# General and Nested Wiberg Minimization: Additional Results

Dennis Strelow
Google
Mountain View, CA
strelow@google.com

## 1. Bundle Adjustment

In this section we detail the bundle adjustment experiments summarized in Section 6.2 of the main paper, considering convergence, speed, the relative robustness of $L_1$ and $L_2$ bundle adjustment, and the real sequence "rover."

### 1.1. Convergence

We performed two experiments with synthetic data, Bundle Adjustment 1 and Bundle Adjustment 2, which differed in the range of errors in the initial estimates. In both, we used three-dimensional points uniformly distributed on a sphere of radius 10 and a camera moving in an ring of radius 20 around the sphere, looking toward the center. We create 300 instances of this problem, varying: the number of images (5, 10, 15, 20, 25, 30); the number of points (10, 20, 30, 40, 50); and the error in the initial estimates of the camera rotation Euler angles, camera translation, and three-dimensional point positions. In Bundle Adjustment 1, we drew the initial errors from $[-\epsilon, \epsilon]$, for 10 $\epsilon$'s varying exponentially between 0.01 and 1.0. In Bundle Adjustment 2, $\epsilon$ varied between 10.0 and 1000.0, producing essentially random initial estimates. As in the factorization experiments in Section 6.1 of the main paper, we marked 20% of the observations as missing and changed another 10% of the observations to be outliers. In each trial, we ran Wiberg and simultaneous for 400 iterations from the same initial estimate.

Figures 1(a) and 1(b) compare Wiberg and simultaneous for Bundle Adjustment 1. In this experiment, both methods converge to the ground truth residual or less in every trial, which is reflected by the spike at Wiberg residual - simultaneous residual = 0 in 1(a). In a few other cases, one method or the other produced a residual less than the ground truth residual, which is possible because of the outliers in the data. In this case, the method has found camera and point estimates that are not the ground truth camera and point estimates but have a lower residual. In this experiment, Wiberg sometimes converged in significantly fewer iterations, shown as the bars on the left of 1(b).

Figures 1(c) and 1(d) similarly compare Wiberg and simultaneous for Bundle Adjustment 2. In this harder experiment, Wiberg usually converged to a lower residual than simultaneous, shown by the bars on the left in 1(c), but not always to the ground truth residual. In this experiment, Wiberg had a more negligible advantage in iterations required for convergence, shown in 1(d).

### 1.2. Speed

To compare the speed of Wiberg and simultaneous iterations, we used synthetic data sets similar to those in Bundle Adjustment 1 and Bundle Adjustment 2, but varying the number of images over (2, 3, 4, 5, 7, 11) and the number of points over (10, 32, 100, 316, 1000). Figure 2 shows the time for the main linear programming solve for each of these problem sizes, for Wiberg and simultaneous. As with factorization, the time for both grows quickly with both the number of images and the number of points. The Wiberg solve is faster only for 2 or 3 images, which the simultaneous solve is faster for 5 or more images. For each number of images, the relative speed of Wiberg and simultaneous is not affected by the number of points.

### 1.3. $L_1$ versus $L_2$

In Figure 3, we show synthetic results comparing the accuracy of $L_1$ and $L_2$ bundle adjustment in the presence of outliers. We used a synthetic dataset with 30 points on the surface of a sphere, and 10 camera positions moving in a ring around the sphere, looking at the center of the sphere. We performed 300 trials, varying the fraction of outliers in the observations (30 values exponentially spaced over 0.01 to 0.3) and the initial error in the estimates (10 values exponentially spaced over $10^{-3}$ to $10^{-1}$). For each fraction of outliers, we averaged the final camera translation error and point error across the ten choices of initial estimate error; the plots show these averages. In these experiments, $L_2$ is sensitive to any outliers in the observations. The camera estimates are immune to outliers up to 16.7% outliers in the observations, but have errors comparable to $L_2$ above that fraction. The point estimates are immune or highly tolerant to outliers up to 11.8% outliers in the observations, then

