

IPC: Incremental Probabilistic Consensus-based Consistent Set Maximization for SLAM Backends

Emilio Olivastri and Alberto Pretto

Abstract—In SLAM (Simultaneous localization and mapping) problems, Pose Graph Optimization (PGO) is a technique to refine an initial estimate of a set of poses (positions and orientations) from a set of pairwise relative measurements. The optimization procedure can be negatively affected even by a single outlier measurement, with possible catastrophic and meaningless results. Although recent works on robust optimization aim to mitigate the presence of outlier measurements, robust solutions capable of handling large numbers of outliers are yet to come. This paper presents IPC, acronym for Incremental Probabilistic Consensus, a method that approximates the solution to the combinatorial problem of finding the maximally consistent set of measurements in an incremental fashion. It evaluates the consistency of each loop closure measurement through a consensus-based procedure, possibly applied to a subset of the global problem, where all previously integrated inlier measurements have veto power. We evaluated IPC on standard benchmarks against several state-of-the-art methods. Although it is simple and relatively easy to implement, IPC competes with or outperforms the other tested methods in handling outliers while providing online performances. We release with this paper an open-source implementation of the proposed method.

I. INTRODUCTION

A reliable, robust, and scalable solution to the SLAM problem is universally recognized as an enabling technology to make a robot truly autonomous. Although it is a problem that has been widely addressed by the research community and industry for more than 30 years, it cannot yet be considered a solved problem. For example, the design of fail-safe, self-tuning SLAM systems is still an open challenge, with many aspects still unexplored [1]. In recent years, the research community converged on modeling the SLAM problem through a graphical representation, in which nodes represent the poses of the robot, while the edges encode the probabilistic spatial relations between nodes. These constraints are computed from estimating the robot's motion, or by observing the environment and are affected by some degree of noise. Pose Graph Optimization (PGO) leverages the graphical representation employed in SLAM to frame and solve the non-convex maximum likelihood estimation (MAP) problem. The solution corresponds to the nodes' configuration that maximizes consistency between the constraints that define the problem. Libraries, such as g2o [2], GTSAM [3], and Ceres [4], are commonly used to solve the PGO problem.

To build a pose graph, it is required a *front-end* module that generates the spatial constraints. In this work, we are

All the authors are with the Department of Information Engineering (DEI), University of Padova, Italy. Emails: [olivastrie; alberto.pretto]@dei.unipd.it.

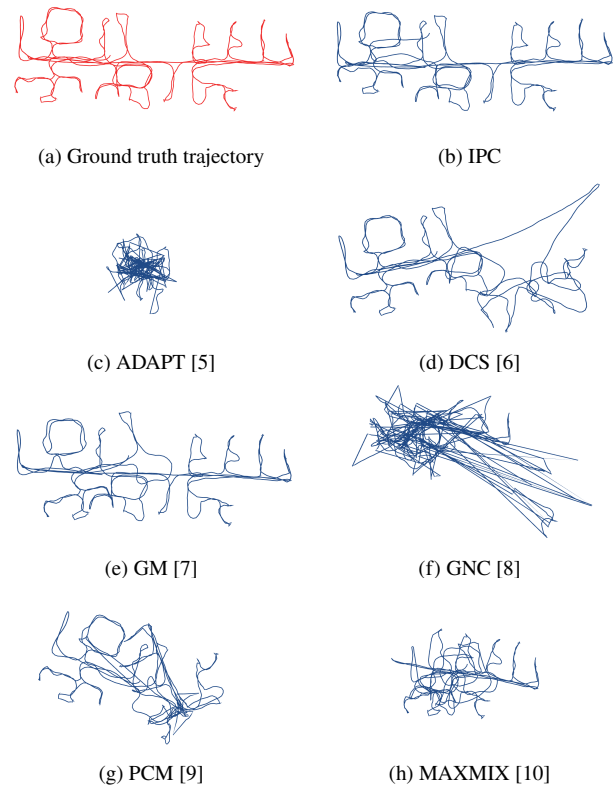


Fig. 1. Trajectory estimated by the different methods on the FR079 dataset corrupted with 70% of outliers.

going to focus on two primary types of constraints: odometry constraints, which relate together consecutive poses, and loop closure constraints, which relate together non-consecutive poses. Most approaches assume these constraints are affected by noise but do not contain outliers. Due to the inherent perceptual aliasing in the environment, these assumptions hold only when exploiting very conservative loop closure detectors, aiming at maximizing the probability of integrating a few inlier measurements at the cost of discarding, along with all outliers, most of the inliers. These behaviors are necessary because, in standard optimization methods, even a single outlier can completely distort the final results. Nevertheless, it would be unrealistic to expect the front-end to never make incorrect assumptions based on isolated observations, given the environment's aliasing effects. Therefore, it should be the back-end's responsibility to handle erroneous measurements, due to its ability in accessing the full exploration history. Recent works on graph-based SLAM proposed stacking methods that are able to handle large amounts of outliers, such as RANSAC [11], GNC [8], Consistent Set Maxi-

mization [9] [12] [13], and M-estimators [5], on top of classic solvers. Our contribution falls in the Consistent Set Maximization category. The proposed approach tries to approximate the solution to the combinatorial problem of finding the maximally consistent set of measurements in an incremental fashion. While the majority of solutions are designed for offline usage, our approach is able to achieve online performances, making it suitable for use during missions rather than solely in post-processing. For each incoming loop closure measurement, we find and optimize only a subgraph for which this constraint has an active influence. The solution to this subproblem is tested by using previously accepted measurements, each one with veto power, to decide whether to accept the new measurement. If accepted, also the subgraph optimization is propagated to the rest of the graph.

