Chapter 5

Experimental Results

5.1 ST-SLAM in a Simulated Large-Scale Outdoor Environment

5.1.1 Performance in Terms of Log Likelihood

Fig. shows the accumulated log likelihood of the best particle for each of the SLAM experiments with different numbers of particles. Along with results for each ST-SLAM experiment, we show, as a baseline for comparison, the accumulated log likelihood of odometry-only SLAM and true-move-only SLAM. The accumulated log likelihood begins improving with more than 10 particles. We need more than 100 particles to achieve better performance than odometry-only SLAM. We observe a saturation trend from 1,000 to 10,000 particles. However, there is still a large gap to reach the performance of true-move-only SLAM.

Fig. 5.2 shows estimated robot's trajectory for each different configuration of SLAM. The trajectories of odometry-only SLAM and ST-SLAM with 10 particles are similar. With 100 particles, the trajectory then becomes better than that of odometry-only SLAM. This transition is consistent with the transition of accumulated log likelihood passing the odometry-only SLAM's performance line. With 1,000 particles, the estimated trajectory is quite on a practical level. The trajectory of ST-SLAM with 10,000 particles resembles that of ST-SLAM with 1,000 particles in global structure, however the trajectory is locally more accurate with 10,000 particles. Even ST-SLAM with 10,000 particles could not estimate correctly the last straight path after turning the third corner. In this part, the robot trajectory veers to the right due to insufficient number of observed landmarks.

5.1.2 Performance in Terms of Move Error

In this section, we evaluate the performance of ST-SLAM by comparing estimated moves and measured true moves. Fig. shows the statistics of move error in terms of Euclidean xy-distance and yaw.

For xy-distance error, we do not observe any superior performance of ST-SLAM over odometry-only SLAM. Increasing the number of particles does not contribute to ST-SLAM's performance in terms of xy-distance. However, the average deviation from odometry-only SLAM is less than 4 cm for any ST-SLAM configuration, which does not cause serious impact on localization and mapping in this large-scale environment.

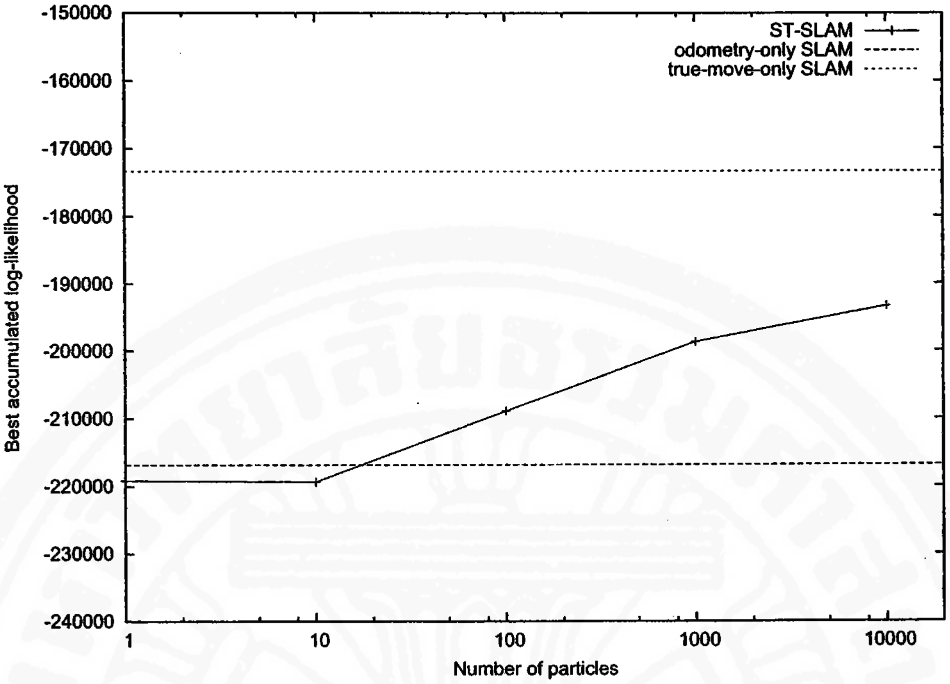

Figure 5.1 Best accumulated log likelihood versus number of particles.

