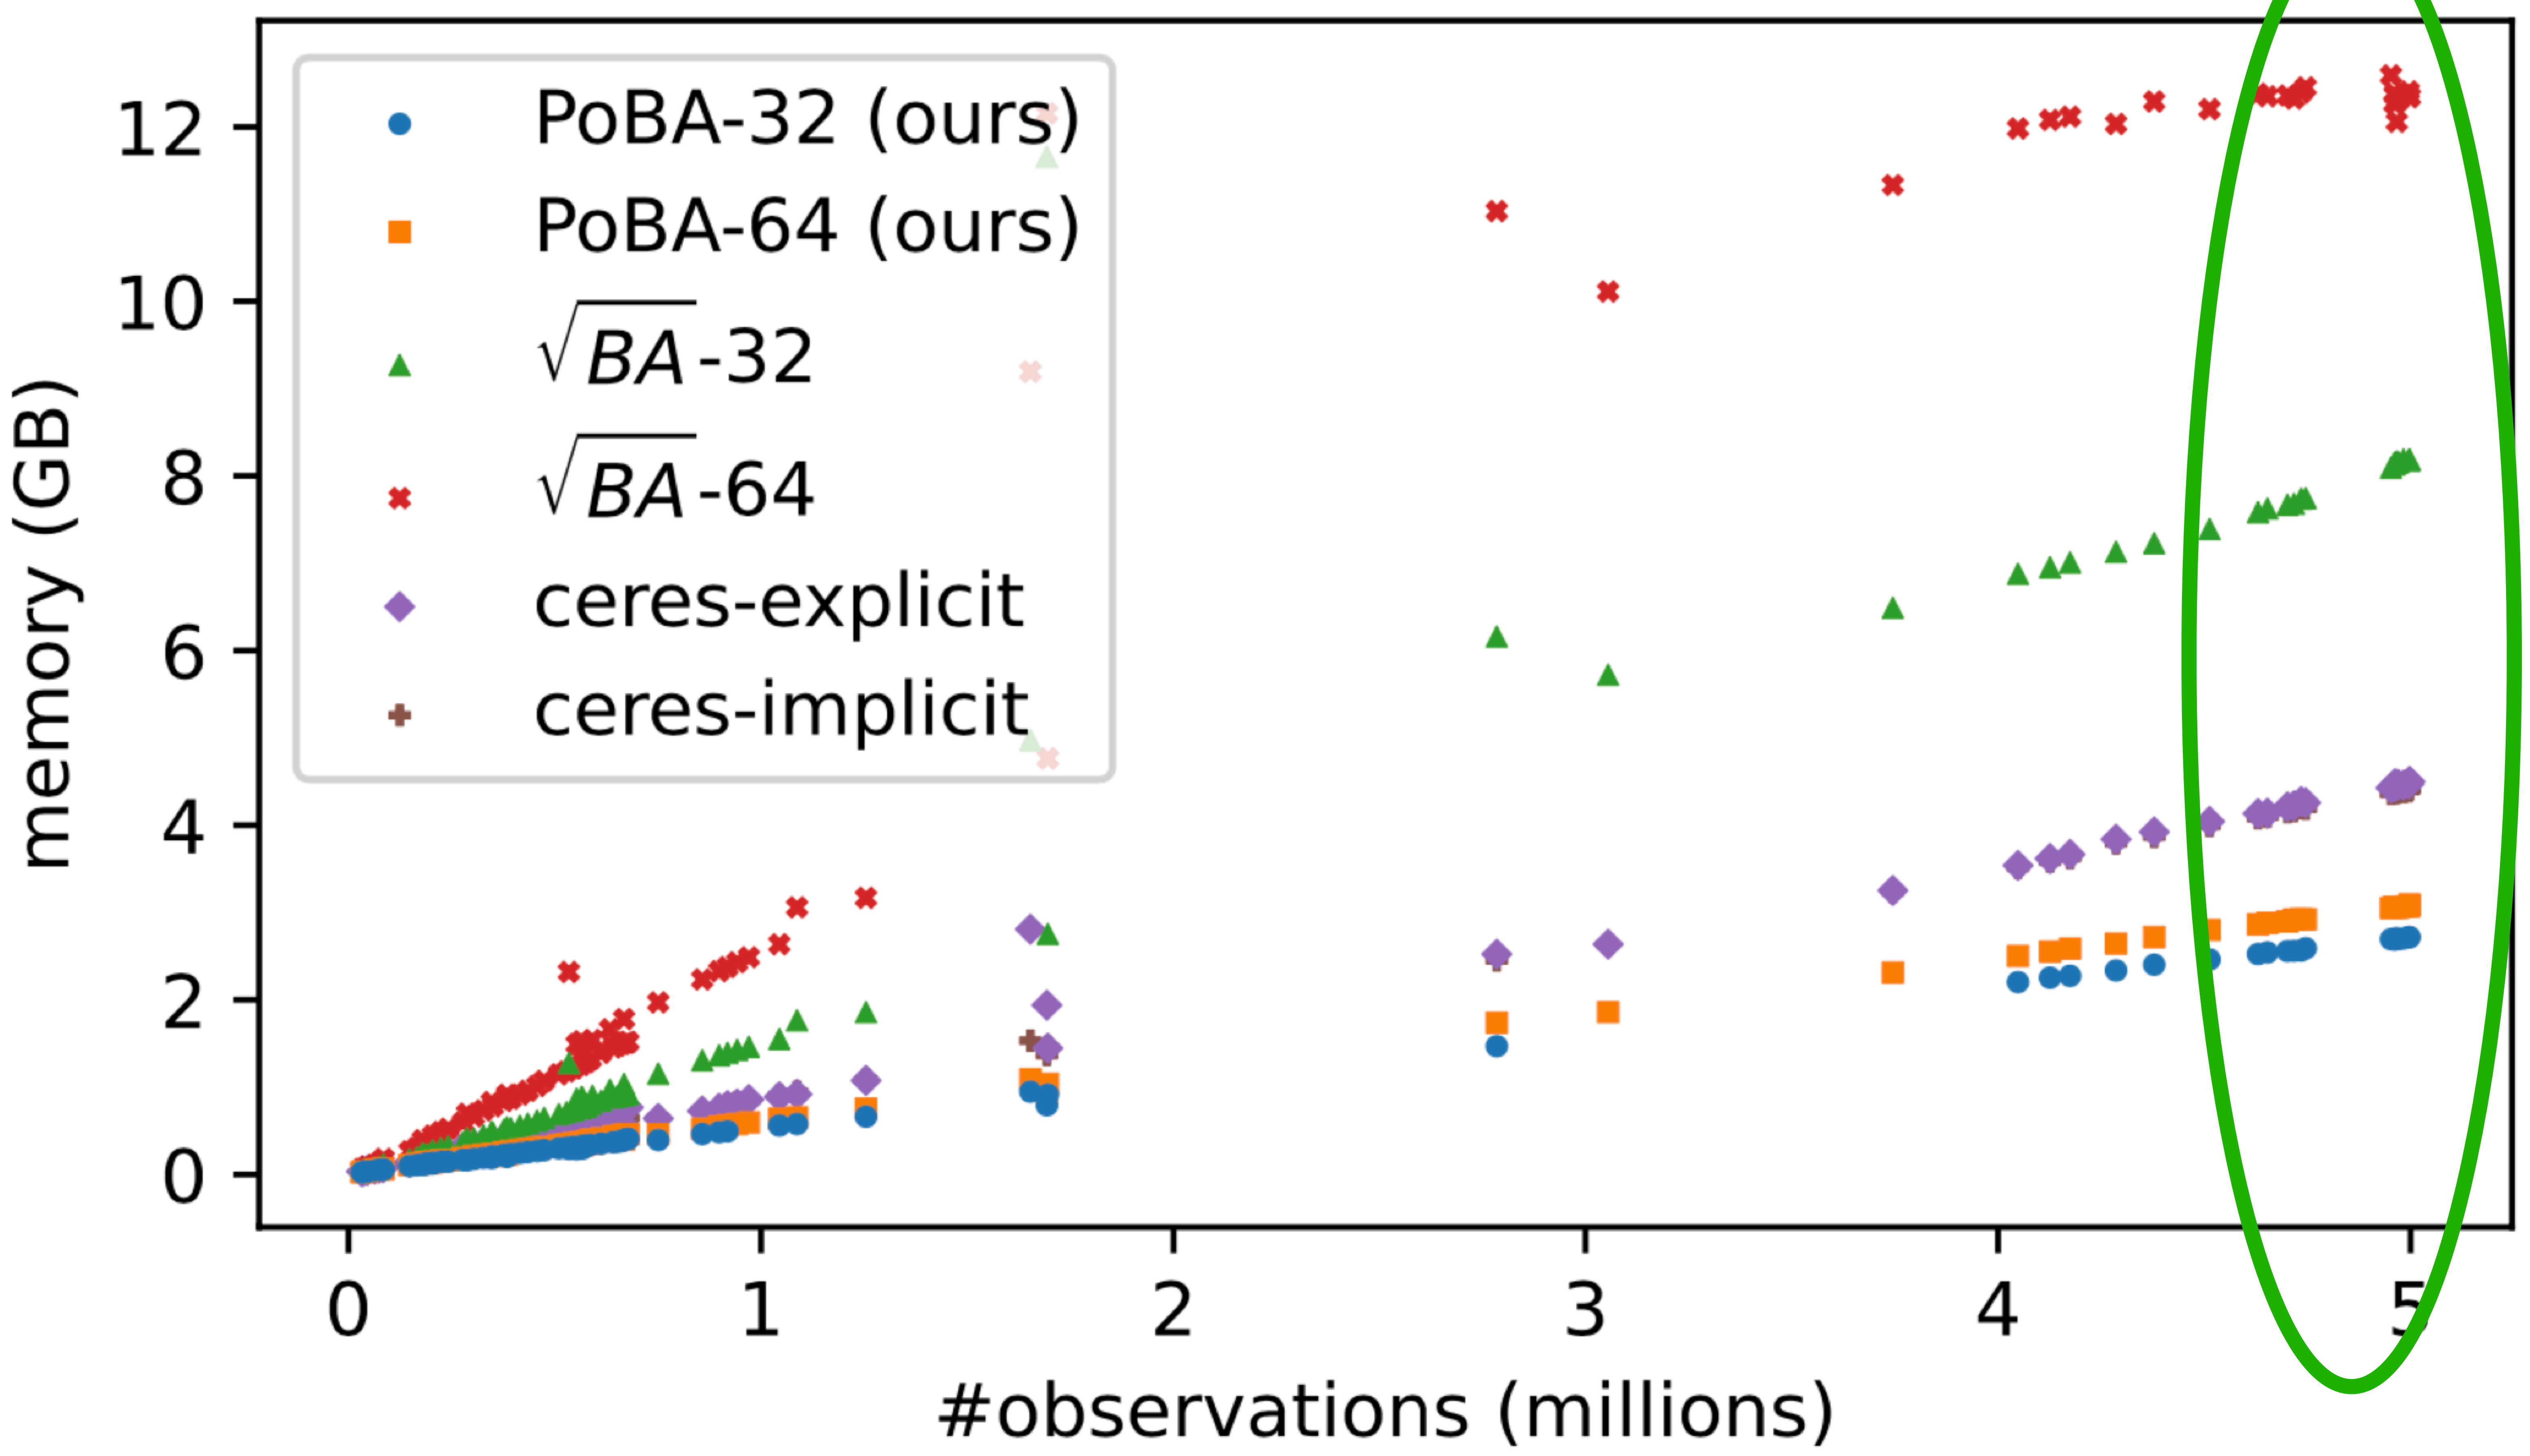

Memory consumption for all BAL problems.

- PoBA requires almost five times less memory than $\sqrt{BA}$ and twice less memory than Ceres.
- By far the **less memory-consuming** solver wrt its challengers.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