The main contributions of this work are the following:

- (i) IPC: An online, consensus-based approach able to incrementally build a good approximation of the maximally consistent set of loop closure measurements;
- (ii) An exhaustive evaluation of the proposed method. The evaluation is performed in standard datasets used for benchmarking state-of-the-art outlier rejection solutions;
- (iii) An open-source implementation of IPC available with this paper at:
<https://github.com/EmilioOlivastri/IPC>

The remainder of this paper is organized as follows:

Sec. II reviews related work on robust optimization and outliers rejection. In Sec. III IPC and the building blocks from which it was defined are presented. In Sec. IV we show a comprehensive evaluation of the current method and we compare it against state-of-the-art solutions on commonly used datasets. Finally, in Sec. V we draw our conclusions.

II. RELATED WORK

It is sufficient to have small percentages of outliers in the optimization process to invalidate the results completely. Robust pose graph optimization tries to mitigate the effect of such elements on the estimation process. M-estimators [7] are one of the most classic techniques used in robust estimation. It involves the application of robust cost functions, such as Huber, Cauchy loss, German-McCloure, and Truncated-Least-Squares, to reduce the weight of the outliers outside of a "trust region", that defines the maximum accepted error. Once, the M-estimator is chosen, it is needed to find an effective threshold for the trusted area, which is a non-trivial task. In order to bypass the problem of selecting the most suitable robust cost function, in [5] an adaptive robust loss function (ADAPT) has been presented. Such loss is a generalization that is able to represent all the different robust cost functions using a control parameter. It is able to automatically determine which one is the most suitable for a particular instance of the problem. Another limitation that these approaches share is the introduction of extra non-linearity, which will likely introduce more local minima to the solution space.

To mitigate the impact of additional non-convexity, approaches

such as Graduated Non-Convexity (GNC) [8] have been proposed. This method involves a stepwise process of progressively introducing non-convexity to the convex approximation of the initial problem, ultimately returning it to its original formulation. When combined with TLS, GNC enables the optimizer to filter out outliers more efficiently compared to a direct application of the M-Estimator. The major drawback is the increase in convergence time. In [14], thanks to the definition of the Scale Invariant Graduated(SIG) Kernel, the authors were able to adapt the GNC paradigm to an incremental solver such as iSAM2 [15], reducing dramatically the number of iterations needed for convergence. Another widely adopted approach is to let the optimizer directly handle the outliers detection by solving an augmented version of the original problem. For example, in [16] latent variables that are able to activate and deactivate measurements are added to the original formulation of the problem. In [6], the authors suggest an equivalent formulation that removes the additional "switchable" variables, reducing the convergence time of the method with respect to [16]. In [10] the presence of outliers was modeled using a multi-modal Gaussian distribution. In particular, the authors proposed a max-mixture distribution to enable the optimizer to select the most promising Gaussian at each iteration.

Another line of research, under the assumption that inlier measurements are jointly consistent, tries to find the set of maximum size of consistent measurements instead of explicitly trying to detect outliers. Computing the maximally consistent set poses a challenge due to its combinatorial complexity, driving the research community to explore effective approximations for its resolution. Joint Compatibility Branch and Bound (JCBB) [13] falls under the category of exact methods that are able to do better than exponential time, but are not viable for either real-time or online applications. In [17] the set of measurements is temporally divided into clusters, that are tested for both intra-consistency and inter-consistency. They approximate the maximum consensus set to all those measurements that passed both tests. In [18] the consistency check was improved by adding a term that was able to mitigate the effect of noisy odometry measurements at the cost of a higher convergence time.

Pairwise Consistency Maximization (PCM) [9] approximates joint consistency with pairwise consistency and by using approximate Max-Clique algorithms to compute the maximum pairwise consistent set. In ROBIN [19] the more general concept of invariance is introduced. It is used to define a new compatibility test from which to build the compatibility graph. Still, to establish effective invariants, it is essential to exploit priors that are inherent to the specific problem being addressed. In Clipper [12] the authors focused on defining different geometric-based invariants for solving diverse registration problems. The method is able to detect the densest compatibility graph while maintaining a relatively low execution time through the projection gradient ascent with a back-tracking line search.

III. METHODOLOGY

We expect the robot to be equipped with (i) an ego-motion estimation front-end (e.g., wheel odometry or visual-inertial odometry) that provides over time the rigid transformations between consecutive poses $\mathbf{x}_a, \mathbf{x}_{a+1} \in SE(2)$ (or $SE(3)$) traversed by the robot; (ii) a place recognition module that can detect medium- and long-term previously seen locations and estimate the transformation between the related previous pose \mathbf{x}_a and the current one \mathbf{x}_b (i.e., a *loop closure measurement* or *constraint*). Given the sequence of odometry and loop closure measurements, our method aims to find the maximum consistent set of measurements while estimating the whole trajectory followed by the robot.