On the other hand, move error in yaw significantly depends on the number of particles used in ST-SLAM. To compare the estimated yaw of odomety-only SLAM with that of ST-SLAM, we use paired t-tests on the mean differences for each move. The null hypothesis is that the mean of the differences is equal to zero. We performed two-tailed tests with a type I error rate of 0.05 comparing the results of odometry-only SLAM and ST-SLAM with each different number of particles. The tests revealed that ST-SLAM with 1 or 10 particles is significantly worse than odometry-only SLAM, ST-SLAM with 100 particles is not significantly different from odometry-only SLAM, and that ST-SLAM with 1,000 or 10,000 particles is significantly better than odometry-only SLAM. The result indicates that we need more than 100 particles for ST-SLAM to achieve better localization performance than odomety-only SLAM. This is consistent with the comparison results in terms of accumulated log likelihood and robot's trajectories that we previously performed.

5.1.3 Reconstructed 3D Point Map

We show the 3D point landmark map obtained with 10,000 particles in Fig. 5.4. The dense parts of the point cloud are points on walls, and corners of grid patterns of buildings in the fort.

5.2 ST-SLAM in a Real Indoor Environment

5.2.1 Performance in Terms of Log Likelihood

Fig. 5.5 shows the accumulated log likelihood of the best particle for each of the SLAM experiments with different numbers of particles and odometric noise levels. Along with results for each ST-SLAM experiment, we show, as a baseline for comparison, the accumulated log likelihood of odometry-only SLAM and true-move-only SLAM. For every configuration of ST-SLAM, the accumulated log likelihood improved as the number of particles increased, finally approaching or even exceeding the accumulated log likelihood of true-move-only SLAM.

For example, in the case where $\alpha = 2.0$, which is a typical degree of odometry error, we needed at least 100 particles to achieve better performance than odometry-only SLAM. With 10,000 particles, the accumulated log likelihood is comparable to that of true-move-only SLAM. This observation is consistent with the quality of the path estimates shown in Fig. 5.6.

In principle, true-move-only SLAM should show the best performance in terms of accumulated log likelihood, but it was not so in our experiments. This is mainly due to two factors: 1) the noise in the 3D point sensor observations making the optimum in terms of log likelihood slightly different from the true motion, and 2) small errors in our measurement of the ground truth motion. Note, however, that the log likelihood for ST-SLAM saturates near the accumulated log likelihood for true-move-only SLAM when the noise is not too large (e.g. when $\alpha < 4.0$).

5.2.2 Performance in Terms of Move Error

In this section, we evaluate the performance of ST-SLAM by comparing estimated moves and measured true moves. Fig. 5.7 and Fig. 5.8 show the mean error of the estimated moves in terms of xy-distance (parallel to the floor) and yaw, respectively, as a function of odometry error. In terms of xy error, ST-SLAM is not an improvement over odometry-only SLAM. However, the average deviation from odometry-only SLAM is less than 1.0 cm, so there is little impact on localization and mapping. On the other hand, the pattern for yaw errors is quite different.

To compare the estimated yaw of odomety-only SLAM with that of ST-SLAM, we use paired t-tests on the mean differences for each move. The null hypothesis is that the mean of the differences is equal to zero. First, we performed eight two-tailed tests (one for each odometric noise level α) with a type I error rate of 0.05 and Bonferroni correction comparing the results of odomety-only SLAM and ST-SLAM with 1,000 particles. The tests revealed that ST-SLAM is significantly better than odometry-only SLAM at $\alpha = 3.0$ and not significantly different from odometry-only SLAM otherwise. Second, we performed another eight tests comparing the results of odometry-only SLAM and ST-SLAM with 10,000 particles. The tests indicated that ST-SLAM is significantly worse than odometry-only SLAM at $\alpha = 0.2$, significantly better than odometry-only SLAM at $\alpha = 2.0$, 3.0 and 5.0, and not significantly different from odometry-only SLAM otherwise. It is important that ST-SLAM achieves superior performance in the range $2.0 \leq \alpha \leq 3.0$, since this range covers the most typical motion errors for indoor mobile robots.

Finally, we find that accurate yaw estimation is crucial for obtaining a high quality landmark map and path reconstruction. Fig. 5.6 shows the final estimated paths projected onto the floor of the lab for each particle set size when $\alpha = 2.0$. For particle set sizes less than 100, we observe catastrophic collapse of the estimated path due to poor estimates of yaw moves. With 1,000 or more particles, however, ST-SLAM succeeds in recovering the global structure of the rig's path (compared to the ground truth path shown in Fig. 4.3).