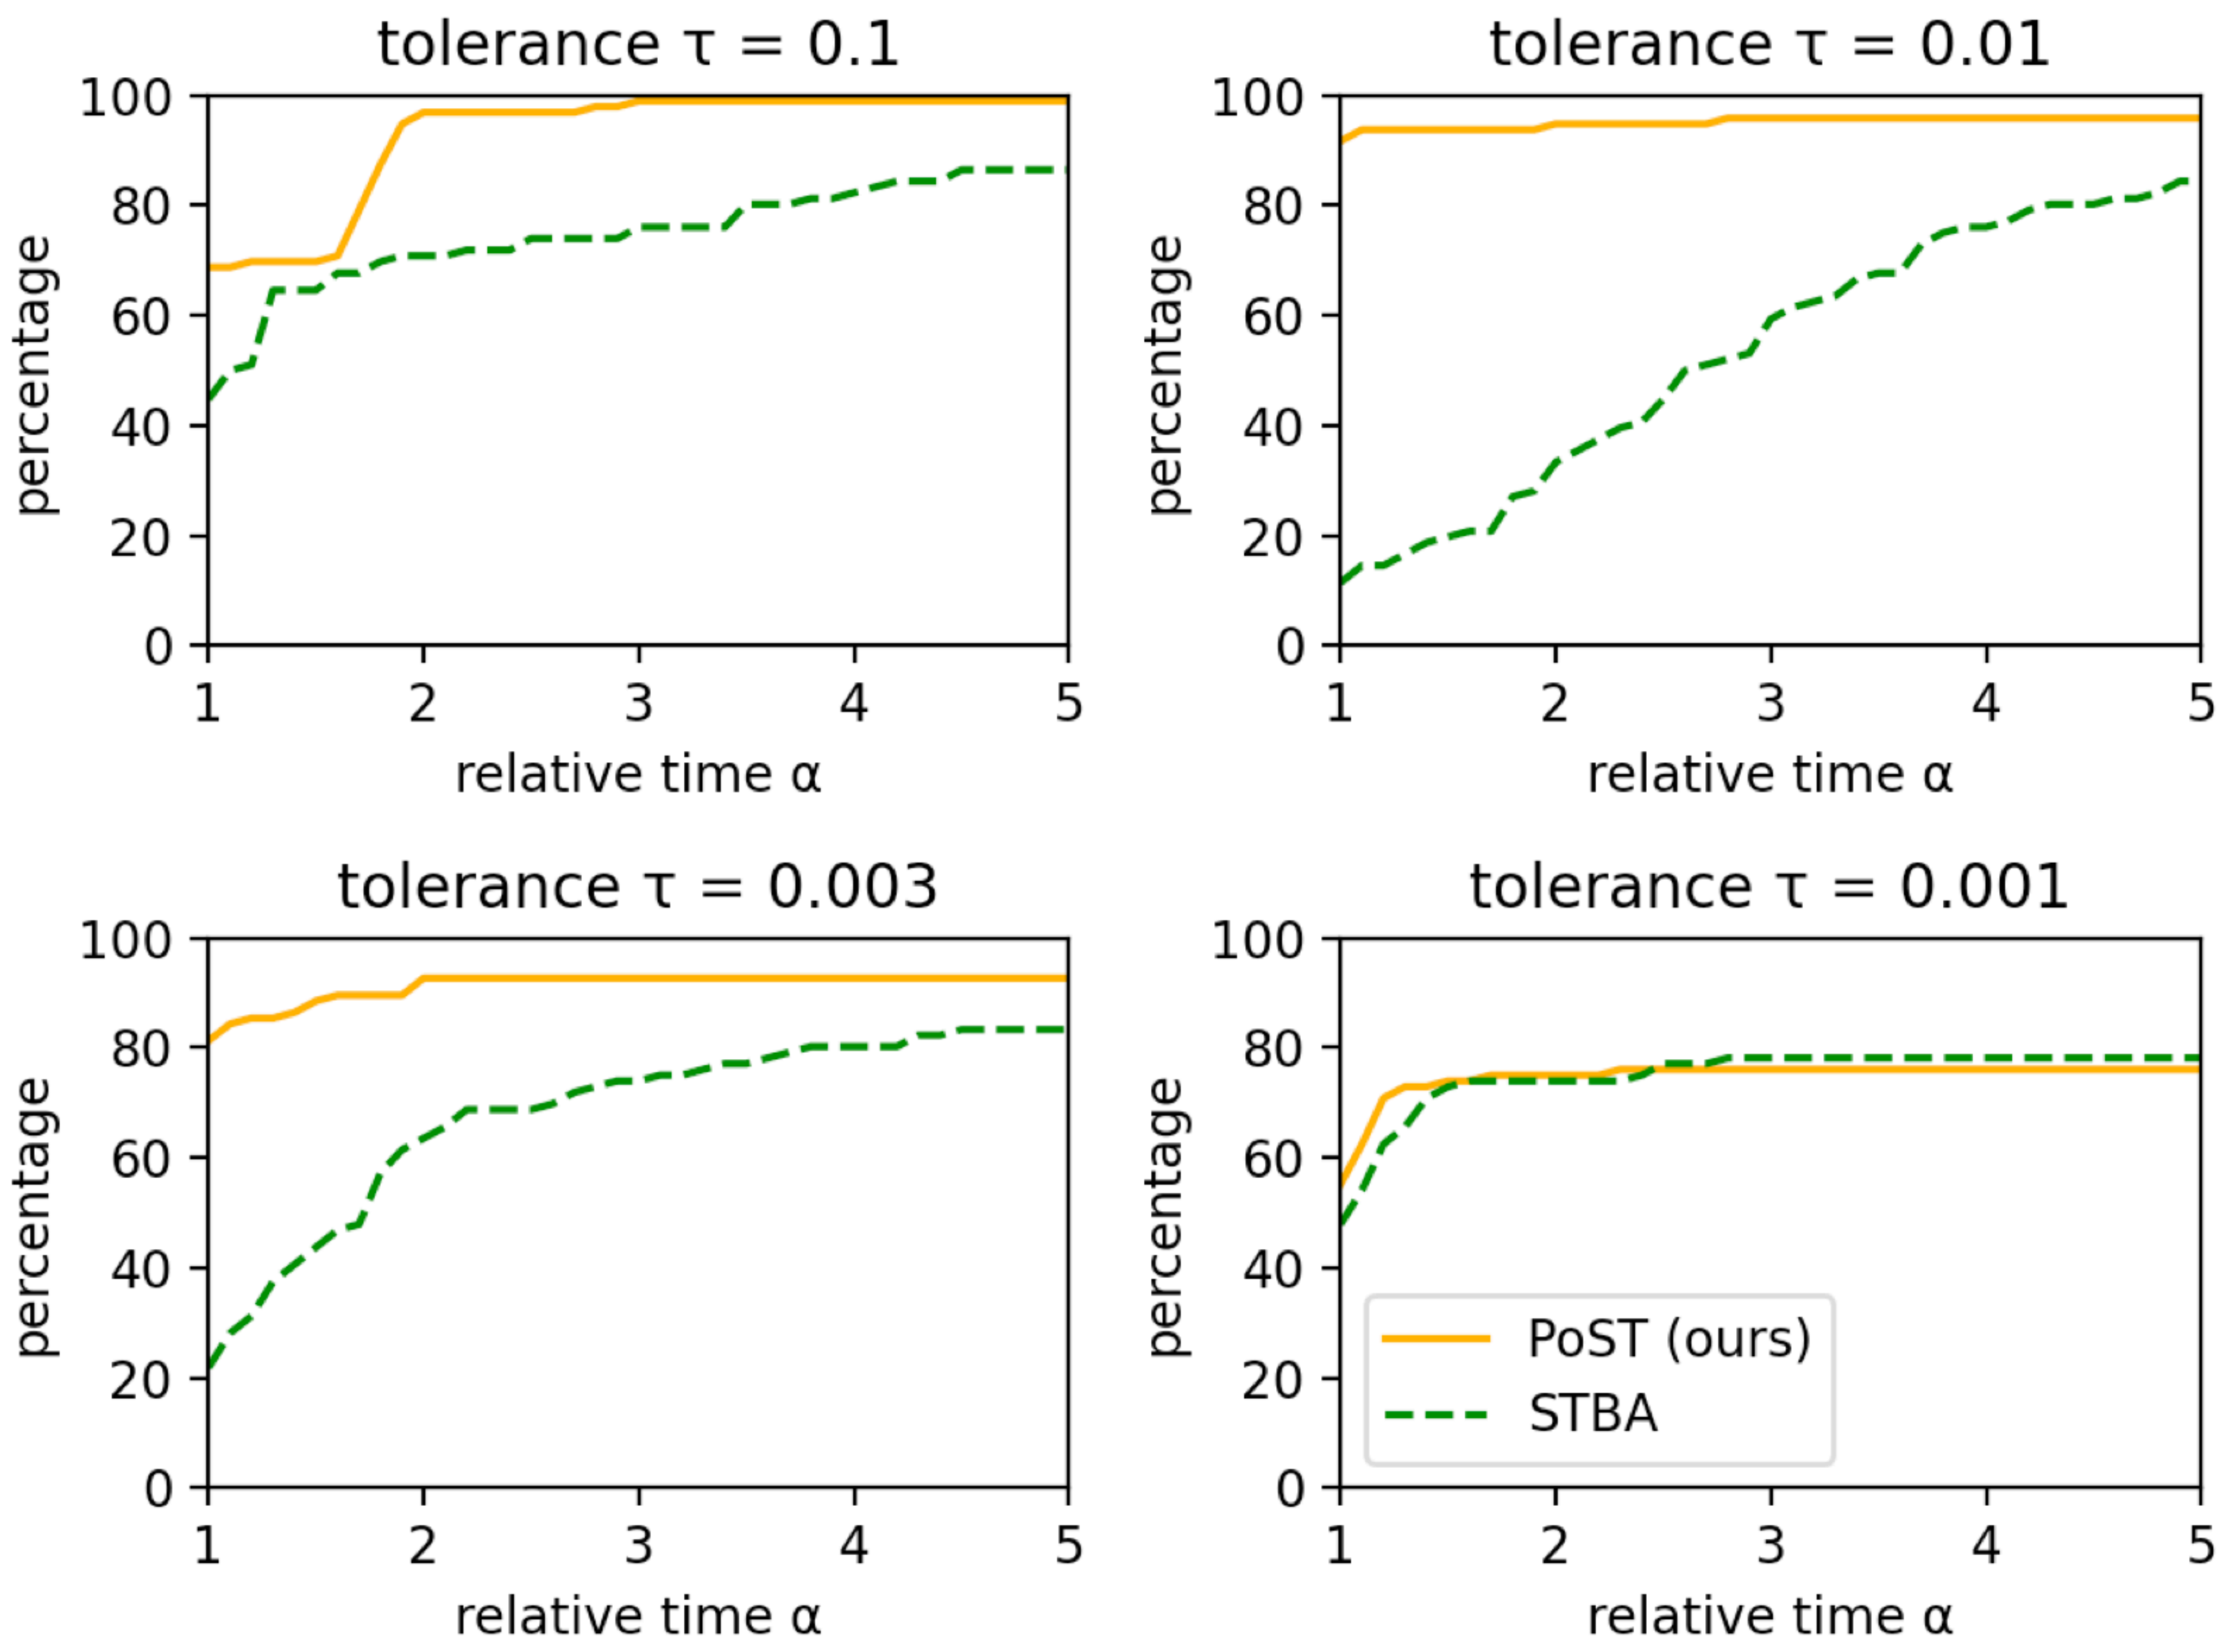

Performance profiles for all 97 BAL problems with a stochastic framework.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

