Published in Proc. of IEEE Intl. Conf. on Emerging Technologies and Factory Automation (ETFA), 2016. DOI: 10.1109/ETFA.2016.7733598

Robust Factor Graph Optimization – A Comparison for Sensor Fusion Applications

Tim Pfeifer, Peter Weissig, Sven Lange and Peter Protzel
Dept. of Electrical Engeneering and Information Technology
TU Chemnitz, Germany
Email: {firstname.lastname}@etit.tu-chemnitz.de

***Abstract*—While many applications of sensor fusion suffer from the occurrence of outliers, a broad range of outlier robust graph optimization techniques has been developed for simultaneous localization and mapping. In this paper we investigate the performance of some of the most advanced algorithms for a simulated wireless localization setting affected by non-Gaussian errors. With this first analysis we can show some of the advantages and disadvantages that are connected with the different concepts behind Max-Mixture, Generalized iSAM, Switchable Constraints and Dynamic Covariance Scaling.**

I. Introduction

While many sensor fusion applications would benefit from a robustness against outliers, a variety of robust graph optimization algorithms have been developed for back-ends of simultaneous localization and mapping (SLAM) systems [1]–[8]. Primary reason of the demand for this robustness in SLAM is the existence of false positive loop closures in almost every place recognition front-end. They are able to cause catastrophic failures, in typical applications like autonomous navigation, if non-robust least squares estimation is applied. The general approach to robustify the optimization process can be summarized as the rejection of outliers based on the assumption that the occurrence of wrong loop closures is unpredictable and arbitrary – but relatively rare. In [9] and [10] the authors compared a subset of available robust algorithms on different real and synthetic datasets, without declaring a clear winner.

Although, all of this research was focused on SLAM problems, the application of factor graphs as a tool for least squares optimization isn't limited to these. There is a broad field of classic sensor fusion problems that formerly was addressed by filter approaches and benefits from a formulation as graph optimization problem like shown in [11]. Even closer related to our work is [12] and [13] where robust back-ends for GPS based applications were proposed.

The foundation of earlier filter or graph based sensor fusion and SLAM was to assume a Gaussian distributed error. While wrong loop closures violate this assumption in SLAM, there are a couple of effects that can cause the corruption of physical sensor data: Ultrasonic and wireless ranging sensors suffer from reflections and wheel based odometry from the possibility to slip. However, these effects don't cause arbitrary or unbounded errors, but can rather be described with a distribution based on physical principles. Some of the robust graph optimization techniques, like Max-Mixture Models [7] or the generalized iSAM algorithm [8] are able to represent such non-Gaussian distributions.

The project is funded by the "Bundesministerium für Wirtschaft und Energie" (German Federal Ministry for Economic Affairs and Energy).

In our work, we will investigate in a first synthetic setting if algorithms that approximate the real error distribution can achieve better results than such that are developed to reject arbitrary outliers.

II. Robust Factor Graph Optimization

The non-linear optimization of Bayesian networks, embodied by factor graphs, is a general technique to find the Maximum A Posteriori estimate for a set of given observations. It involves the search for a state $\boldsymbol{X}^*$ that maximizes the probability $P(\boldsymbol{X}|\boldsymbol{Z})$, for given measurements $\boldsymbol{Z}$ using a non-linear least squares estimation:

$$\boldsymbol{X}^* = \underset{\boldsymbol{X}}{\operatorname{argmin}} \sum_i \| \underbrace{e(X_i, Z_i)}_{e_i} \|^2_{\Sigma_i} \tag{1}$$

Each of the sum's elements represents a single factor in this graph which can have a different non-linear error function $e(X_i, Z_i)$. To optimize (1), a batch solver like Gauss-Newton or an incremental one like provided by iSAM2 [14] can be used.

A. Reject Outliers vs. Expect Outliers

Based on the SLAM problem a recent approach to deal with measurements that violate the Gaussian assumption is to identify the non-Gaussian outliers within the back-end and exclude them from the optimization process. In general, these algorithms handle the classification of outliers and the following consequences differently. Realizing, Reversing, Recovering (RRR) [5] and the EM-Algorithm [6] take final decisions binary. Switchable Constraints (SC) [1] and Dynamic Covariance Scaling (DCS) [2] instead, give each measurement a continuous weight to express the affiliation to the class of outliers. Because binary decisions seems to be unsuited for distributions where the value range of outliers overlap with the Gaussian proportion, we restrict our first comparison to SC and DCS.