5.2.3 Reconstructed 3D Point Map

We show the 3D point landmark map obtained with 10,000 particles in Fig. 5.9. The dense parts of the point cloud are objects on the desks in the laboratory.

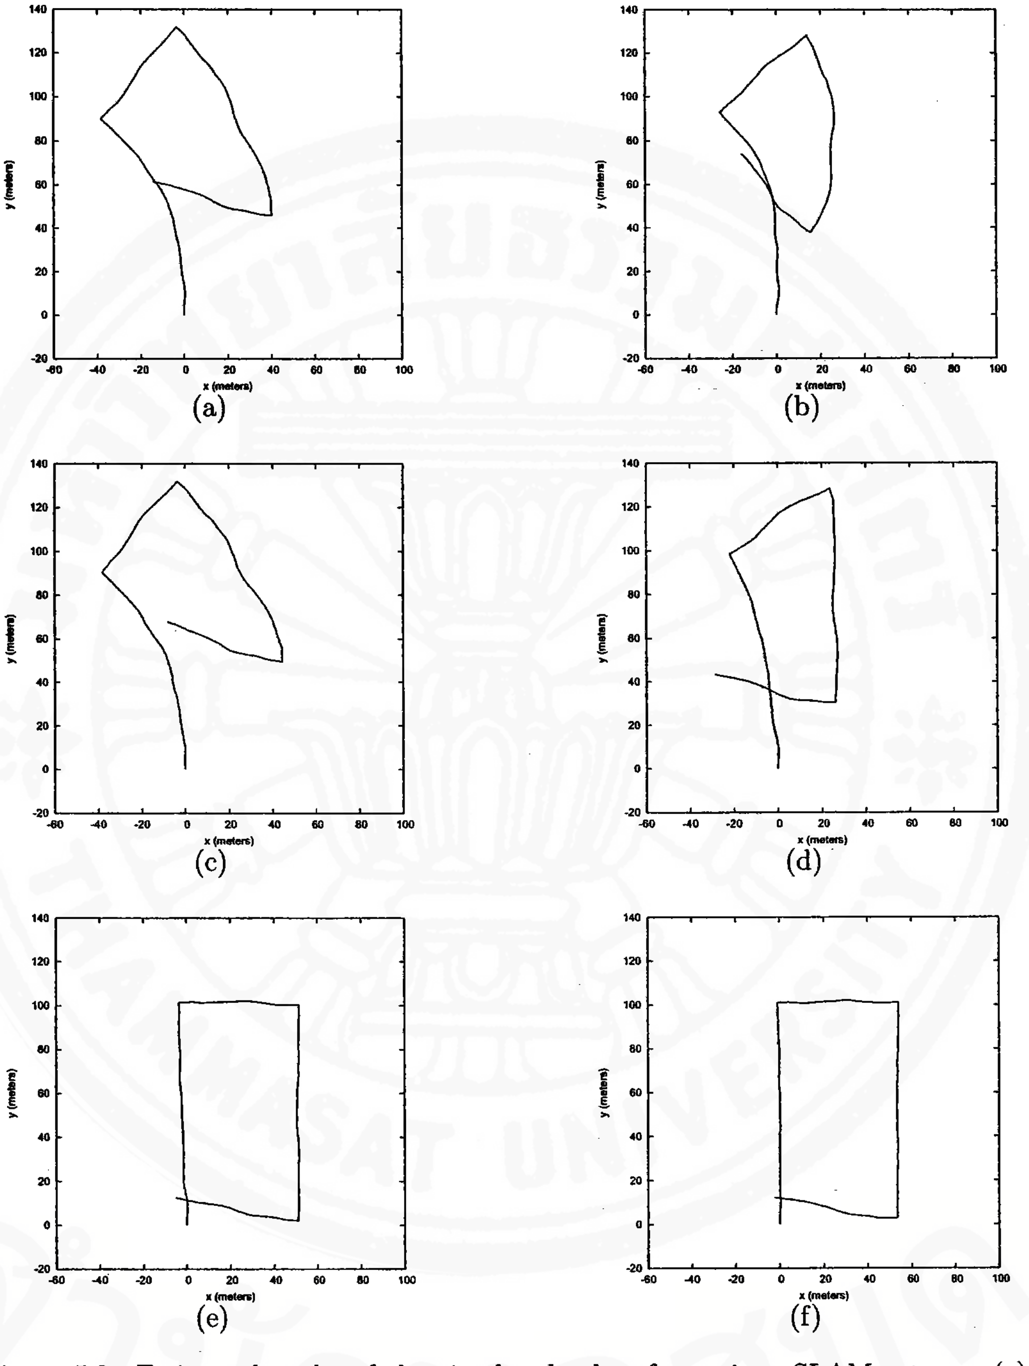

Figure 5.2: Estimated paths of the simulated robot for various SLAM setups. (a) Odometry-only SLAM. (b–f) ST-SLAM with 