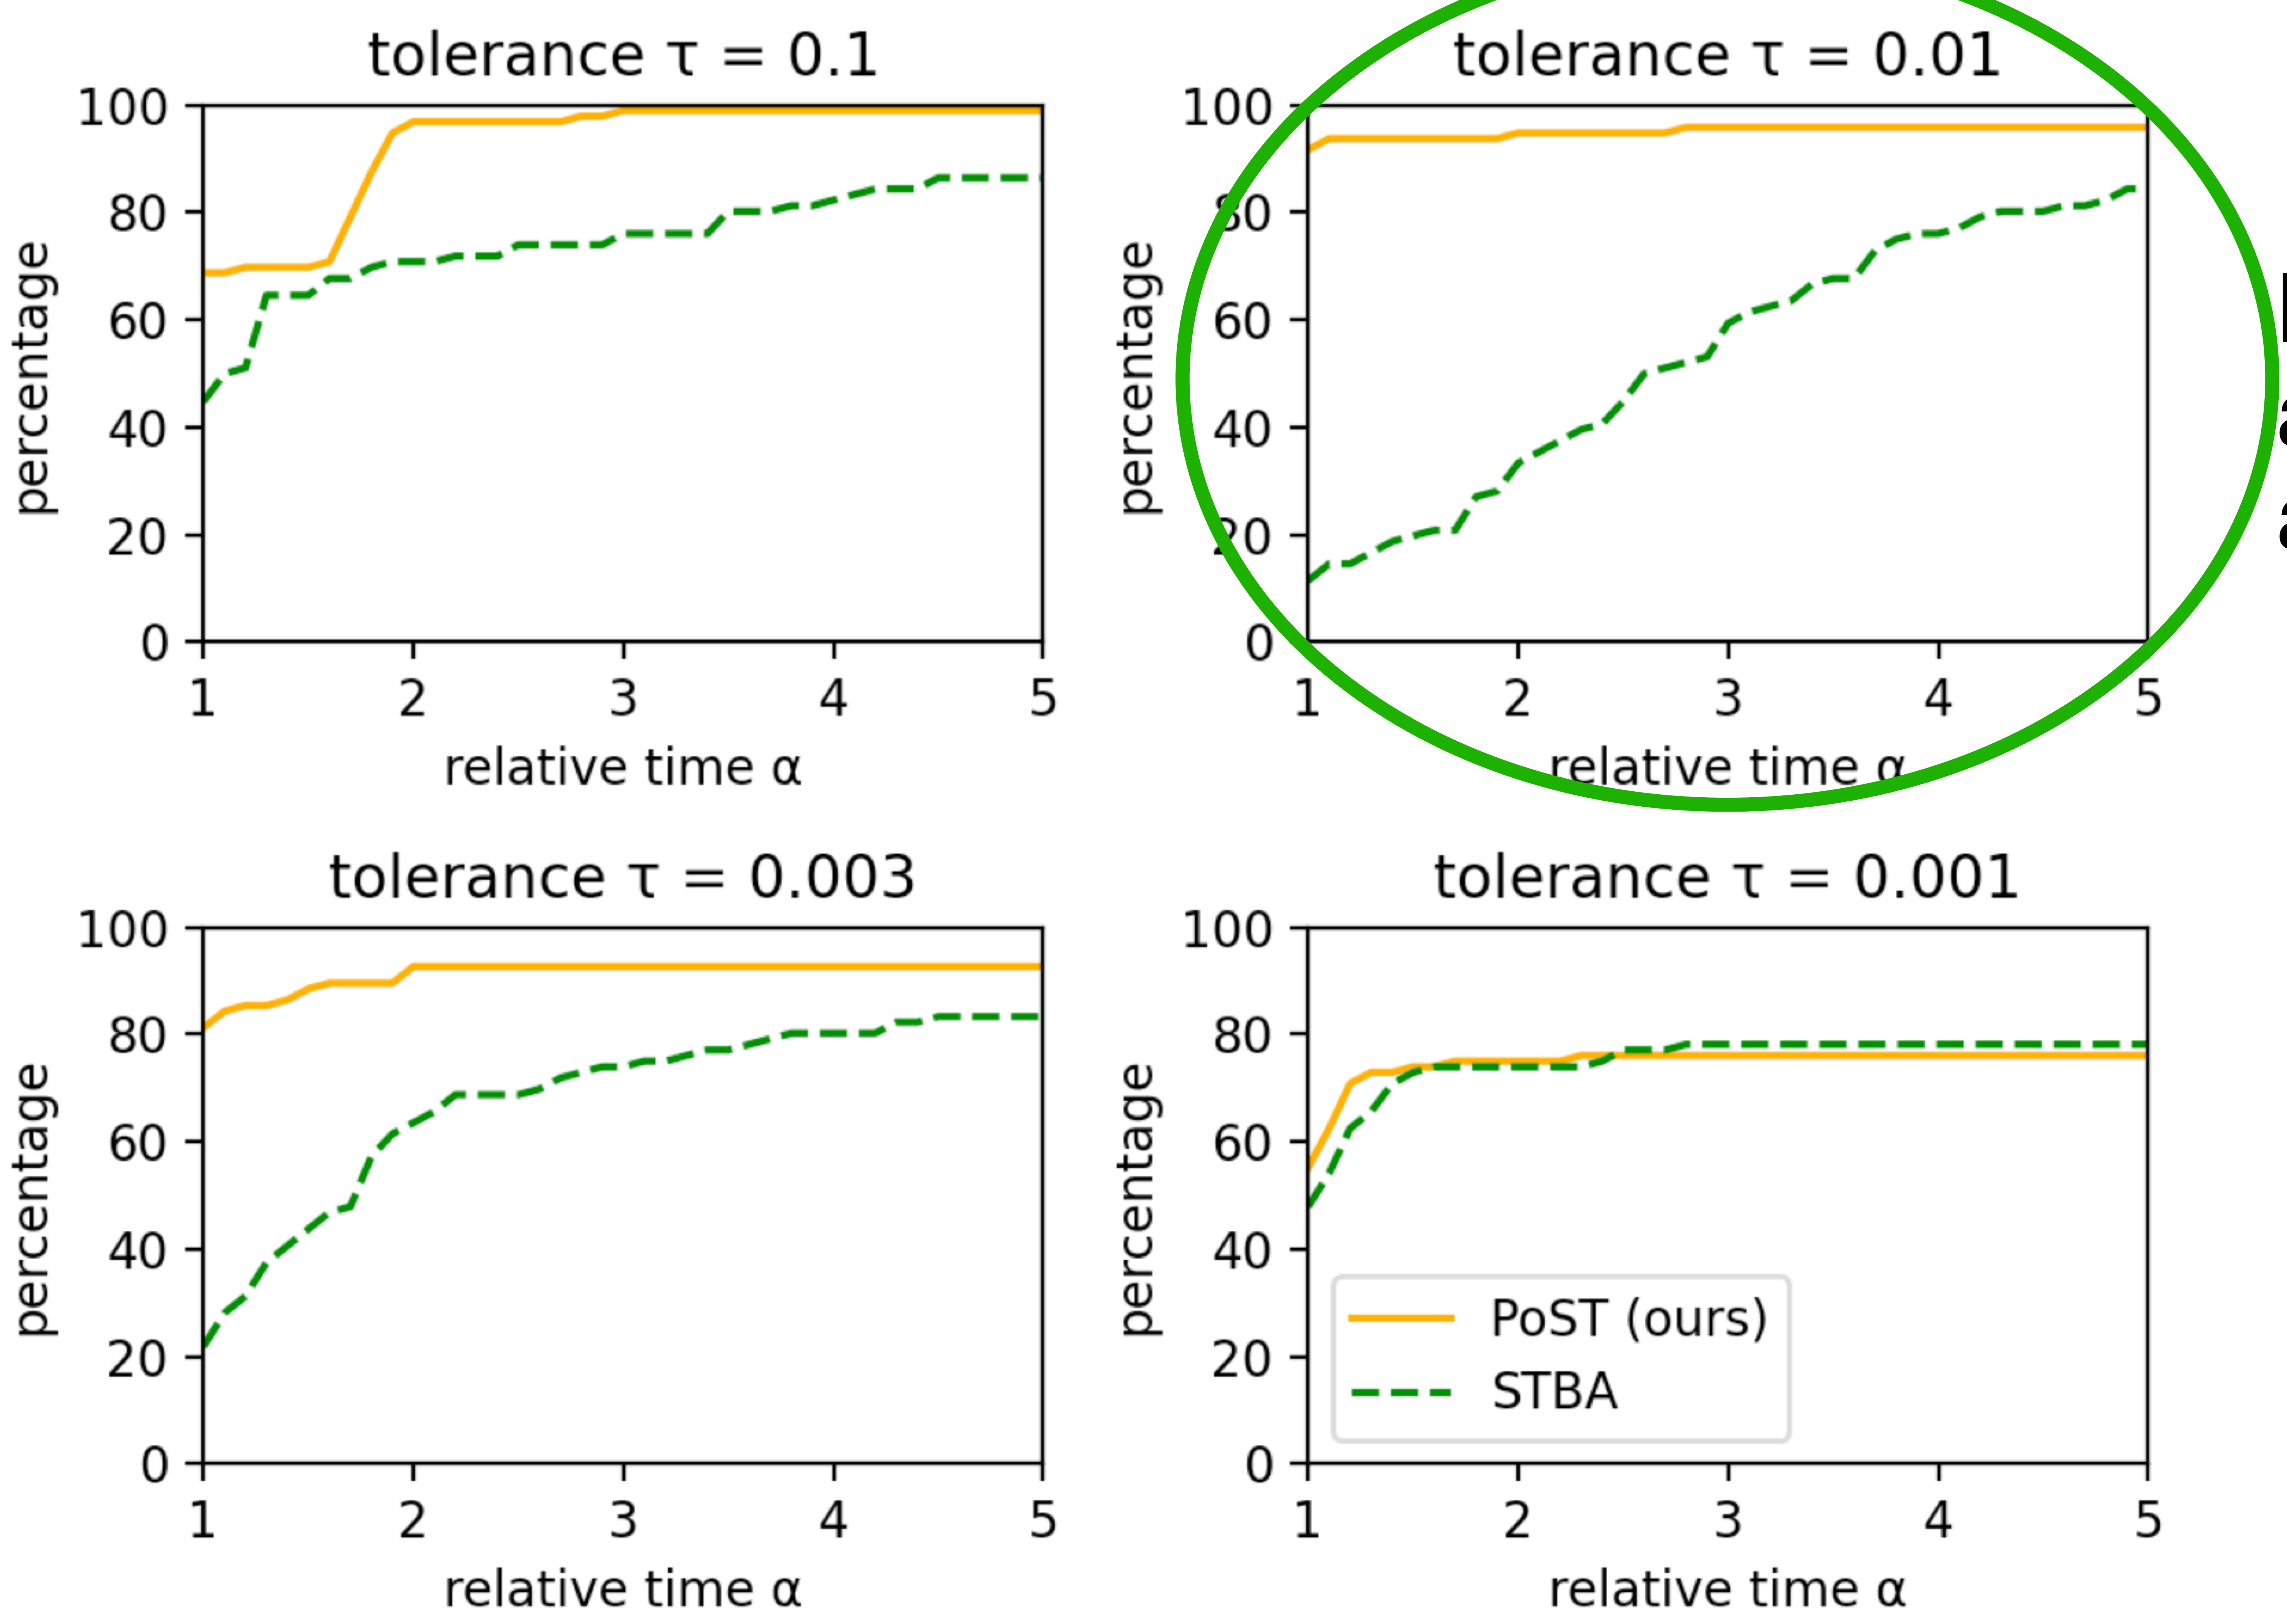

Performance profiles for all 97 BAL problems with a stochastic framework.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