A similar way, but with different consequences is to describe the entity of measurements as one complete distribution which changes the estimation equation (1). While the true distribution is often unknown, a Gaussian Mixture model (GMM) can be utilized to describe an empirical approximation. We use Max-Mixture (MM) [7] and the generalized iSAM algorithm (GiSAM) [8] to implement the estimation equation for such a model.

B. Switchable Constraints

In [1] a set of switch variables $\boldsymbol{S}$ with i components s_i, representing the affiliation of each measurements Z_i to the group of non-outliers, is added to the state.

$$\boldsymbol{X}^*, \boldsymbol{S}^* = \underset{\boldsymbol{X},\boldsymbol{S}}{\operatorname{argmin}} \sum_i \|e_i \cdot s_i\|^2_{\Sigma_i} + \sum_i \| \underbrace{1 - s_i}_{e^{SP}} \|^2_{\Xi_i} \tag{2}$$

The estimation equation (2) includes an additional prior e^{SP} to prevent the trivial solution $s_i = 0$. Central idea of SC is to let $\boldsymbol{S}$ be part of the optimized values, so the optimizer can find the best possible choice of each switch variable. The only free parameter Ξ_i is the covariance of the switch prior.

C. Dynamic Covariance Scaling

DSC [2] was proposed as an enhancement of SC which share the same error function but offers the following closed form solution of s_i.

$$s_i = \min\left(1, \frac{2\Phi_i}{\Phi_i + \|e_i\|^2_{\Sigma_i}}\right) \text{ with } \Phi_i = \Xi_i^{-1} \tag{3}$$

With Φ_i there is also only one free parameter.

D. Maximum-Mixture

The authors of [7] proposed an algorithm to approximate the sum operator in a sum of Gaussians error model (4) with a maximum selector.

$$P(\boldsymbol{X}|\boldsymbol{Z}) = \prod_i \sum_n \frac{w_n}{\sqrt{2\Sigma_{in}^2}} \cdot \exp\left(-\frac{(e_i - \mu_n)^2}{2\Sigma_{in}}\right) \tag{4}$$

After applying the negative logarithm, the maximum operator becomes a minimum inside the least squares estimator:

$$\boldsymbol{X}^* = \underset{\boldsymbol{X}}{\operatorname{argmin}} \sum_i \min_n \left(\frac{1}{2}\|e_i - \mu_n\|^2_{\Sigma_{in}} - \ln\left(\frac{w_n}{\sqrt{2\Sigma_{in}^2}}\right)\right) \tag{5}$$

Instead of just one, MM involves $3 \times n$ free parameters. While Ξ_i and Φ_i are complete new artificial parameters that are hard to determine, μ_n, Σ_n and w_n might be easy to estimate for a given sensor with a static or predictable error distribution. Otherwise the distribution has to be computed on-line.

E. Generalized iSAM2

Generalized iSAM2 [8] is a general method to incorporate an arbitrary probability density function $P(X_i|Z_i)$ into the estimation equation as follows:

$$\boldsymbol{X}^* = \underset{\boldsymbol{X}}{\operatorname{argmin}} \sum_i \left\| \sqrt{-\ln \frac{P(X_i|Z_i)}{c_i}} \right\|^2 \tag{6}$$

Under the Condition:

$$c_i > \max P(X_i|Z_i) > 0 \text{ for all } X_i \tag{7}$$

Similar to MM, we use a sum of Gaussian but keep the sum operator for each factor (8). So we can achieve an exact solution instead of approximating it with a maximum selector.

$$P(X_i|Z_i) = \sum_n \frac{w_n}{\sqrt{2\Sigma_{in}^2}} \cdot \exp\left(-\frac{(e_i - \mu_n)^2}{2\Sigma_{in}^2}\right) \tag{8}$$

Based on (8) the parameter c_i has to satisfy the condition:

$$c_i > \sum_n \frac{w_n}{\sqrt{2\Sigma_{in}^2}} \tag{9}$$

The required parameters are identical to Max-Mixture.

III. Experiments

Our experimental set-up mimics a wireless localization system which suffers from non-Gaussian errors. Multipath effects through reflections are a typical physical phenomenon that is causes such outliers. This occurs with GPS in urban canyons as well as with wireless indoor localisation in cluttered environments and lead to large distortions in classic Gaussian estimators. Therefore, we define a scenario which is based on the same physical principles to evaluate the performance of the different algorithms under artificial conditions. Along with SC, DCS, MM and generalized iSAM we use a non-robust factor graph as baseline of our Benchmark. All experiments were performed with the GTSAM framework.