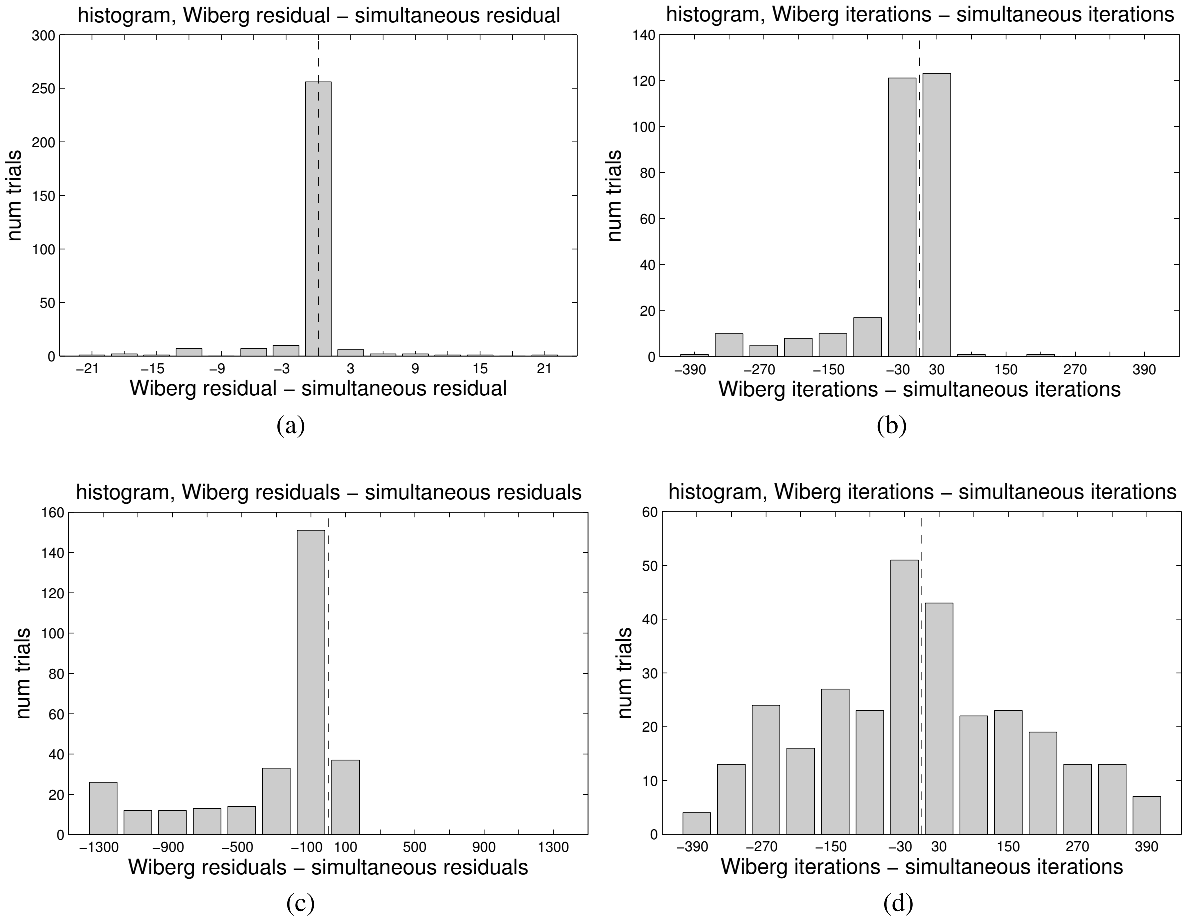


Figure 1. Bundle Adjustment 1 and Bundle Adjustment 2 residuals and iterations until convergence (Section 1.1). For Bundle Adjustment 1, which had small errors in the initial estimates, (a) shows the difference in final residuals for Wiberg and simultaneous, while (b) shows the difference in the number of iterations for convergence for Wiberg and simultaneous. Histograms (c) and (d) show the same results for Bundle Adjustment 2, which had larger initial errors in the estimates.

approach errors comparable to $L_2$ above that fraction. The higher resistance of the camera translations to outliers reflects the fact that each camera translation depends on 30 observations whereas each point depends on only 10.