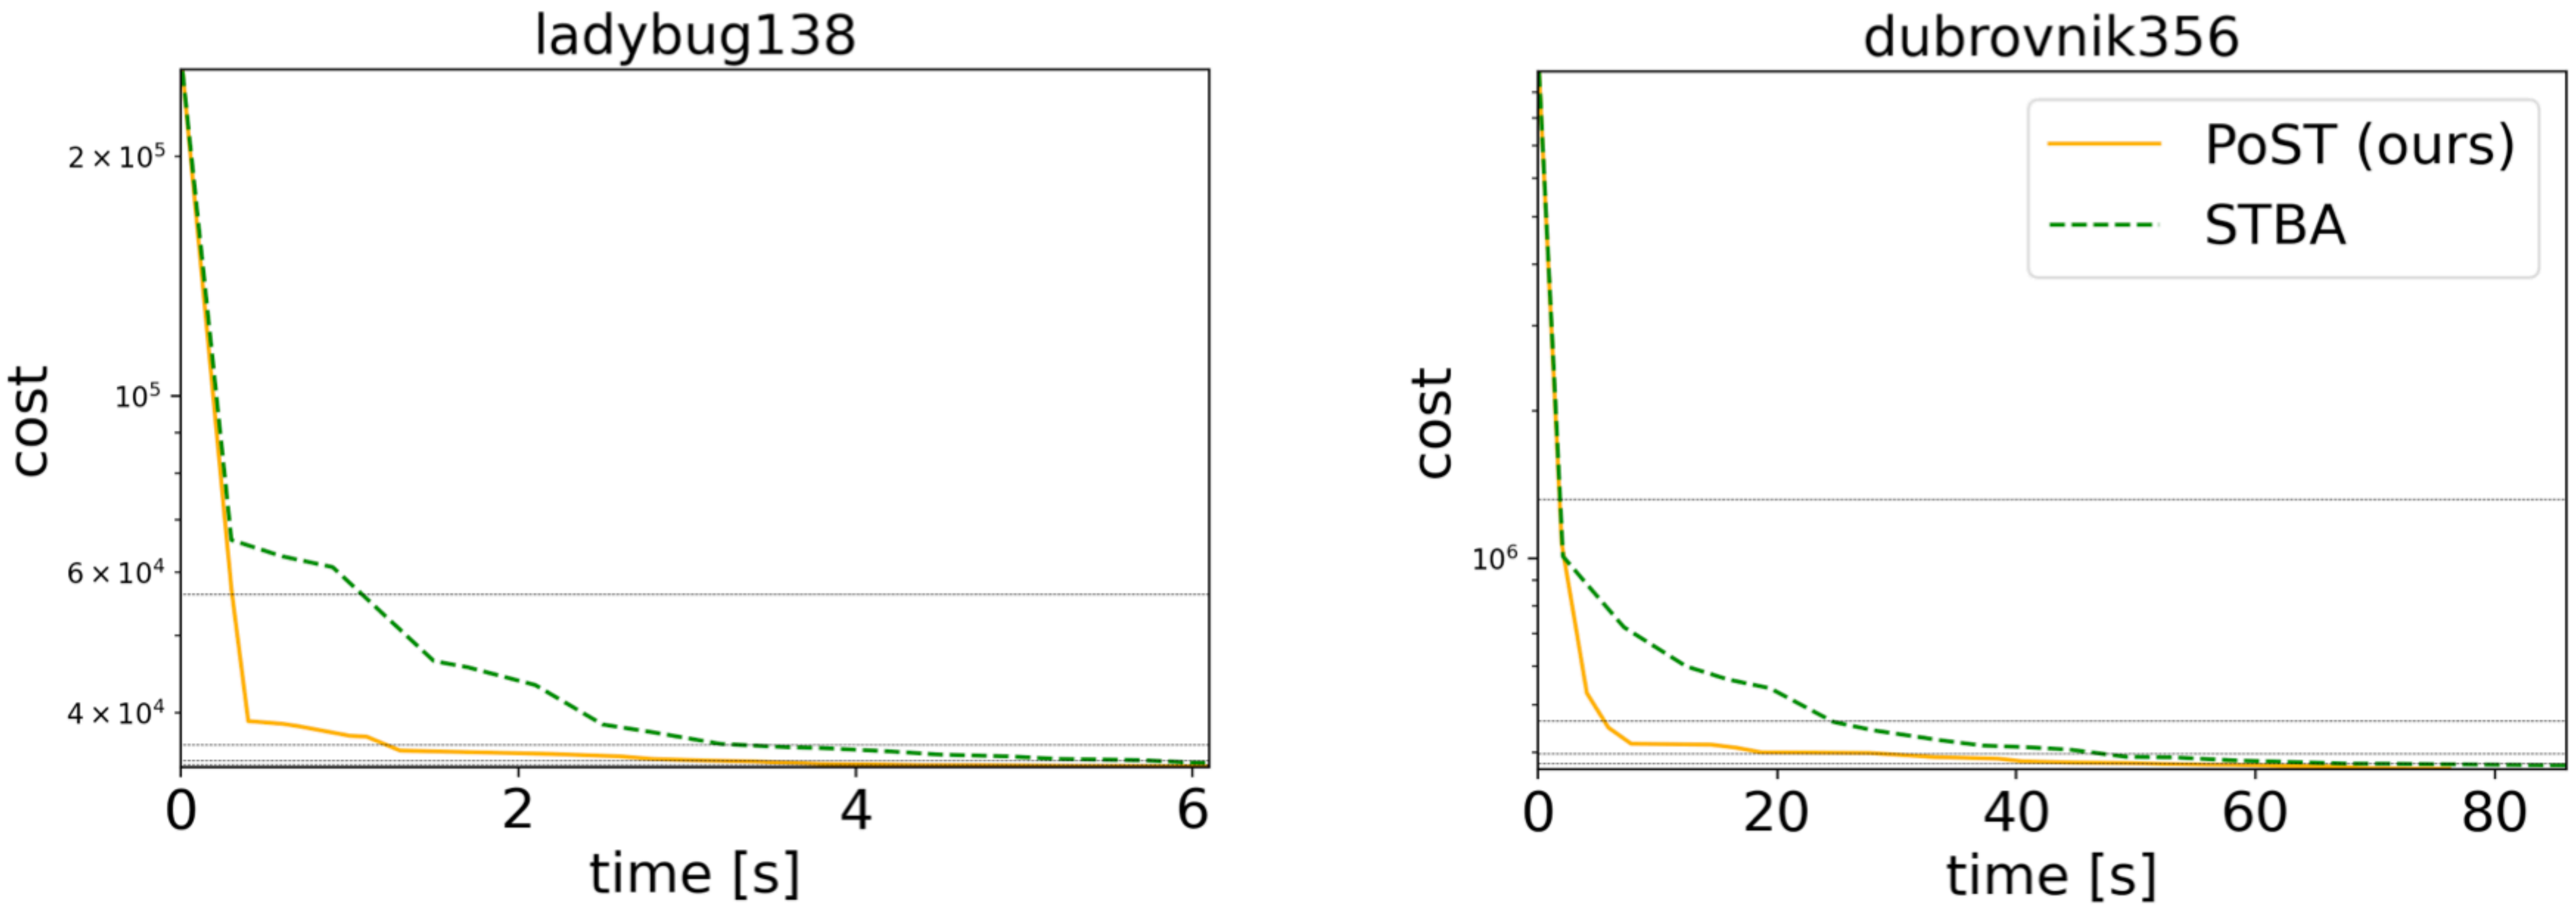

Convergence plots for BAL problems with 138 poses (left) and 356 poses (right) in a stochastic framework.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