A. Dataset

Our experiment is based on the Manhattan-World dataset [15]. We decided to use synthetic instead of real GPS measurements to focus our work on the multi-modal multipath error only. Therefore, we will have full ground truth throughout all simulations. Since the dataset contains a factor graph based of

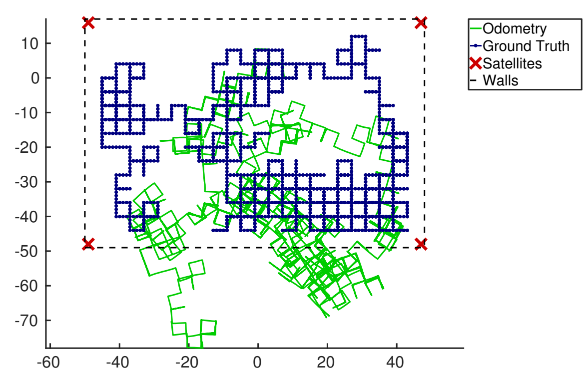

Fig. 1: Test data based on Manhattan-World dataset ”m3500”. The ground truth path (blue line with circle markers) and the virtual satellites (red crosses) are used to calculate a set of four range measurements for each position. Afterwards Gaussian noise is added. For a specific proportion of measurements the calculated distance represent a reflection on the virtual walls (dotted line) instead of the the shortest connection.

odometry and loop-closures, we transformed it to a ranging scenario to fit our needs. We started with the bare odometry and created ground truth data knowing that all coordinates are integers. Both trajectories can be seen in figure 1 together with the virtual satellites. The virtual walls around the trajectory are used to generate non-line-of-sight (NLOS) measurements.

For each determined NLOS ranging the distance is calculated using the shortest possible reflection on the outer walls.

With all this, we now create range measurements from each satellite to each ground truth pose. To vary the error distribution we are able to change the amount of Gaussian noise and the likelihood of reflected measurements independently.

B. Error metric and parametrization

Our comparative metric is the root mean square error (RSME) of all 2D positions which is similar to the ATE [16] that is used in both SLAM related comparisons [9] and [10]. To evaluate the robustness of each algorithm, we perform several simulations with increasing NLOS ratios from 0 up to 50 percent of all measurements. This exceed the outlier proportion of former publications by far, although the impact of each is significantly smaller due the absolute character of position measuring. To minimize the influence of a specific random initialization on the error distribution we reiterated each experiment 20 times with different random seeds and calculated our error metrics over all runs.

Through the perfect ground truth, the error distribution can be precalculated and approximated with a GMM like shown in (2). The mixture model's parameters can be applied to Max-Mixture as well as generalized iSAM. Despite DCS and SC work with their only parameter fixed, they also benefit from the off-line GMM estimation through their Gaussian component which is set to the LOS part of the GMM. The non-robust graph uses the covariance of a fitted 1-component Gaussian model for optimal performance.

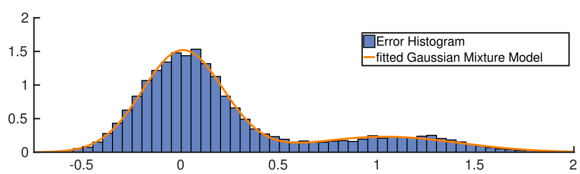

Fig. 2: The resulting error distribution (blue) of a run with 25% NLOS measurements can be easily fitted with a 2-component Gaussian mixture model (orange).

The GTSAM implementation of iSAM2 which we use, provides a Gauss-Newton optimizer as well as the Dog-Leg trust region method. According to our observation the choice of the optimizer can considerably impact the results, so we decided to apply the optimizer with the best RSME results independently for each algorithm. Therefore, all algorithms except generalized iSAM use Gauss-Newton instead of Dog-Leg. This irregularity might be caused by the Jacobian (see Figure 3) that changes its sign and crosses zero which is caused by the exact GMM's implementation. While Gauss Newton gets stuck when the Jacobian comes close to zero, the approximated quadratic term of Dog-Leg can still improve the estimation.