### 1.4. Rover Sequence

Figure 2(a) in the main paper shows an example image from a real sequence, "rover," with tracked points shown as black dots. The camera looks out from the back of a rover while the rover executes three large loops. The sequence includes about 700 images and about 10,000 three-dimensional points.

We used the Wiberg bundle adjustment algorithm and because of the large problem size, we used the primal solve strategy described in Section 4.3. The algorithm correctly recovers the structure and motion, and the estimated motion is shown in Figure 2(b) in the main paper. The result is locally correct and consistent with the global motion estimates from GPS and odometry. As an initial estimate in this example, we used the estimate from an $L_2$ Kalman filter with large errors added to the camera and point estimates. We picked this example because it cannot be solved using factorization and affine structure-from-motion – perspective effects are extremely strong because of the large difference in distance to the near and far points.

## 2. Projective Bundle Adjustment

### 2.1. Convergence

To compare the convergence of the Wiberg and simultaneous algorithms for projective bundle adjustment, we conducted four tests, Projective Bundle Adjustment 1-4 (PBA 1-4). The synthetic data and initial errors for these are similar to Bundle Adjustment 1 and 2 above. In each, we again used three-dimensional points uniformly distributed on a sphere of radius 10 and a camera moving in an ring of ra-

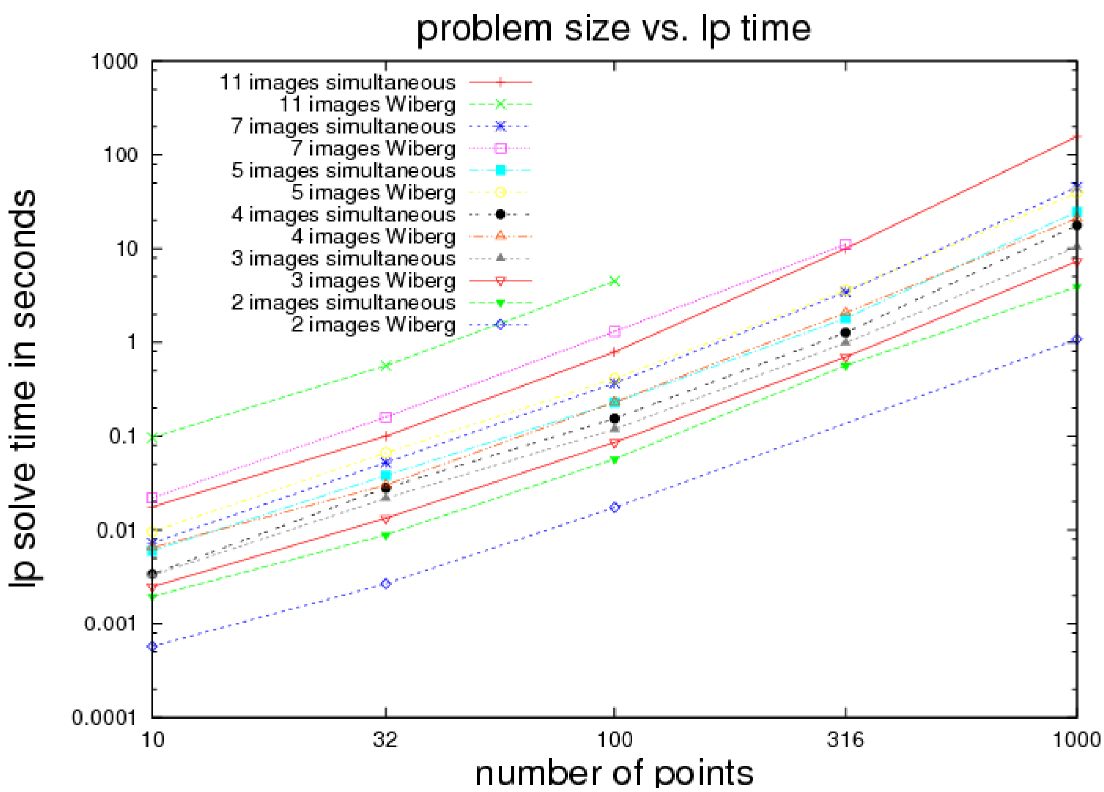


Figure 2. The average time per linear programming solve for the Wiberg and simultaneous methods, by problem size (Section 1.2). Wiberg solves are faster for problems with 2 or 3 images, and simultaneous solves are faster for problems with 5 images or more. Both axes are logarithmic.