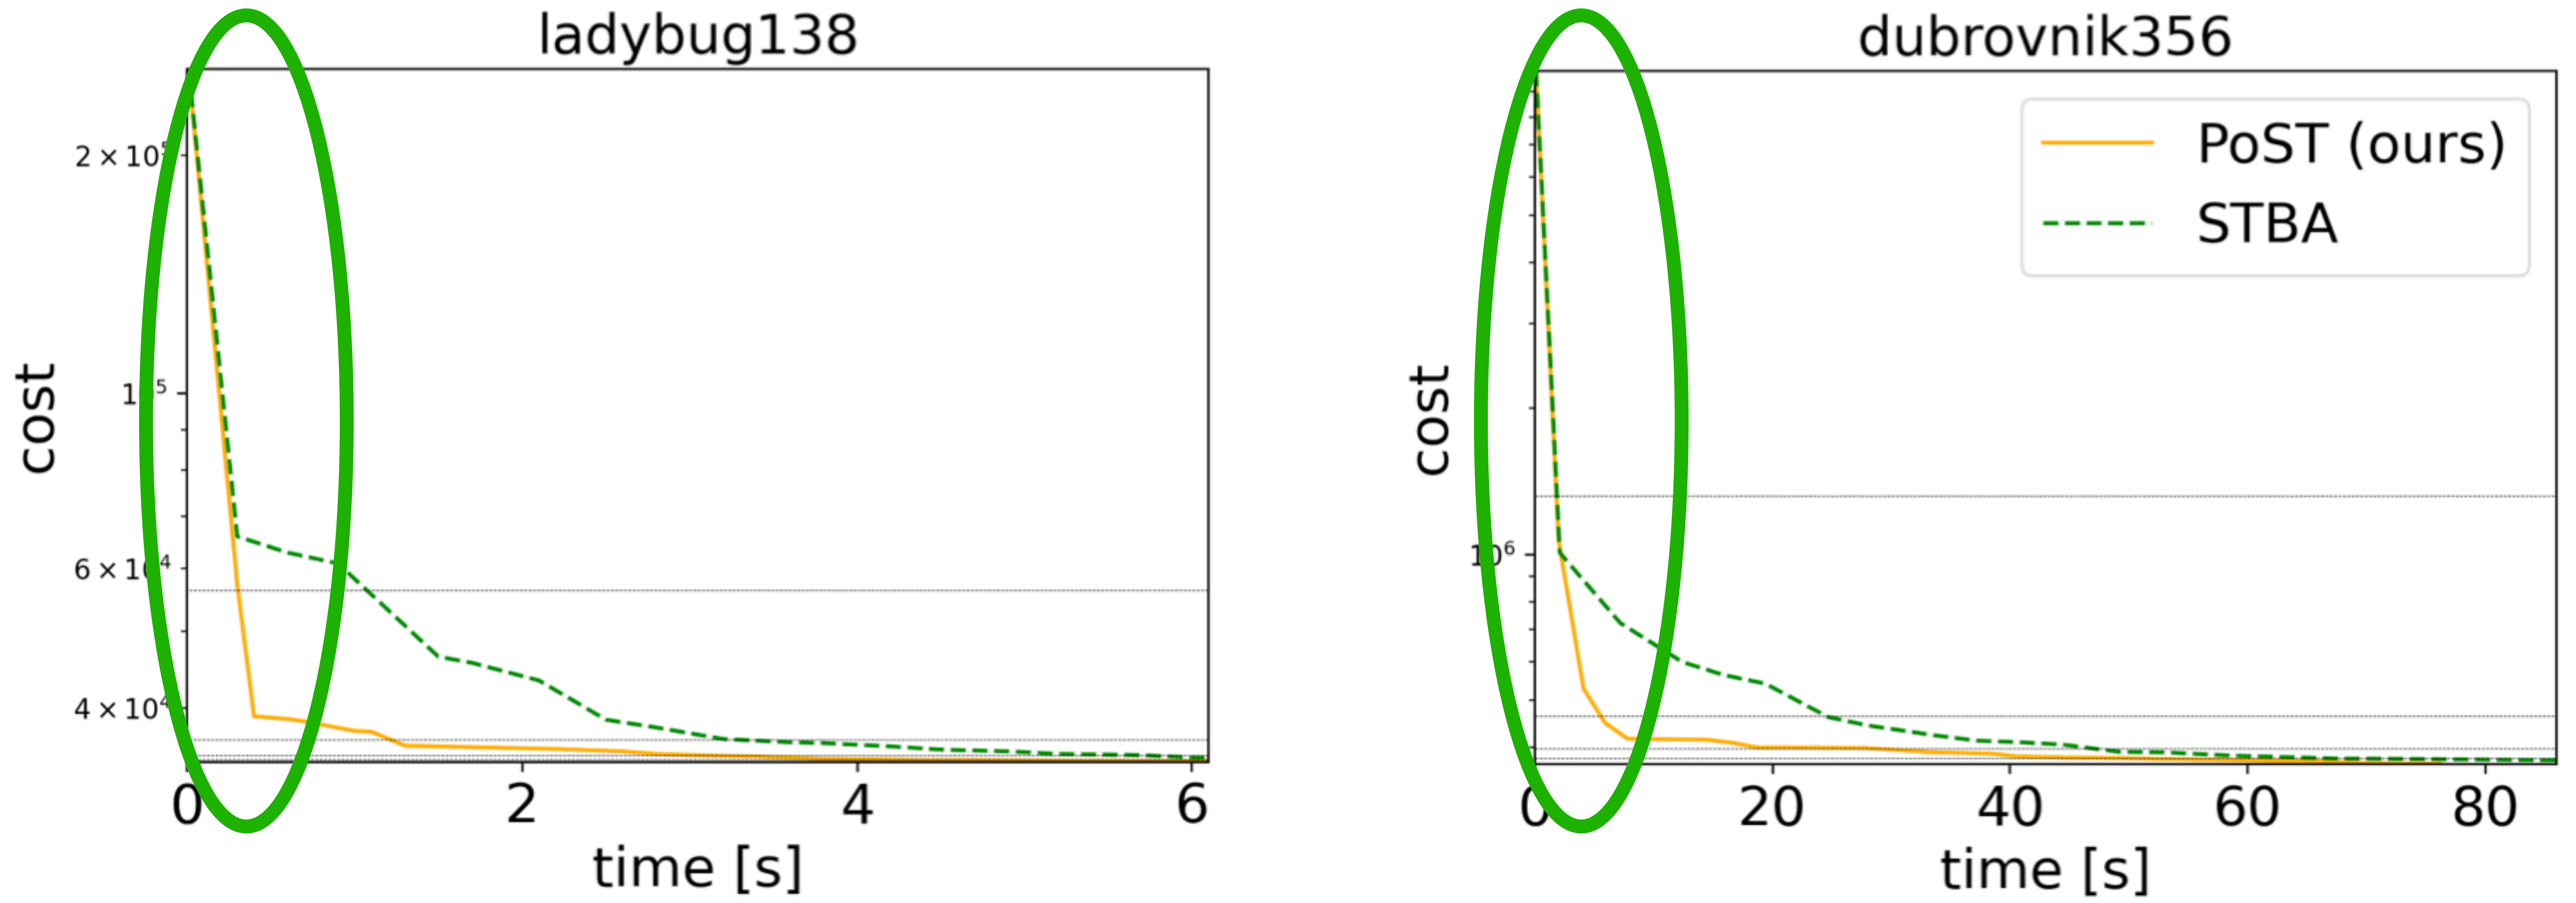

Convergence plots for BAL problems with 138 poses (left) and 356 poses (right) in a stochastic framework.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Results