1, 10, 100, 1,000, and 10,000 particles, respectively.

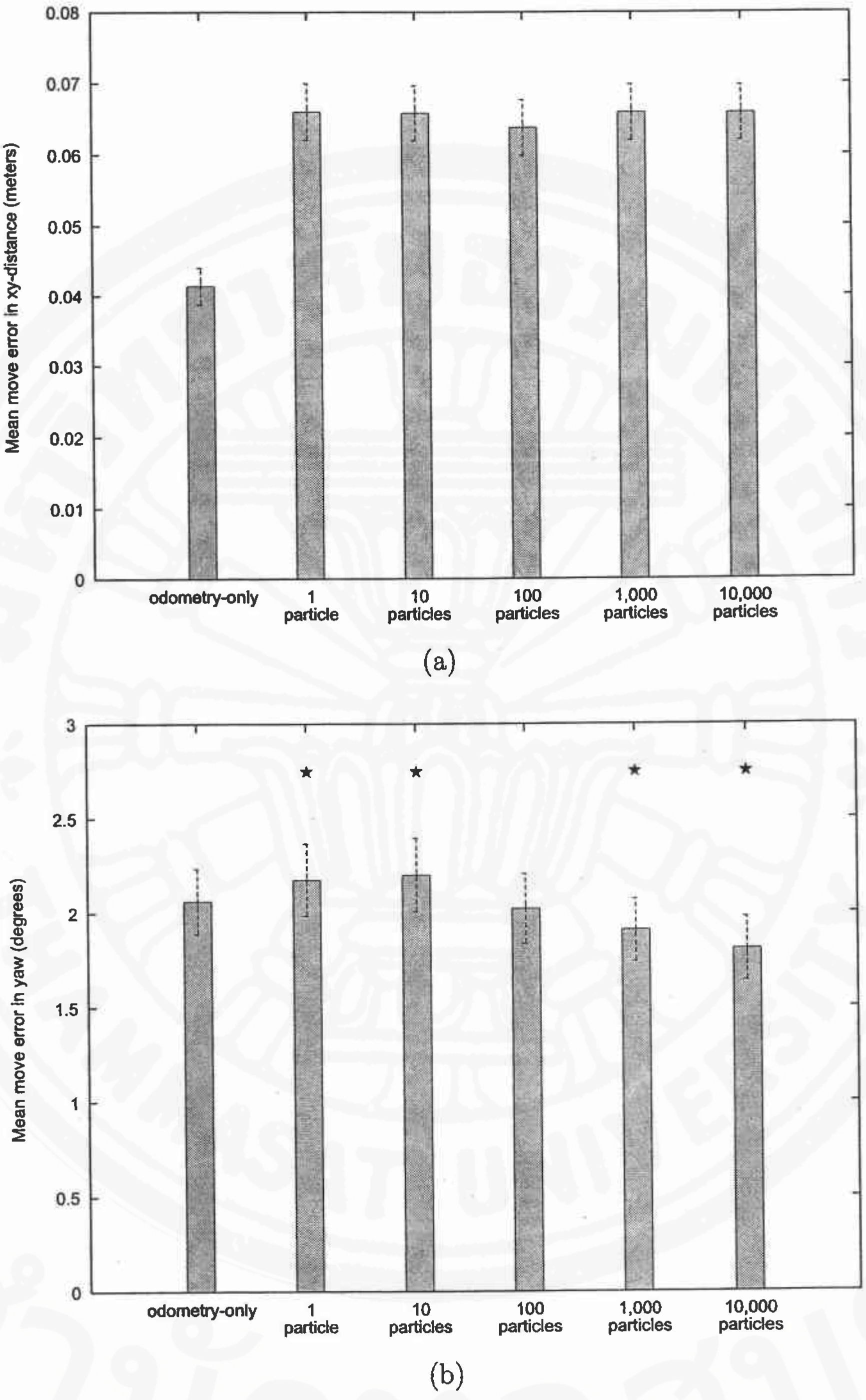

Figure 5.3: Move errors for various SLAM setups. Error bars are 95% confidence intervals on the mean error over the experiment. (a) xy-distance error. (b) Yaw error. $\star$: significantly different from odometry-only SLAM with a type I error rate of 0.05.