dius around the sphere, looking toward the center. We create 300 instances of this problem, varying: the number of images (10, 12, 14, 17, 21, 25); the number of points (7, 10, 15, 22, 33); and the error in the initial estimates of the camera rotation Euler angles, camera translation, and three-dimensional point positions. Unlike Bundle Adjustment 1 and 2 above, we did not remove observations or add outliers in PBA 1-4. We ran each method for 100 iterations.

The four experiments varied in the camera ring radius and the initial errors. In Projective Bundle Adjustment 1 and 2, the radius of the camera ring is 20, the same as Bundle Adjustment 1 and 2 above. In Projective Bundle Adjustment 3 and 4, we create a wider range of projective depths by moving the camera closer to the sphere, to 10.1. We varied the initial estimate errors between 0.001 and 0.1 in Projective Bundle Adjustment 1 and 3, and between 1.0 and 100.0 in Projective Bundle Adjustment 2 and 4. In all cases, we used 1.0 as the initial estimate for projective depths, which is a standard initial estimate for projective factorization and corresponds to an affine projection assumption.

Figures 4 summarizes the convergence for PBA 1 and 2, and Figure 5 summarizes the convergence for PBA 3 and 4. The left of each figure shows histograms of the differences in final residuals between the two methods, while the right of each figure shows histograms of the differences in the number of iterations required for convergence to a strict threshold. Comparing the four graphs, we see that Wiberg is able to converge more effectively to strict residual thresholds, reflected in Wiberg's better convergence in PBA1 and PBA3, while simultaneous can produce better residuals in the case of extremely large initial errors, although not typically converging to zero residual in these cases; and Wiberg handles large variation in projective depths better, as shown in PBA3. The right in each figure shows that Wiberg consistently requires fewer iterations to reach its final residual, within a small tolerance.

### 2.2. Speed

Figure 6 shows the time required for the outer linear programming solve for both Wiberg and simultaneous. The plots show the solve time in seconds for various problem sizes, with an experimental configuration similar to Projective Bundle Adjustment 2 but with a wider range of problem sizes. In this case, Wiberg is faster up to 51 points, with simultaneous becoming faster before problem sizes reach 100 points. In this example we performed one minimization of

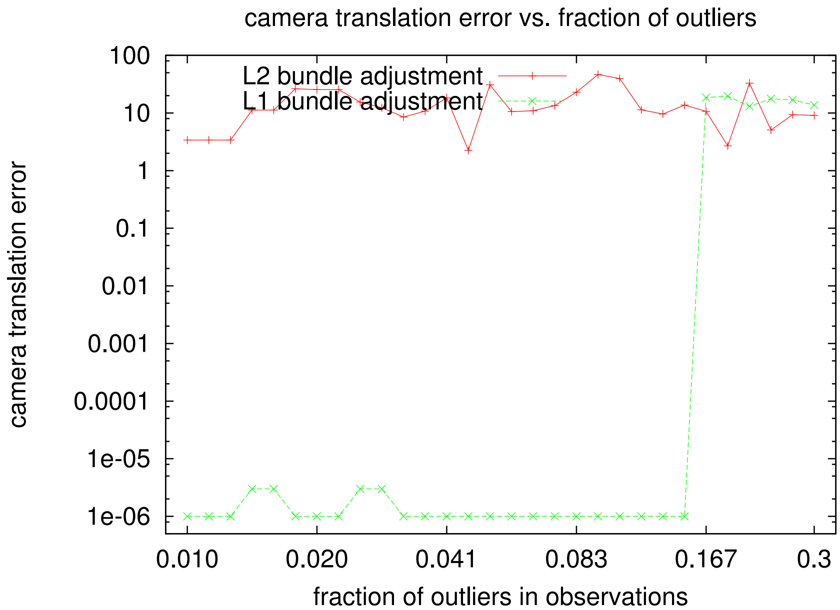


(a)