A. Assumptions and Heuristics

Our method builds upon the following assumptions:

A1: Odometry measurements are inliers. The quality of the odometry can change based on the front-end that is used, but as stated in [20], they are usually not corrupted by outliers. These measurements are reliable constraints that can be used to enforce maintaining the local shape of the trajectory.

A2: Prevalence of inliers. Given the inherent aliasing effect in the environment, it is normal for the place recognition module to produce a certain percentage of outliers. On the other hand, it is reasonable to expect that the majority of loop closure measurements are inliers. Furthermore considering **A1** and the fact that the frequency of odometry measurements is generally higher than the one of loop closure measurements, the prevalence of inlier measurements is guaranteed.

As generally done in robust optimization, in IPC we apply the following heuristics:

H1: Start from a simpler version of the problem, for which its solution will be meaningful for the more complex version of the problem. As stated in [21] and later in [22], a simpler version of a problem exhibits a broader and more stable convergence basin, thus ensuring better convergence. Consequently, the solution obtained for the reduced problem can be used as the initial guess for the full problem, also improving the convergence of the latter.

H2: The nature of the compatibility test is to reject as many outliers as possible while preserving the inliers. The distinction between outliers and inliers is not always observable as stated in [20], so instead of trying to detect which measurements are outliers, consistency can be verified using Eq. (4).

B. Problem Formulation

A pose graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ uses nodes in \mathcal{V} to represent the discretized poses \mathbf{x}_i of the robot (i.e., the followed trajectory) while the edges $(a, b) \in \mathcal{E}$ that connect the nodes represent the spatial relations (i.e., the measurements, or constraints) between them. The full state can be represented using the vector $\mathbf{x} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top]^\top$, where the subscript defines the temporal order and n is the total number of poses. The goal of PGO is to find a solution \mathbf{x}^* that minimizes the following




Fig. 2. An example of pose graph with nodes represented by light blue circles. Each error term is associated with the corresponding edge. The coordinate system is usually fixed in the first node \mathbf{x}_0 . Light green and purple ovals represent examples of simple independent subgraphs (i.e., they only include one loop) while the dotted red oval represents more complex independent subgraphs, that include both multiple loops with crossing edges and internal loops.

objective function:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{(a,b) \in \mathcal{E}} \mathbf{e}_{ab}(\mathbf{x}_a, \mathbf{x}_b)^T \Omega_{ab} \mathbf{e}_{ab}(\mathbf{x}_a, \mathbf{x}_b) \quad (1)$$

where Ω_{ab} is the information matrix associated with the measurement collected between \mathbf{x}_a and \mathbf{x}_b and the error is defined as follows :

$$\mathbf{e}_{ab}(\mathbf{x}_a, \mathbf{x}_b) = \mathbf{h}(\mathbf{x}_a, \mathbf{x}_b) \boxminus \mathbf{z}_{ab}. \quad (2)$$

where $\mathbf{h}(\mathbf{x}_a, \mathbf{x}_b)$ is a function that computes the rigid transformation between the two poses, \mathbf{z}_{ab} is the odometry or loop closure measurement, and \boxminus is the difference operator for the manifold domain. We use the notation \mathbf{x}_i to both refer to a robot pose and the corresponding node, and \mathbf{e}_{ab} to both refer to an error term and the corresponding edge of \mathcal{G} . Following [17], we partition the edges in two sets, $\mathcal{E} = \mathcal{E}_o \cup \mathcal{E}_l$, where $\mathcal{E}_o = \{(a, b) \in \mathcal{E} \mid b = a + 1\}$ is the set of edges associated with the odometry measurements, and $\mathcal{E}_l = \{(a, b) \in \mathcal{E} \mid b \neq a + 1\}$ is the set of edges associated with the loop closure measurements.

To make the notation more fluid the following shortcuts are used:

- $\mathbf{e}_a = \mathbf{e}_{a, a+1}(\mathbf{x}_a, \mathbf{x}_{a+1})$ for the odometry errors;
- $\mathbf{e}_{ab} = \mathbf{e}_{ab}(\mathbf{x}_a, \mathbf{x}_b)$ for the loop closure.