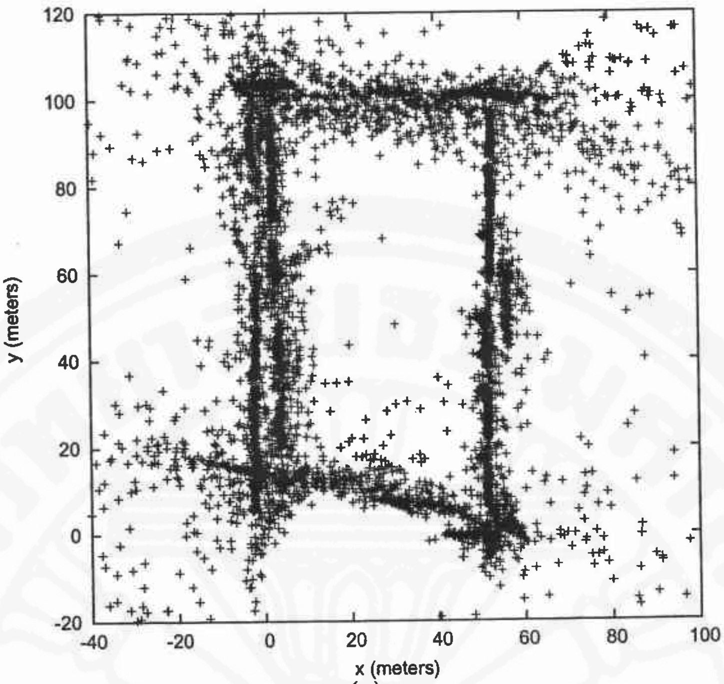

(a)

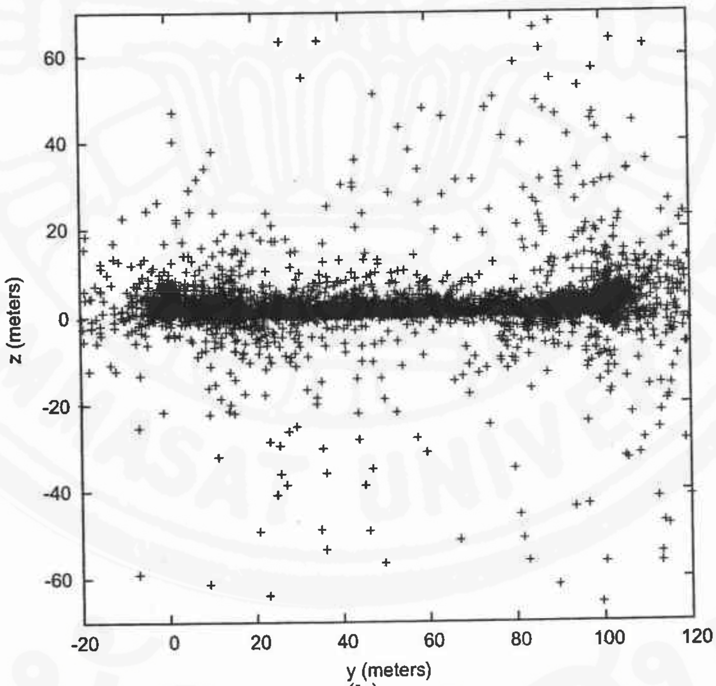

(b)

Figure 5.4: 2D projections of the 3D point landmark map of the simulation environment obtained by ST-SLAM with 10,000 particles. (a) Top view with the 3D point landmarks projected onto the xy-plane. (b) Side view with the 3D point landmarks projected into the yz-plane.