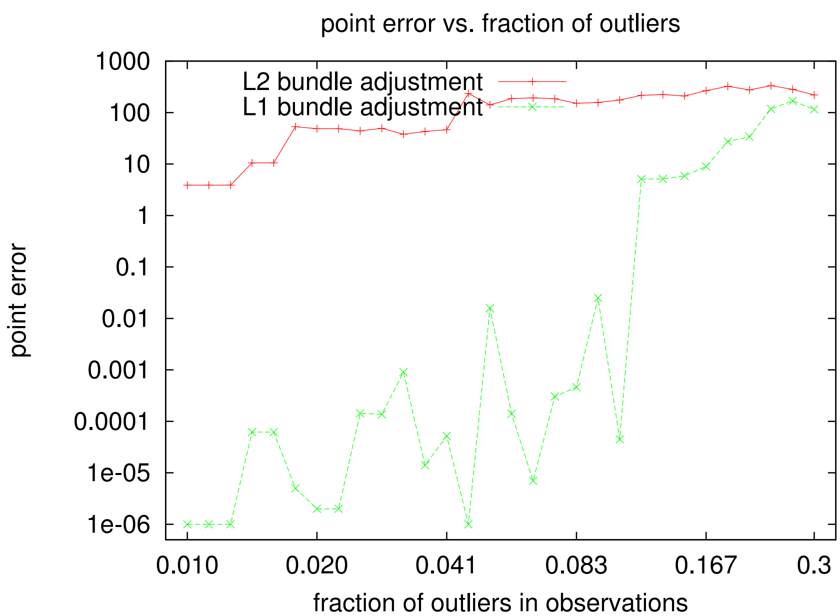


(b)

Figure 3. $L_1$ and $L_2$ translation (a) and point (b) errors when there are outliers in the observations (Section 1.3). In this experiment the $L_2$ estimates are sensitive to any number of outliers, while the $L_1$ translations and points are largely immune to up to 16.7% and 11.8% outliers, respectively. Both axes are logarithmic.

each problem size shown for 10 iterations and averaged the linear programming time across the 10 iterations.