It allows us to rewrite Eq. (1) as:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{(a, a+1) \in \mathcal{E}_o} \mathbf{e}_a^T \Omega_{a, a+1} \mathbf{e}_a + \sum_{(a, b) \in \mathcal{E}_l} \mathbf{e}_{ab}^T \Omega_{ab} \mathbf{e}_{ab} \quad (3)$$

C. Independent Subgraphs and Solution Propagation

In the simplest case, a loop closure measurement (e.g., the green edge between \mathbf{x}_1 , and \mathbf{x}_5 in Fig. 2) creates in the pose graph \mathcal{G} a loop without edge crossing or internal loops; all nodes that compose this loop form a subgraph \mathcal{G}' (in the example, \mathcal{G}' is composed of the set $\{\mathbf{x}_1, \dots, \mathbf{x}_5\}$, highlighted in purple) that can be optimized independently from the other nodes in $\mathcal{G} \setminus \mathcal{G}'$. A subgraph defined in this way always defines a time interval that includes all nodes, and related edges, from index a to index b , with $a < b$ and $b - a > 1$. More formally, we can define $\mathcal{G}' = \mathcal{G}(a, b) = (\mathcal{V}(a, b), \mathcal{E}(a, b))$ where $\mathcal{V}(a, b) = \{\mathbf{x}_i \mid a \leq i \leq b\}$ and $\mathcal{E}(a, b) = \{\mathbf{e}_{ij} \mid a \leq i, j \leq b\}$. If no other node $\mathbf{x}_k \in \mathcal{G} \setminus \mathcal{G}(a, b)$ is able to generate a loop with any of $\mathbf{x}_i \in \mathcal{G}(a, b)$, we define $\mathcal{G}(a, b)$ as an *independent*

Algorithm 1: IPC

```

1  $CnS, \mathcal{E}_o, \mathcal{E}_l, C \leftarrow \emptyset$ 
2  $\mathcal{V} \leftarrow \{\mathbf{x}_0\}$ 
3 upon receiving  $\mathbf{z}_{ij}$  then
4   // odometry measurement
5   if  $j = i + 1$  then
6      $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i \boxplus \mathbf{z}_{i,i+1}$  // integrate motion
7      $\mathbf{e}_i \leftarrow \text{edge}(\mathbf{x}_i, \mathbf{x}_{i+1}, \mathbf{z}_{i,i+1}, \Omega_{i,i+1})$ 
8      $\mathcal{E}_o \leftarrow \mathcal{E}_o \cup \{\mathbf{e}_i\}$  // update edges
9      $\mathcal{V} \leftarrow \mathcal{V} \cup \{\mathbf{x}_{i+1}\}$  // update variables
10  // loop closure measurement
11  else
12    // find initial  $\mathcal{G}^I$ 
13     $\mathbf{e}_{ij} \leftarrow \text{edge}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{z}_{ij}, \Omega_{ij})$ 
14     $\mathcal{V}^I, \mathcal{E}^I \leftarrow \text{find } \mathcal{G}_{ij}^I$ 
15    // find cross edges until independent
16     $a \leftarrow i$ 
17    while  $\exists \mathbf{x}_k \in \mathcal{G} \setminus \mathcal{G}^I \mid \exists \mathbf{e}_{kg} \in CnS, a < g < j$  do
18       $\mathcal{E}^I \leftarrow \mathcal{E}^I \cup \{\mathbf{e}_k, \dots, \mathbf{e}_a\}$ 
19       $\mathcal{E}^I \leftarrow \mathcal{E}^I \cup \{\mathbf{e}_{kg}\}$ 
20       $\mathcal{V}^I \leftarrow \mathcal{V}^I \cup \{\mathbf{x}_k, \dots, \mathbf{x}_{a-1}\}$ 
21       $a \leftarrow k$ 
22     $\mathcal{E}_{so}, \mathcal{E}_{sl} \leftarrow \mathcal{E}^I$ 
23    solve Eq. (6) // PGO
24    if  $\mathbf{e}_{ab}^T \Omega_{ab} \mathbf{e}_{ab} < \chi_{\alpha, \delta}^2, \forall \mathbf{e}_{ab} \in \mathcal{E}^I$  then
25       $CnS \leftarrow CnS \cup \{\mathbf{e}_{ij}\}$ 
26       $\mathcal{E}_l \leftarrow \mathcal{E}_l \cup \{\mathbf{e}_{ij}\}$ 
27      propagate solution to  $\mathbf{x}_k \in \mathcal{V} \setminus \mathcal{V}^I$  as described
28      in Sec. III-C
29    else
30      restore( $\mathcal{V}^I$ )

```

subgraph $\mathcal{G}^I(a, b)$ that can be optimized independently from all the other nodes (see Fig. 2). The result of this optimization can be *propagated* to all other nodes $\mathbf{x}_k \in \mathcal{G} \setminus \mathcal{G}^I$ simply by applying the estimated transformation updates of the node in \mathcal{G}^I they are connected to. For example, after optimizing the subgraph highlighted in purple in Fig. 2, the result of this optimization can be propagated to the nodes $\mathbf{x}_k \in \mathcal{G}(6, 13)$ by adding the transformation update applied to \mathbf{x}_5 .

Independent subgraphs can contain multiple loops and cross edges: for example in Fig. 2 the arrival of a loop closure between node 7 and node 13 causes the fusion between the two previous independent subgraphs represented by the light green ovals.