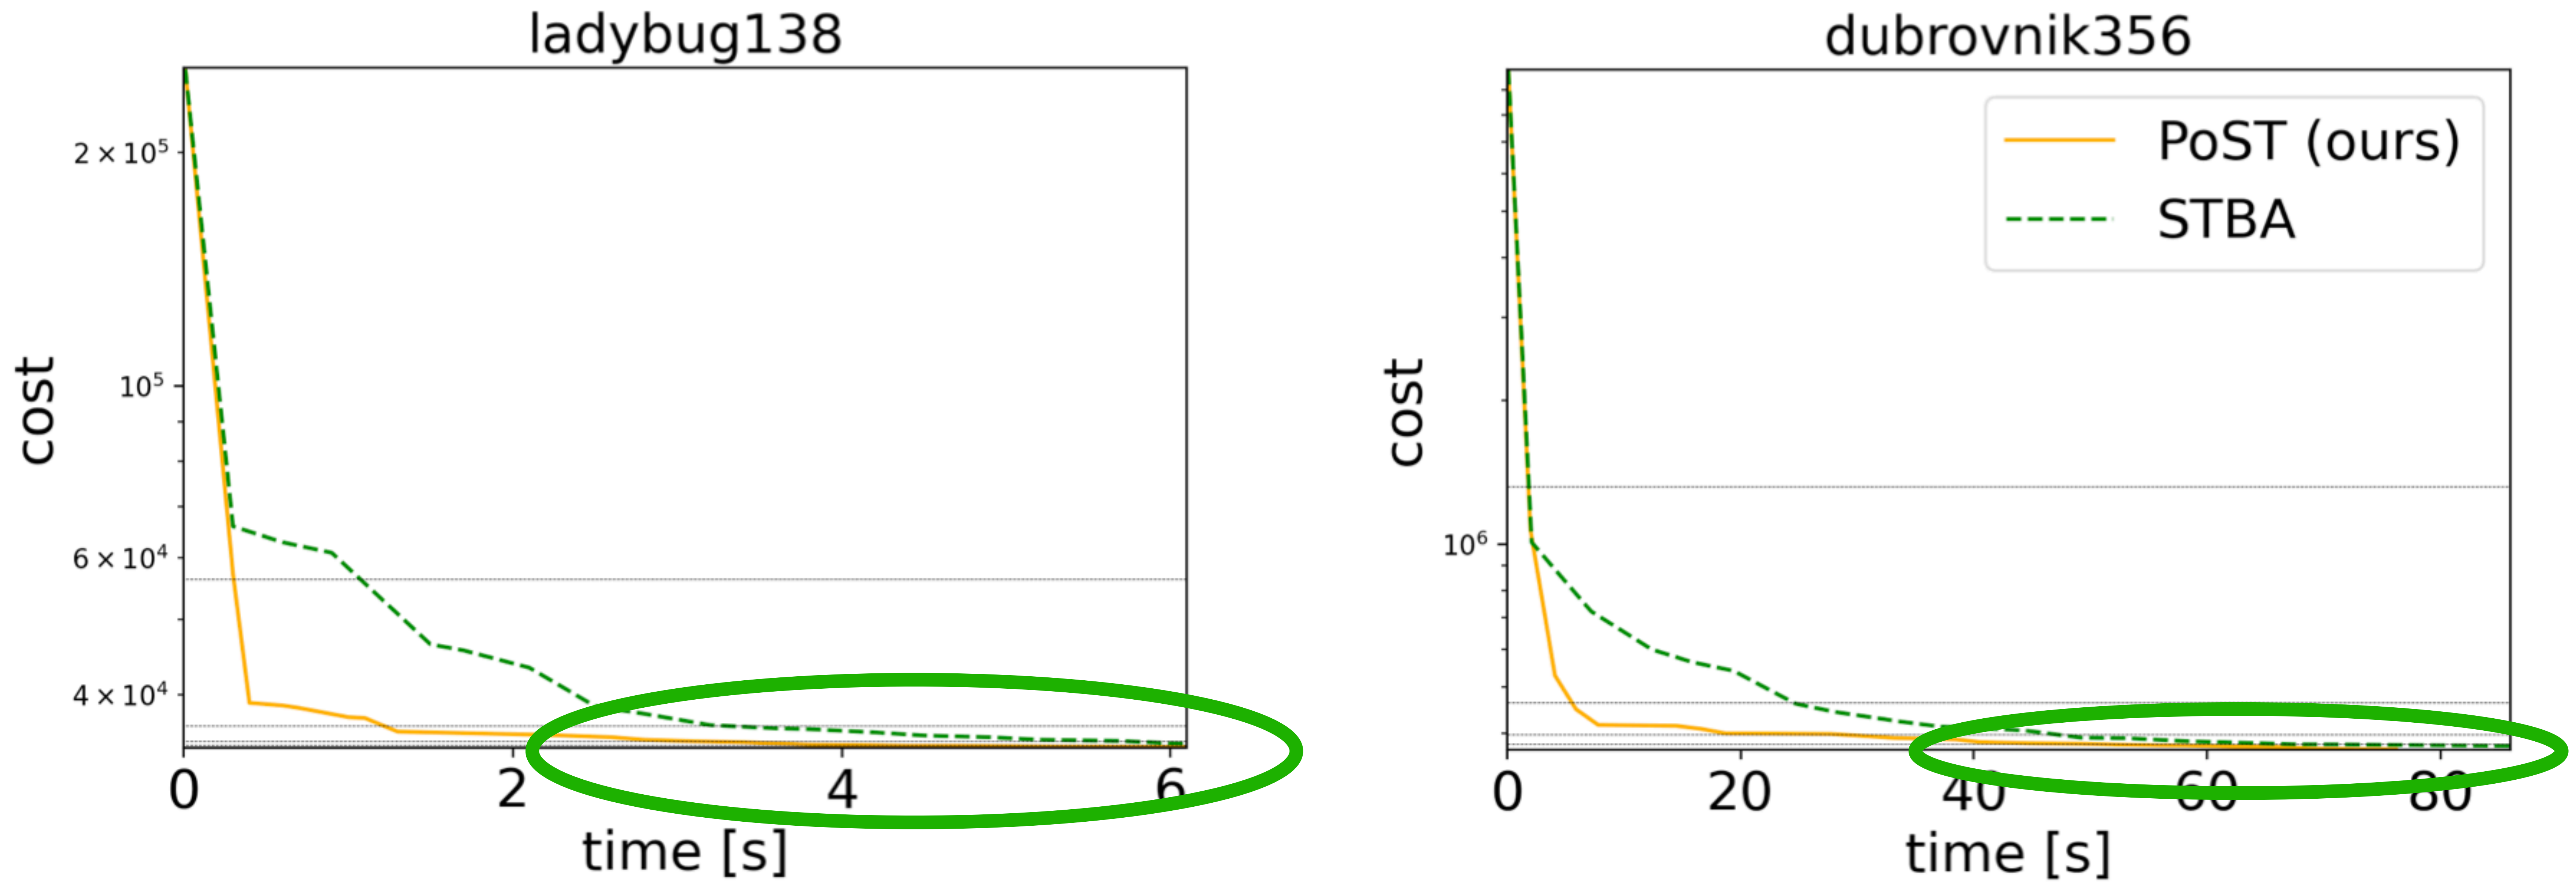

Convergence plots for BAL problems with 138 poses (left) and 356 poses (right) in a stochastic framework.

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23

Conclusion

- Introduce a **new class of large-scale BA solvers** built on a **power expansion** of the inverse Schur complement.
- Prove the **theoretical validity** and the **convergence** of PoBA.
- Experimentally show that PoBA outperforms competitive iterative solvers in terms of **speed**, **accuracy**, and **memory-consumption**.

Open-source implementation:

https://github.com/simonwebertum/poba

Contact:
sim.weber@tum.de

Weber et al., *Power Bundle Adjustment for Large-Scale 3D Reconstruction*, CVPR23