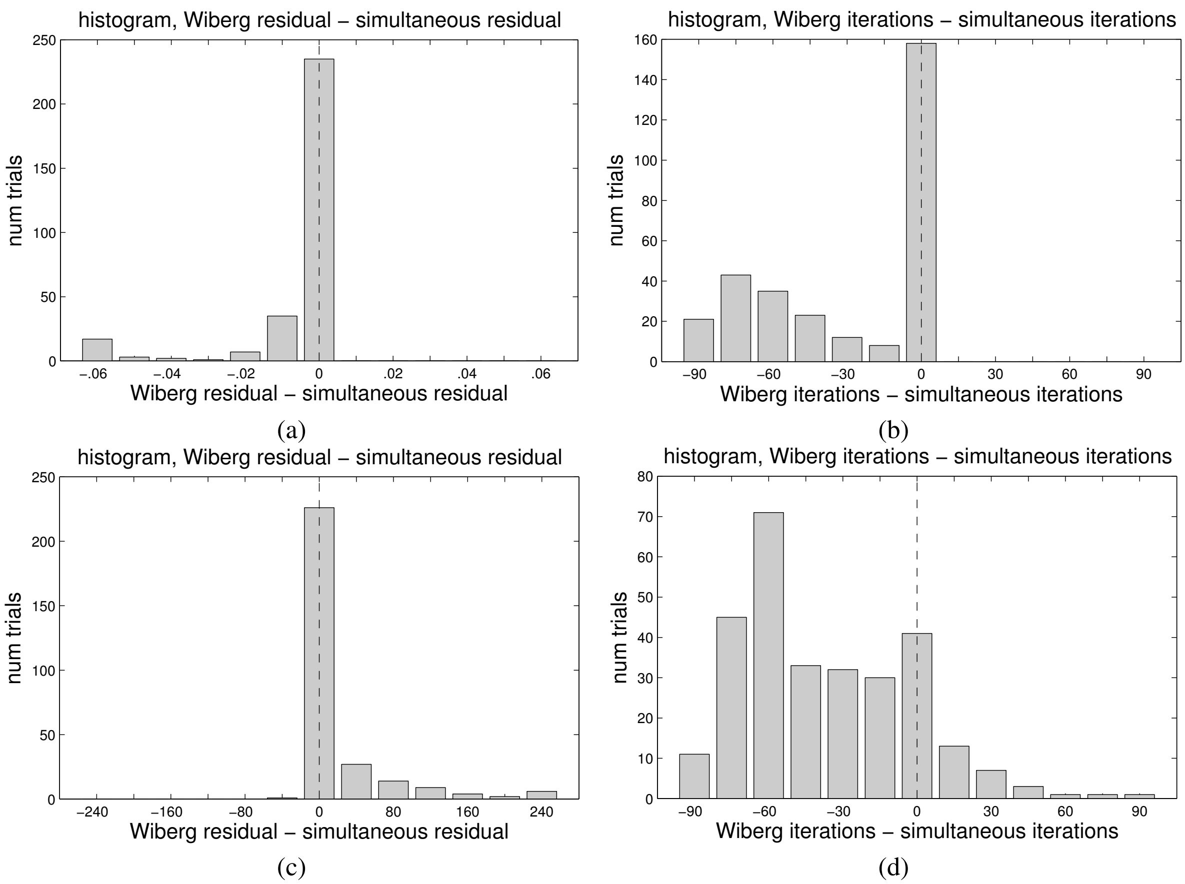


Figure 4. Results for Projective Bundle Adjustment (PBA) 1 and 2 (Section 2.1). The histograms (a) and (c) show the difference in residuals between Wiberg and simultaneous for PBA 1 and PBA 2, while (b) and (d) show the difference in number of iterations required for convergence to a strict threshold.