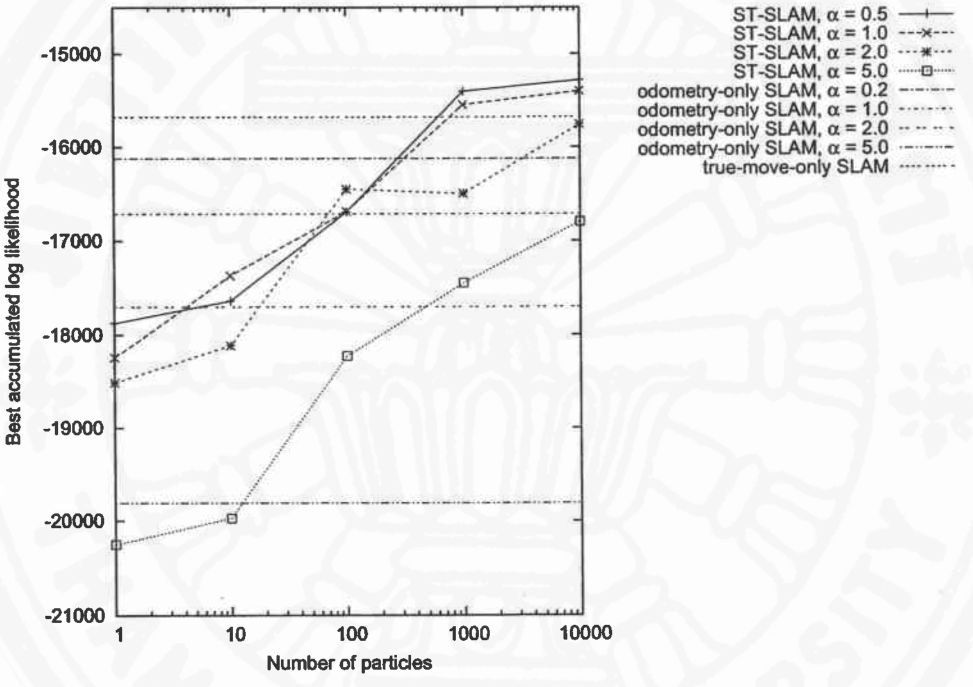

Figure 5.5: Best accumulated log likelihood versus number of particles. α is the odometric noise used to calculate the measured motions $\boldsymbol{u}_{1:t}$.