D. The IPC Algorithm

The IPC algorithm is fundamentally very simple and exploits and updates a set CnS (called **consensus set**) of loop closure measurements that have been accepted in previous iterations. The following steps briefly outline the IPC algorithm, where we report in square brackets the used assumptions and heuristics (see Sec. III-A):

- Integrates all odometry measurements and generates a new node when necessary (for example, if the distance to the previous node is above a certain threshold)[**A1**];
- For each incoming loop closure measurement \mathbf{z}' , find

the (minimal) independent subgraph \mathcal{G}' (see Sec. III-C) that contains such constraint by considering odometry constraints and only the loop closure constraint in CnS collected so far;

- Solve the PGO problem only for the independent subgraph \mathcal{G}' [**H1**];
- Accept and include \mathbf{z}' in CnS only if *all* measurements in \mathcal{G}' are consistent (i.e., *agree*, see below) with the new solution obtained in the previous point [**A1, A2, H1, H2**];
- If \mathbf{z}' has been accepted, propagate the obtained solution to the remaining nodes as described in Sec. III-C.

To check whether a measurement agrees with a PGO solution, we exploit the χ^2 test [13]:

$$\mathbf{e}_{ab}^T \Omega_{ab} \mathbf{e}_{ab} < \chi_{\alpha, \delta}^2, \forall \mathbf{e}_{ab} \in \mathcal{E}' \quad (4)$$

where α is the confidence parameter and δ the degrees of freedom of the error and \mathcal{E}' is the edgset of \mathcal{G}' . The error term is computed here by using as robot pose the solution obtained by the PGO carried out on \mathcal{G}' . If one of the tests fails, the new loop closure measurement \mathbf{z}' is *not* accepted, hence not included in CnS .

We report the pseudocode of our method in Alg. 1, which is described step by step below. For the sake of simplicity, we assume here to create a new node for each new odometry measurement.

As soon as a new measurement \mathbf{z}_{ij} is received by the system, the first operation to carry out is to identify if it is an odometry or a loop closure measurement. If an odometry measurement is received, then $\mathbf{z}_{ij} = \mathbf{z}_{i,i+1}$ (line 4), a new variable \mathbf{x}_{i+1} is initialized integrating the motion estimate $\mathbf{z}_{i,i+1}$ from \mathbf{x}_i (line 5), and then added to \mathcal{V} . The final operation is to update the set \mathcal{E}_o with the new edge \mathbf{e}_i relative to the measurement $\mathbf{z}_{i,i+1}$ (line 7).

In the scenario a loop closure measurement \mathbf{z}_{ij} is received (line 9), its relative edge \mathbf{e}_{ij} is created. Then, \mathbf{e}_{ij} is used to identify the minimal independent subgraph $\mathcal{G}^I(a, j)$ (line 11-17) such that :

$$\exists \mathbf{x}_k \in \mathcal{G} \setminus \mathcal{G}^I(a, j) \mid \exists \mathbf{e}_{kg} \in CnS, a < g < j \quad (5)$$

Once, $\mathcal{G}^I(a, j)$ has been identified, its edges are decomposed in the following manner $\mathcal{E} = \mathcal{E}_o \cup \mathcal{E}_l$. This step is fundamental because being the odometry measurements the "trustworthy" ones, a scaling factor $s > 1$ is applied to their information matrix to increase the importance of minimizing their error during the optimization using Eq. (6) (line 19):

$$\sum_{a \in \mathcal{E}_o} \mathbf{e}_a^T s \Omega_{a, a+1} \mathbf{e}_a + \sum_{(a, b) \in \mathcal{E}_l} \mathbf{e}_{ab}^T \Omega_{ab} \mathbf{e}_{ab} \quad (6)$$

It allows the solver to prioritize maintaining the local shape of the trajectory, actively reducing the acceptance rate of outliers. The scaling factor is fixed and doesn't play a role in the optimization procedure, thus avoiding an increase in the problem's complexity, as opposed to SC [16] where additional extra variables were added. In case Eq. (4) is not satisfied, it means that with the addition of \mathbf{e}_{ij} the optimizer is not able to find a configuration of the nodes for which the new edge is consistent with the measurements in its \mathcal{G}^I and the elements

of \mathcal{V}^I are restored to its state previous the optimization. Instead, if Eq. (4) is satisfied, a consistent configuration is found leading to the acceptance of the loop closure to the CnS (line 21). The current solution obtained for the nodes \mathcal{V}^I is then propagated to the nodes that didn't participate in the optimization $\mathcal{V} \setminus \mathcal{V}^I$ (line 23).

IV. EXPERIMENTS

We implemented the proposed method (called IPC in all experiments) in C++ building upon the g2o [2] framework, exploiting the Dog-Leg optimization steps and setting $s = 3$ for all experiments. We compared IPC with 7 recent alternative, state-of-the-art approaches, reported below along with the acronym used in the plots and the used implementation. Most of these methods aim to mitigate the anomalous problem in loop closure detection. For Max-Mixture method [10] (MAXMIX in the plots) in the experiment we used the implementation provided in OpenSLAM¹. We implemented inside g2o the Adaptive Kernel [5] robust kernel (ADAPT in the plots). For Pairwise consistent measurement set maximization method [9] (PCM in the plots), Kimera's implementation is used [23]. The remaining methods are provided by the gtsam library [3]: Graduated Non-Convexity [8], Dynamic Covariance Scaling [6], Huber and German-McClure [7], will be referred to in the plots respectively as GNC, DCS, HUBER, and GM. For all methods we kept all their parameters constant, reporting for each one of them the results obtained with the best set of parameters. Experiments were carried out in a cluster PC using a single Intel(R) Xeon(R) Gold 5220 CPU.