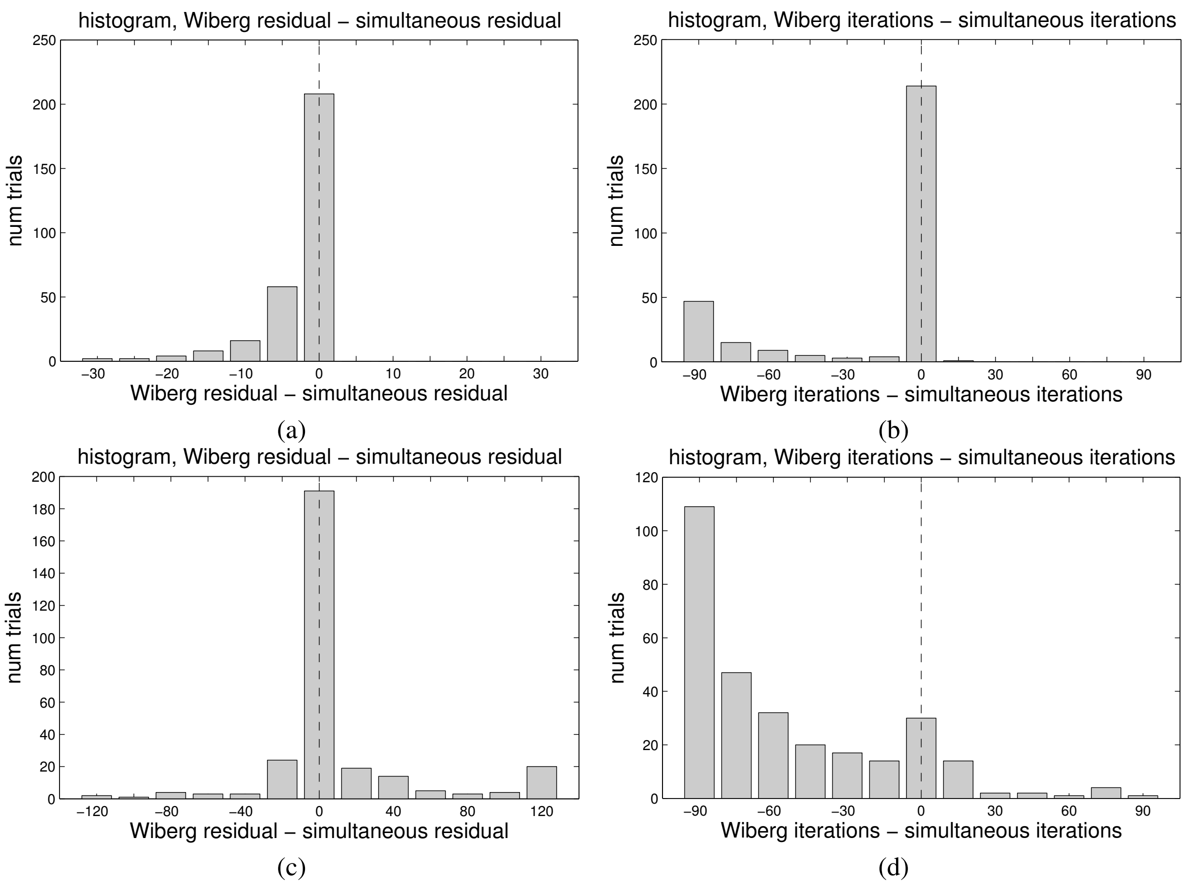


Figure 5. Results for Projective Bundle Adjustment (PBA) 3 and 4 (Section 2.1). The histograms (a) and (c) show the difference in residuals between Wiberg and simultaneous for PBA 3 and PBA 4, while (b) and (d) show the difference in number of iterations required for convergence to a strict threshold.

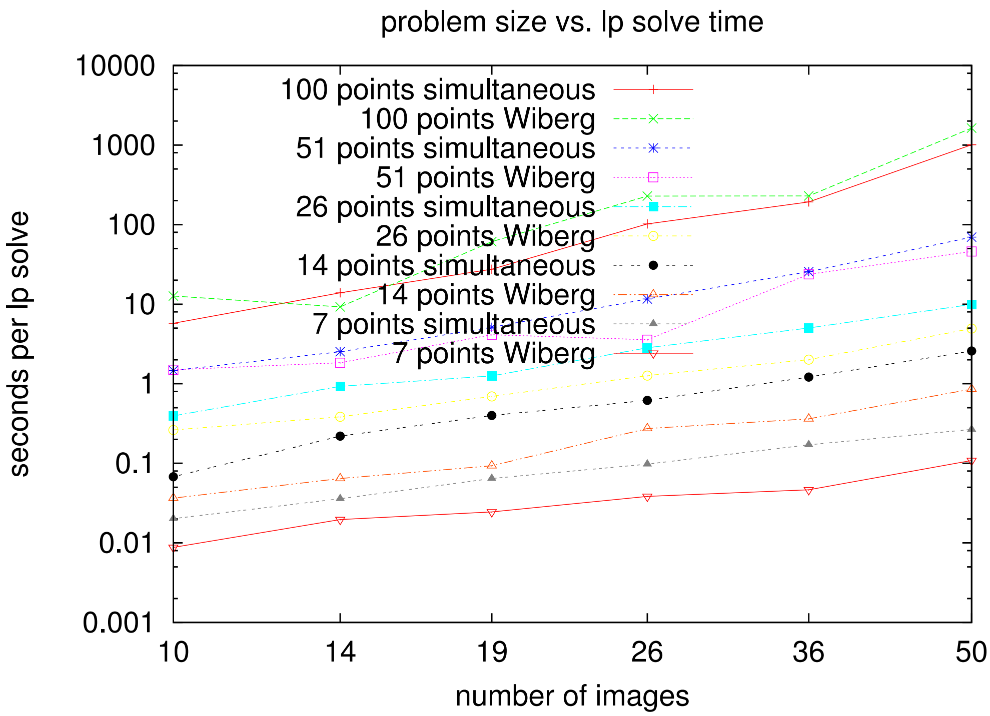


(a)

Figure 6. The average time per linear programming solve for Wiberg and simultaneous projective bundle adjustment, by problem size (Section 2.2). Wiberg solves are faster for problems with up to 51 points, and simultaneous solves are faster for problems with 100 points or more. Both axes are logarithmic.