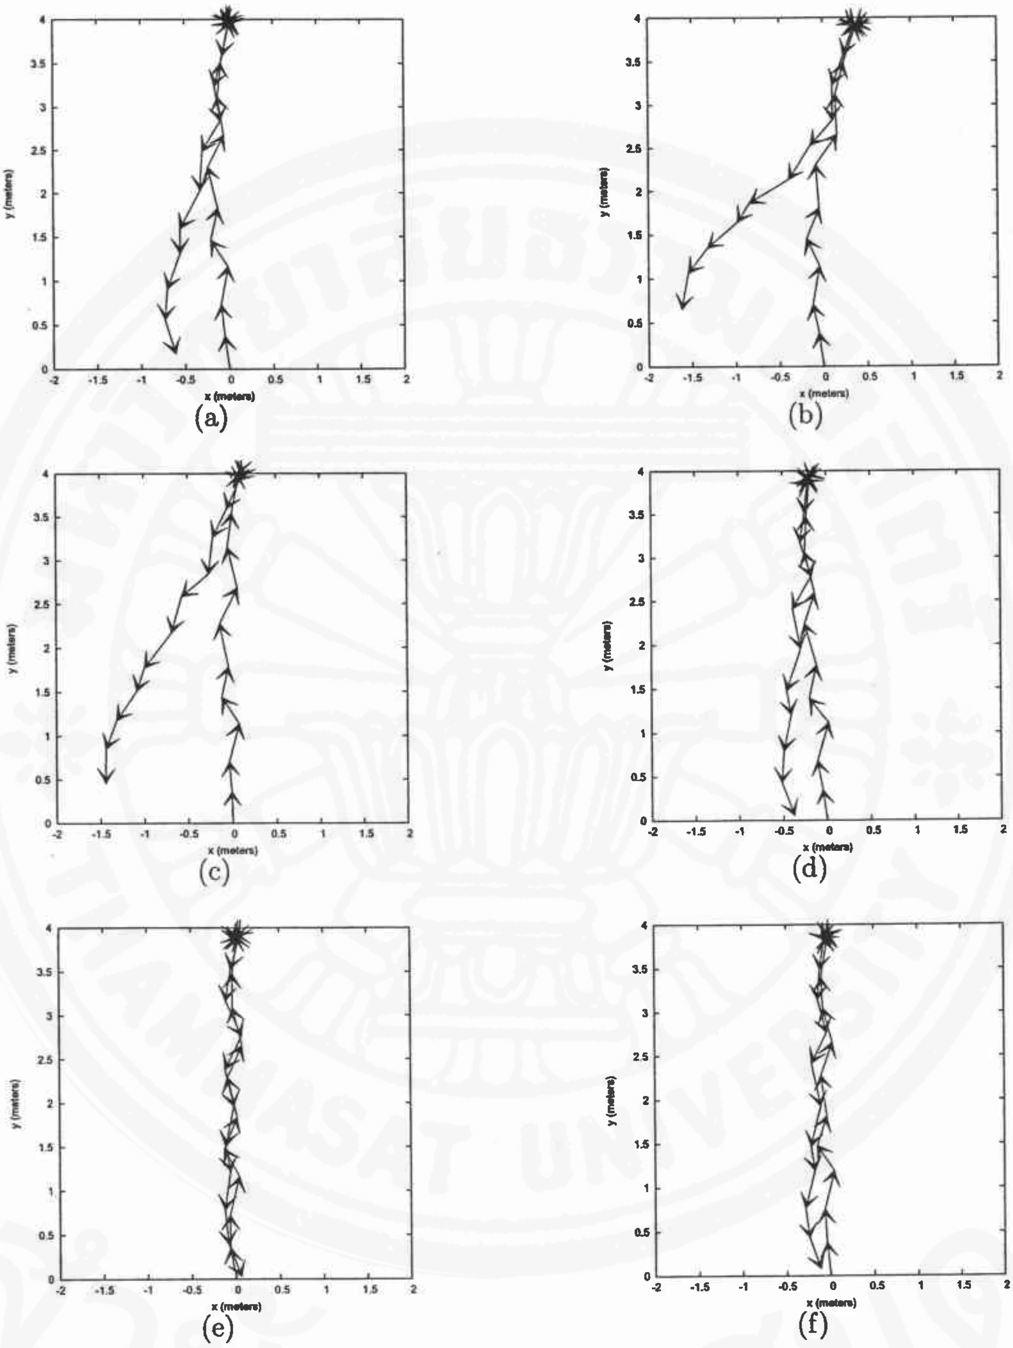

Figure 5.6: Estimated paths of the rig for various SLAM setups with odometric noise level $\alpha = 2.0$. (a) Odometry-only SLAM. (b–f) ST-SLAM with 1, 10, 100, 1,000, and 10,000 particles, respectively.

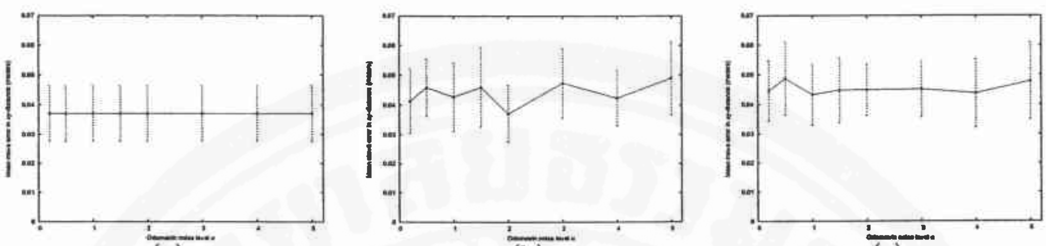

(a) (b) (c)

Figure 5.7: xy-distance move error versus odometric noise level α. (a) Odometry-only SLAM. (Since we only vary yaw error, α has no effect on the xy-distance move error.) (b) ST-SLAM with 1,000 particles. (c) ST-SLAM with 10,000 particles. Error bars are 95% confidence intervals on the mean error over the experiment.