A. Datasets

The datasets used for evaluation are currently the standard for evaluating the performances of back-end optimizers and are described in [24] [25]. They can be divided into two categories, synthetic and real. The Intel, Mit, Csail, Freiburg Building (Fig. 1), and Freiburg University Hospital are real datasets, while M3500 is synthetic. In both cases, the datasets are outlier-free. The loop closure edges in the unaltered datasets were used as the ground truth for the inliers. Outliers were added, using the Vertigo package [26], as a percentage of the total inliers present in the dataset in order to have balanced results.

B. Metrics

The **Precision**, **Recall**, and **F1 score** metrics are used in the evaluation, the latter defined from the former:

$$Prec = \frac{TP}{TP + FP}, Rec = \frac{TP}{TP + FN}, F1 = 2 * \frac{Prec * Rec}{Prec + Rec}$$

where **TP**, **FP**, **TN** **FN** stand for a loop closure that is a true and false positive and true and false negative, respectively. Finally, the solution obtained using SE-Sync [27], which is a fast non-minimal solver for PGO, on the outlier free version of the dataset was used as ground truth trajectory for the evaluation of both **Absolute Trajectory Error** (ATE) and **Relative Pose Error** (RPE) [28], that will be used to measure the accuracy of estimated trajectory.

¹<https://openslam-org.github.io/>

	50%			100%		
	F1 ↑	RPE ↓	TIME ↓	F1 ↑	RPE ↓	TIME ↓
GNC	0.747	32.20	155.86	0.604	35.53	349.04
ADAPT	0.873	1.867	5457.07	0.738	1.856	12576.27
MAXMIX	0.77	0.324	70.99	0.66	0.43	137.29
DCS	0.796	26.12	4.58	0.66	37.684	8.28
GM	0.795	23.64	7.07	0.65	38.33	9.59
HUBER	0.706	21.44	10.04	0.55	27.32	8.29
PCM	0.53	25.01	8289.04	0.371	28.54	12655.41
IPC	0.91	0.05	321.37	0.89	0.068	442.733

TABLE I

SUMMARY OF THE MOST IMPORTANT METRICS IN FIG. 3 AND IN FIG. 4 FOR 50% AND 100% OF OUTLIERS.

C. Evaluation

The introduced outliers range from 10% to 100% of the ground truth inliers, with increments of ten percent. To ensure fair and robust testing, we generated ten distinct trajectories for each outlier percentage. A percentage of 100% means that the number of outliers is equal to the number of inliers. As done in [14], for methods that don't explicitly classify measurements, an inlier is defined as a measurement having a χ^2 error less than χ_{th}^2 with $\alpha_{th} = 0.95$. For all methods, the different CnS s are evaluated at the end of the optimization of the full problem. One small note to take into account is that results for both GNC and PCM were taken after a fixed number of iterations. It was done due to their large convergence time and the exhaustive amount of testing that was carried out for each method.

Fig. 3 shows the performances for precision, recall, and F1 score averaged on all the used datasets. The precision is one of the most crucial parameters for a correct termination for the optimization. Its value should ideally approach one since even minor outlier contributions have the potential to severely compromise the trajectory estimate. MAXMIX and IPC are the methods that reject most of the outliers across the different datasets, closely followed by PCM. Conversely, for GNC, DCS, and GM when the percentage of outliers exceeds 50%, the precision drops below 0.5, strongly reducing the probability that the solver reaches a good solution. Simply having high precision is not sufficient, because having poor recall performances will negatively impact the elimination of the accumulated drift due to the high number of inliers discarded, also resulting in poor trajectory estimation. Still, the acceptable bound is not as strict as the one used for precision. As shown in Fig. 3, all methods have an almost constant recall trend as the percentage of outliers increases. GM and DCS demonstrate the highest recall rates, followed by ADAPT and IPC which provide a consistent recall of over 80%. Nonetheless, IPC outperforms all these methods in the accuracy metric by a large margin, as shown in the RPE column in Tab. I. The F1 score, which relates both precision and recall together, is able to give a global view of the actual methods' performances. It can be seen that the methods report a similar trend with respect to the precision. It is an expected behavior given the quasi-constant performance of the recall. ADAPT slightly dominates for smaller percentages of outliers, while IPC is able to keep stable and reliable results for all the different percentages, resulting in having the best performances for higher percentages of outliers.