IV. Results

Figure 4 summarizes the results after 20 runs with our simulated localization problem for each NLOS ratio. The random error distributions are identical for each algorithm

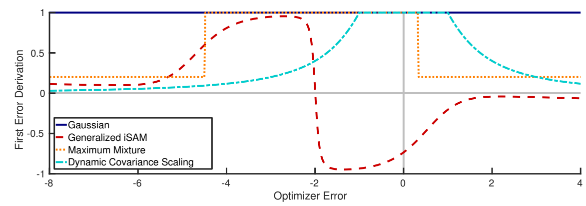

Fig. 3: The Jacobians of the compared algorithms for an one-dimensional example. The GiSAM implementation with a GMM leads to first error derivation with a changing sign. SC's Jacobian can not be plotted due its two-dimensional structure.

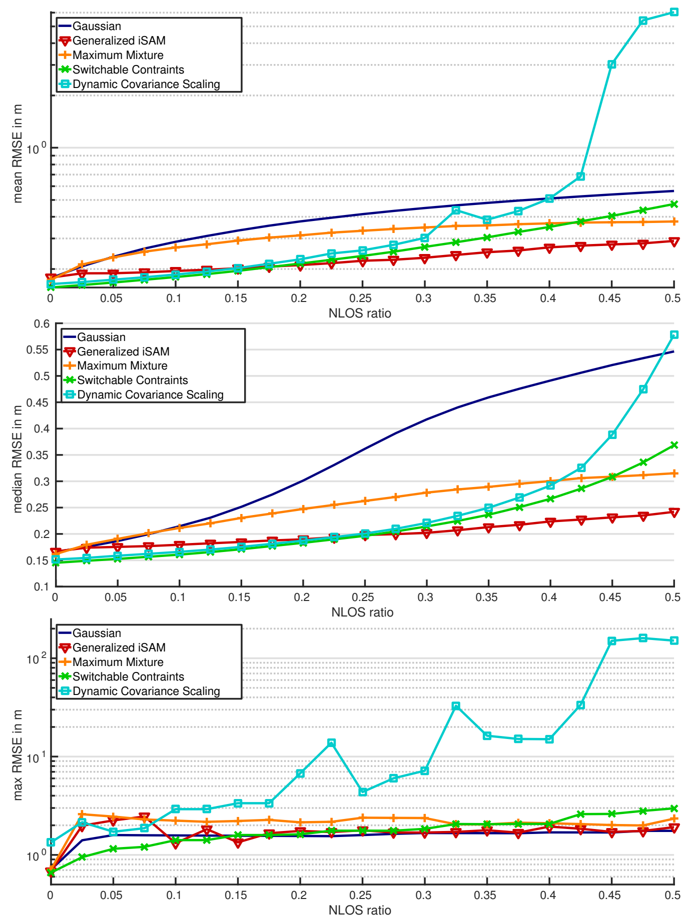

Fig. 4: Comparison of the resulting mean/median/max RSME for the tested algorithms with different NLOS ratios. Due to the huge range of values, mean and max error are scaled logarithmically. Except for the results of DCS all outperform the non-robust solution over the whole parameter set. Altogether generalized iSAM shows the best performance especially for higher ratios of reflected measurements.

but were recalculated every iteration. Although we show the median and maximum RSME too, the mean error is our main performance criterion.

A. Switchable Constraints

For smaller NLOS ratios up to 25%, SC keeps the RSME quite low and provides the best median performance together with DCS. Additionally the maximum error in this range is the smallest overall. Even in the run with 50% outliers it decrease the error significantly compared to the non robust solution.

This is remarkable if you remember that SC is based on the decision if a measurement is a outlier or not, which is difficult under this conditions.

B. Dynamic Covariance Scaling

While achieving a low mean RMSE in low NLOS conditions close to SC, the algorithm fails at higher ratios in such a drastic way it raises doubts that it's identical to SC. A deeper look to the results supports the idea, that the inconsistent performance could be caused by local minima in which the optimizer gets stuck. The strong down weighting of greater errors causes a very fast descending first derivation which inhibits the optimizer in his progress. Moreover, in our experiment it tends to be not able to recover from this minima during further iterations. Especially for high NLOS ratios, the time consuming optimization process for the additional variables in SC seems to prevent this phenomenon. But for final conclusions further work on this topic is required.

C. Maximum-Mixture

In its coarse trend the Max-Mixture approach resembles the non robust Gaussian factor graph with a much smaller influence of the ascending NLOS proportion. Although it can't improve the performance for small ratios, the robustness for high ones around 50% exceed even SC. The maximum error however is slightly increased compared to the Gaussian baseline but the effect to the mean RSME seems bounded.

D. Generalized iSAM2

The GiSAM algorithm with a mathematically exact implementation of a GMM appears to be the best in this scenario. Generalized iSAM can't reach results of SC and DCS at NLOS ratios between 0% and 10% but shows remarkable low error values at higher ratios. Also the maximum RSME is close to the Gaussian factor graph.