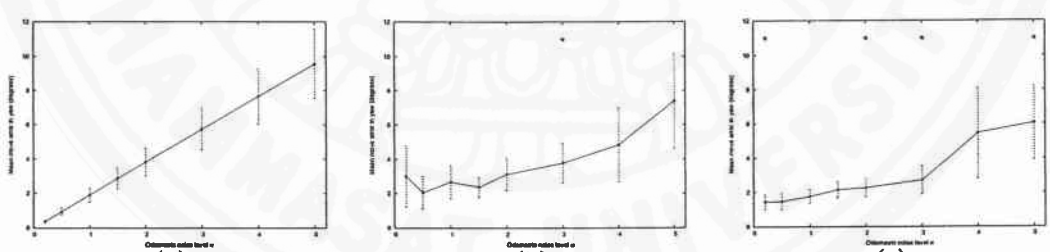

(a) (b) (c)

Figure 5.8: Yaw move error versus odometric noise level α. (a) Odometry-only SLAM. (b) ST-SLAM with 1,000 particles. (c) ST-SLAM with 10,000 particles. Error bars are 95% confidence intervals on the mean error over the experiment. $\star$: significantly different from odometry-only SLAM at corresponding α with a type I error rate of 0.05 and Bonferroni correction.

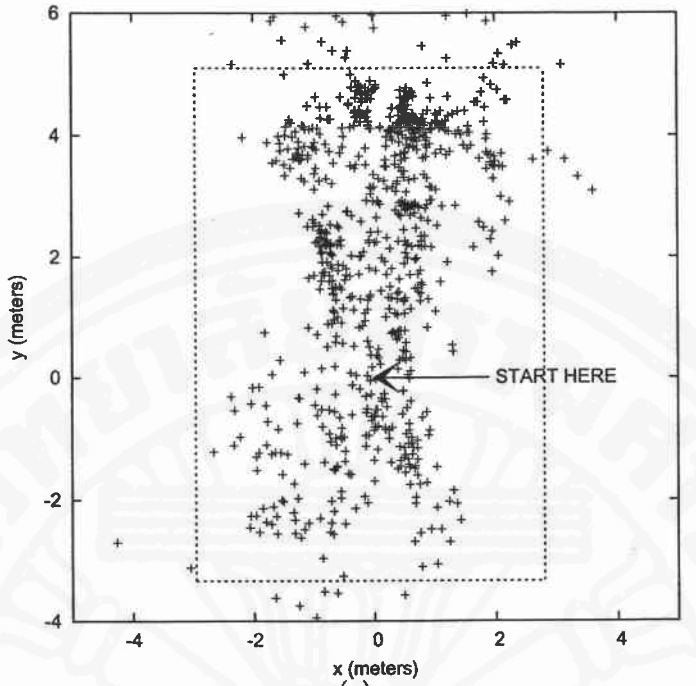

(a)

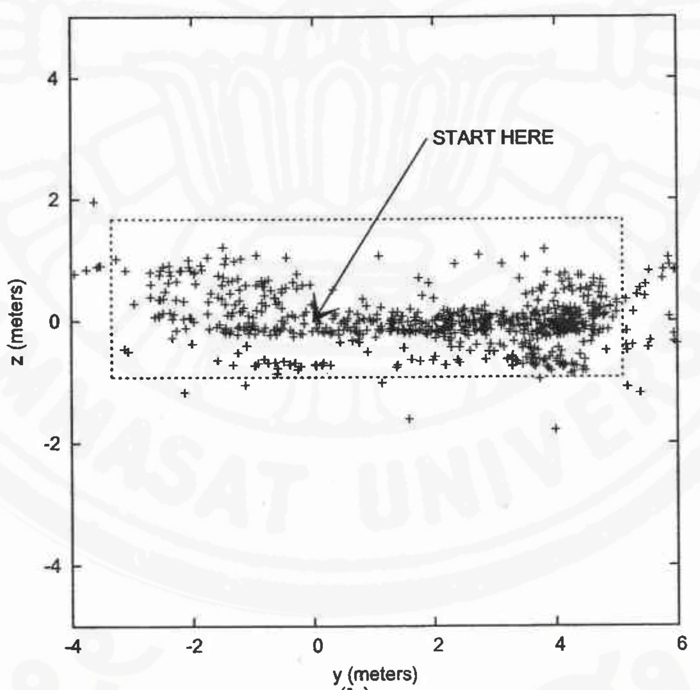

(b)

Figure 5.9: 2D projections of the 3D point landmark map obtained by ST-SLAM with 10,000 particles. Odometric noise level $\alpha = 2.0$. The dashed rectangle is the boundary of the lab. The initial position of the tripod is marked by a label. (a) Top view with the 3D point landmarks projected onto the xy-plane. (b) Side view with the 3D point landmarks projected into the yz-plane.