Fig. 4 shows respectively the convergence time, the ATE,

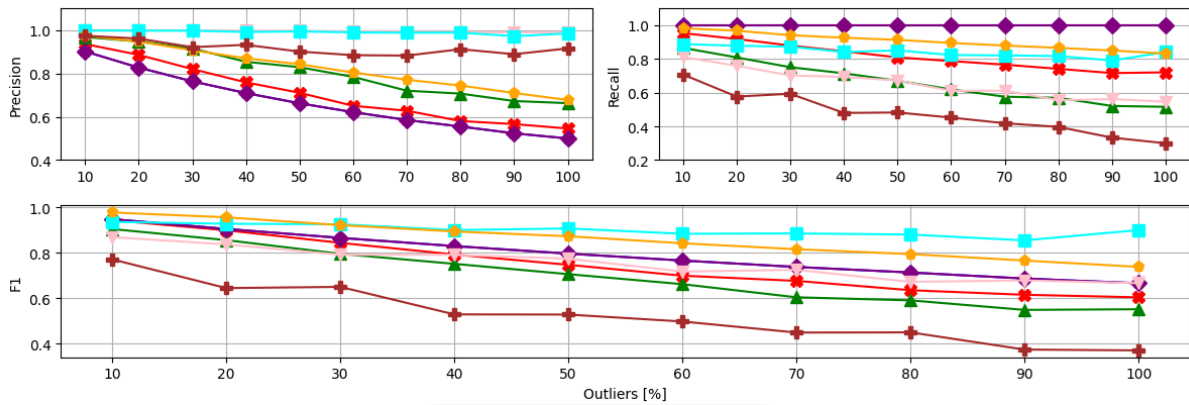


Fig. 3. Top left, top right, and bottom images are respectively the precision, the recall, and F1 score performances of the methods for varying outlier percentages. Legend of the graphs: HUBER (▲), GM(●), IPC(□), DCS(◆), MAXMIX(▽), ADAPT(○), PCM(⊕), GNC(⊗). GM is completely covered by DCS in both precision and recall plots, while MAXMIX is nearly entirely covered by IPC in the precision plot.

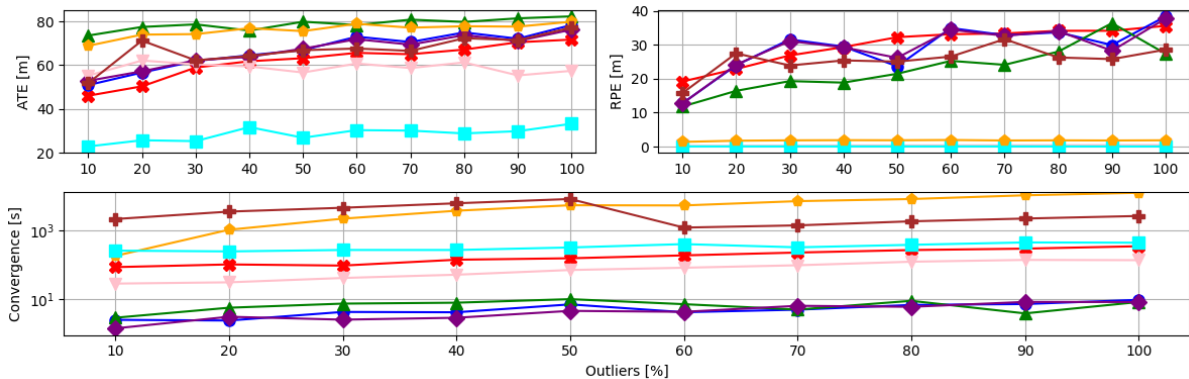


Fig. 4. Top left, top right, and bottom images are respectively the ATE, the RPE, and convergence time of the methods for varying outlier percentages. Legend of the graphs: HUBER (▲), GM(●), IPC(□), DCS(◆), MAXMIX(▽), ADAPT(○), PCM(⊕), GNC(⊗). The MAXMIX plot for the RPE metrics is totally covered by the ADAPT plot.

and the RPE metrics. We can see that for the ATE almost all methods have similar results, with IPC clearly overtaking the other methods. ATE highlights how globally brittle the optimization is even in the presence of a small number of outliers, while RPE measures the local consistency of the estimated trajectories. IPC, MAXMIX, and ADAPT have the best RPE of the different methods.

To compare the offline methods' time performances with IPC, it is used the time needed to validate all the loop closure constraints provided in the different datasets. As expected, it is not able to compete with the robust estimators and DCS in terms of speed, but it provides comparable results with more complex methods such as GNC and MAXMIX, even if they were specifically designed to work offline.

In Tab. II for each dataset is reported the Average Convergence Time per Constraint (ACTxC), that is the average convergence time on the subproblem identified by a loop closure constraint. The worst results are the ones for the M3500, because the majority of outliers captured a vast section of the trajectory, leading to predominantly solving subproblems of high dimensionality. Also, in case a false positive of such dimension is accepted, it would negatively impact the clusters' size and the future constraints' convergence time. Still, the results obtained fit the online deployment of the method.

TABLE II
IPC'S ACTxC FOR EACH DATASET AS THE OUTLIER PERCENTAGE VARIES. THE REPORTED RESULTS ARE EXPRESSED IN SECONDS.