V. Conclusion

SC, MM and GiSAM have shown their capability to improve factor graph based sensor fusion with non-Gaussian errors. SC with rejecting as well as MM/GiSAM with modelling outliers are able to decrease the RMSE compared to the non-robust graph. While the rejection works better at small outlier ratios the modelling does it at high. Generalized iSAM provides the best total performance but has an awfully complex Jacobian which works only with the Dog-Leg optimizer. Our further research will try to improve this through the integration of other mathematical models for error distributions.

DCS on the other hand seems to have serious issues with high NLOS ratios. A high amount of outliers that are bounded their range is nothing that occur in typical SLAM scenarios, so it's possible that previous comparisons like [10] doesn't noticed this difference to SC. In addition the structure of range only measurements leads to probability distributions in shape of intersecting circles with multiple local minima in case of DCS. Based on the results of our experiment the algorithm seems to be not suitable for this kind of problem.

Also, we have to notice that the maximum error of the most robust algorithms is higher than the one of the non-robust factor graph. This effect occurs at all outlier ratios and could be a problem for critical applications where a single wrong estimation can cause damage to the system or its environment. The recording of a real dataset is another next step to verify the characteristics of our simulation and this first results.

Finally it should be mentioned that SC and DCS were tested with fixed parameters and MM and GiSAM require a estimation of the underlying distribution. To combine an on-line distribution estimation with a factor graph would be another important improvement. The first step in this direction is already done with [13]. However, their work addresses only unimodal zero-mean distributions with a huge amount of off-line learning.

References

[1] N. Sünderhauf, "Robust optimization for simultaneous localization and mapping," Ph.D. dissertation, Technischen Universität Chemnitz, 2012.

[2] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard, "Robust map optimization using dynamic covariance scaling," in *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 62–69.

[3] P. Agarwal, G. Grisetti, G. Diego Tipaldi, L. Spinello, W. Burgard, and C. Stachniss, "Experimental analysis of dynamic covariance scaling for robust map optimization under bad initial estimates," in *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, 2014, pp. 3626–3631.

[4] P. Agarwal, "Robust graph-based localization and mapping," Ph.D. dissertation, PhD thesis, University of Freiburg, Germany, 2015.

[5] Y. Latif, C. Cadena, and J. Neira, "Robust loop closing over time for pose graph slam," *The International Journal of Robotics Research*, vol. 32, no. 14, 2013.

[6] G. H. Lee, F. Fraundorfer, and M. Pollefeys, "Robust pose-graph loop-closures with expectation-maximization," in *Proc. of Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 556–563.

[7] E. Olson and P. Agarwal, "Inference on networks of mixtures for robust robot mapping," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 826–840, 2013.

[8] D. M. Rosen, M. Kaess, and J. J. Leonard, "Robust incremental online inference over sparse factor graphs: Beyond the gaussian case," in *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 1025–1032.

[9] N. Sünderhauf and P. Protzel, "Switchable constraints vs. max-mixture models vs. rrr a comparison of three approaches to robust pose graph slam," in *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 5198–5203.

[10] Y. Latif, C. Cadena, and J. Neira, "Robust graph SLAM back-ends: A comparative analysis," in *Proc. of Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2014, pp. 2683–2690.

[11] S. Lange, N. Sünderhauf, and P. Protzel, "Incremental smoothing vs. filtering for sensor fusion on an indoor UAV," in *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 1773–1778.

[12] N. Sünderhauf, M. Obst, S. Lange, G. Wanielik, and P. Protzel, "Switchable constraints and incremental smoothing for online mitigation of non-line-of-sight and multipath effects," in *Proc. of Intelligent Vehicles Symposium (IV)*, 2013.

[13] R. Morton and E. Olson, "Robust sensor characterization via max-mixture models: Gps sensors," in *Proc. of Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 528–533.

[14] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree," *Intl. Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.

[15] E. Olson, J. Leonard, and S. Teller, "Fast iterative alignment of pose graphs with poor initial estimates," in *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, 2006, pp. 2262–2269.

[16] A. Bonarini, W. Burgard, G. Fontana, M. Matteucci, D. G. Sorrenti, and J. D. Tardos, "RAWSEEDS: Robotics Advancement through Web-publishing of Sensorial and Elaborated Extensive Data Sets," in *Proc. of IROS'06 Workshop on Benchmarks in Robotics Research*, 2006.