Dataset	20%	40%	60%	80%	100%
MIT	0.053	0.050	0.050	0.050	0.044
INTEL	0.056	0.051	0.058	0.054	0.056
M3500	0.577	0.542	0.717	0.601	0.623
CSAIL	0.043	0.042	0.043	0.042	0.043
FRH	0.098	0.095	0.100	0.098	0.109
FR079	0.042	0.042	0.043	0.041	0.043

V. CONCLUSIONS

We presented IPC, a consensus-based method to incrementally build the maximum coherent set of measurements. IPC demonstrates comparable or better performance with respect to recent state-of-the-art approaches while providing online capabilities. In contrast to most methods, IPC shows great generalization capabilities by showing consistent results across the different datasets without changing its parameters, and its only control parameter s is fairly simple to configure. Currently, IPC's primary limitation is its inability to rectify past erroneous decisions. Therefore, a possible future direction could involve investigating the specific contexts where IPC is more prone to making errors and also integrating an offline method on top of it that revises and possibly discards the previously integrated measurements, in order to gain the benefits of both methodologies.

REFERENCES

- [1] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. Leonard, “Past, present, and future of simultaneous localization and mapping: Towards the robust-perception age,” *IEEE Transactions on Robotics*, vol. 32, no. 6, p. 1309–1332, 2016.
- [2] G. Grisetti, R. Kümmerle, H. Strasdat, and K. Konolige, “g2o: A general framework for (hyper) graph optimization,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 9–13.
- [3] F. Dellaert and G. Contributors, “borglab/gtsam,” May 2022. [Online]. Available: <https://github.com/borglab/gtsam>
- [4] S. Agarwal, K. Mierle, and T. C. S. Team, “Ceres Solver,” 3 2022. [Online]. Available: <https://github.com/ceres-solver/ceres-solver>
- [5] J. T. Barron, “A general and adaptive robust loss function,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4331–4339.
- [6] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard, “Robust map optimization using dynamic covariance scaling,” in *2013 IEEE International Conference on Robotics and Automation*. Ieee, 2013, pp. 62–69.
- [7] M. J. Black and A. Rangarajan, “On the unification of line processes, outlier rejection, and robust statistics with applications in early vision,” *International journal of computer vision*, vol. 19, no. 1, pp. 57–91, 1996.
- [8] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone, “Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1127–1134, 2020.
- [9] J. G. Mangelson, D. Dominic, R. M. Eustice, and R. Vasudevan, “Pairwise consistent measurement set maximization for robust multi-robot map merging,” in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 2916–2923.
- [10] 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.
- [11] J. Matas and O. Chum, “Randomized ransac with td, d test,” *Image and vision computing*, vol. 22, no. 10, pp. 837–842, 2004.
- [12] P. C. Lusk, K. Fathian, and J. P. How, “Clipper: A graph-theoretic framework for robust data association,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 828–13 834.
- [13] J. Neira and J. D. Tardós, “Data association in stochastic mapping using the joint compatibility test,” *IEEE Transactions on robotics and automation*, vol. 17, no. 6, pp. 890–897, 2001.
- [14] D. McGann, J. G. Rogers, and M. Kaess, “Robust incremental smoothing and mapping (risam),” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 4157–4163.
- [15] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, “isam2: Incremental smoothing and mapping using the bayes tree,” *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [16] N. Sünderhauf and P. Protzel, “Switchable constraints for robust pose graph slam,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 1879–1884.
- [17] 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, pp. 1611–1626, 2013.
- [18] F. Wu and G. Beltrame, “Cluster-based penalty scaling for robust pose graph optimization,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6193–6200, 2020.
- [19] J. Shi, H. Yang, and L. Carlone, “Robin: a graph-theoretic approach to reject outliers in robust estimation using invariants,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 820–13 827.
- [20] L. Carlone, A. Censi, and F. Dellaert, “Selecting good measurements via 1 relaxation: A convex approach for robust estimation over graphs,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 2667–2674.
- [21] G. Grisetti, R. Kümmerle, and K. Ni, “Robust optimization of factor graphs by using condensed measurements,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 581–588.
- [22] T. Guadagnino, L. Di Giammarino, and G. Grisetti, “Hipe: Hierarchical initialization for pose graphs,” *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 287–294, 2021.
- [23] A. Rosinol, A. Violette, M. Abate, N. Hughes, Y. Chang, J. Shi, A. Gupta, and L. Carlone, “Kimera: From slam to spatial perception with 3d dynamic scene graphs,” *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1510–1546, 2021.
- [24] L. Carlone and A. Censi, “From angular manifolds to the integer lattice: Guaranteed orientation estimation with application to pose graph optimization,” *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 475–492, 2014.
- [25] L. Carlone, R. Aragues, J. A. Castellanos, and B. Bona, “A fast and accurate approximation for planar pose graph optimization,” *The International Journal of Robotics Research*, vol. 33, no. 7, pp. 965–987, 2014.
- [26] N. Sünderhauf and P. Protzel, “Vertigo (versatile extensions for robust inference using graph optimization) source code,” 2019. [Online]. Available: <https://openslam-org.github.io/vertigo.html>
- [27] D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard, “Se-sync: A certifiably correct algorithm for synchronization over the special euclidean group,” *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 95–125, 2019.
- [28] Z. Zhang and D. Scaramuzza, “A tutorial on quantitative trajectory evaluation for visual (inertial) odometry,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 7244–7251.