MIT Open Access Articles

LOCUS: A Multi-Sensor Lidar-Centric Solution for High-Precision Odometry and 3D Mapping in Real-Time

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: 10.1109/LRA.2020.3044864

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/134380>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

LOCUS: A Multi-Sensor Lidar-Centric Solution for High-Precision Odometry and 3D Mapping in Real-Time

Matteo Palieri^{1,2}, Benjamin Morrell¹, Abhishek Thakur³, Kamak Ebadi¹, Jeremy Nash¹, Arghya Chatterjee⁴,
 Christoforos Kanellakis⁵, Luca Carlone⁶, Cataldo Guaragnella², Ali-akbar Agha-mohammadi¹

Abstract—A reliable odometry source is a prerequisite to enable complex autonomy behaviour in next-generation robots operating in extreme environments. In this work, we present a high-precision lidar odometry system to achieve robust and real-time operation under challenging perceptual conditions. LOCUS (Lidar Odometry for Consistent operation in Uncertain Settings), provides an accurate multi-stage scan matching unit equipped with a health-aware sensor integration module for seamless fusion of additional sensing modalities. We evaluate the performance of the proposed system against state-of-the-art techniques in perceptually challenging environments, and demonstrate top-class localization accuracy along with substantial improvements in robustness to sensor failures. We then demonstrate real-time performance of LOCUS on various types of robotic mobility platforms involved in the autonomous exploration of the Satsop power plant in Elma, WA where the proposed system was a key element of the CoSTAR team’s solution that won first place in the Urban Circuit of the DARPA Subterranean Challenge.

I. INTRODUCTION

Robotic systems are rapidly entering in all aspects of human life. In particular, robots are being deployed in increasingly complex environments for a broad spectrum of applications ranging from mining [1] and search-and-rescue [2], to industrial monitoring [3] and planetary exploration [4]. In these scenarios, darkness, presence of obscurants (e.g. fog, dust, smoke), lack of prominent perceptual features in self-similar areas, and slippery terrains (leading to jerky sensory motion) are common features that pose severe perceptual challenges to robotic operation. In this work, we focus on developing an accurate and reliable odometry estimation method (i.e., estimating the robot movement) which is a key

This work was supported by the Jet Propulsion Laboratory - California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). This work was partially funded by the Defense Advanced Research Projects Agency (DARPA). ©2020 All rights reserved.

¹Palieri, Morrell, Ebadi, Nash, and Agha-mohammadi are with NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA matteo.palieri@poliba.it

²Palieri and Guaragnella are with the Department of Electrical And Information Engineering, Polytechnic University of Bari, IT matteo.palieri@poliba.it

³Thakur is with Aptiv, Troy, MI, USA abhi.dtull@gmail.com

⁴Chatterjee is with Bangladesh University of Engineering and Technology, Dhaka, Bangladesh. arghyame20buet@gmail.com

⁵Kanellakis is with Luleå University of Technology, Luleå, Sweden. christoforos.kanellakis@ltu.se

⁶Carlone is with Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA. lcarlone@mit.edu

Fig. 1. Testing of the proposed lidar odometry system, LOCUS, in the DARPA Subterranean Challenge. a) Urban Circuit test environment with a Husky wheeled robot, showing unavoidable rubble. b) Tunnel Circuit test environment with a Husky wheeled robot carrying two lidars, showing self-similar areas with low lidar observability. c) Tunnel dataset ground truth map. d) Example Urban Circuit test with a legged robot carrying one lidar. Both robots run LOCUS. e) Map produced by LOCUS on one Husky robot in the Urban Beta course. f) Urban Alpha ground truth map.

requirement to enable advanced autonomous behaviours in such perceptually-challenging conditions.

As lidar sensors are less sensitive to illumination variations and provide high-fidelity, long-range 3D measurements, they have been commonly used for robotic odometry estimation in the last decade. Typically, lidar odometry (LO) algorithms estimate the ego motion of the robot by comparing and registering consecutive lidar acquisitions. When it comes to perceptually-challenging settings, lidars are commonly fused with additional sensing modalities to improve accuracy. However, in these settings, potential failures in different sensing modalities can degrade, or drastically compromise the odometry performance.

In this paper, we present LOCUS (Lidar Odometry for Consistent operation in Uncertain Settings), a lidar odometry system that (i) enables accurate real-time operation in extreme and perceptually-challenging scenarios, and (ii) is robust to intermittent and faulty sensor measurements. LOCUS has been a key element of the CoSTAR team’s solution [5] that won first place at the Urban Circuit of the DARPA Subterranean Challenge (SubT Challenge), where robots are tasked to autonomously explore complex GPS-denied underground environments (e.g. Fig. 1).

A. Related Works

LO algorithms can be categorized by the representation type and the number of points (or features) used to align

lidar-scans, including (i) feature-based methods, (ii) grid-based methods, and (iii) dense methods.

Feature-based methods: Feature-based methods rely on extracting and matching salient features across consecutive lidar-scans to estimate the ego motion of the robot. Possible features can include planar and edge features [6]–[9], ellipsoidal surfels [10] and ground features [11]. These features can be matched by proximity [12], type [6], or descriptor [10], depending on the algorithm and feature type.

Grid-based methods: Probability grid methods map lidar-scans into grids and compute the occupancy probability densities that can later be matched using Newton’s method [13].

Dense methods: Dense methods work with a large subset of lidar-scan points. As using the full point cloud can be computationally expensive for real-time operation, most approaches select a subset of points for scan matching. The Generalized Iterative Closest Point (GICP) [14] is a common dense point-based scan matching method, where local surface normal information is used to better address the measurement noise in scan matching by using both point-to-point and point-to-plane matching, with planes being evaluated in local neighborhoods. LOCUS falls into this category.

Scan-to-scan alignment: The computation of the optimal alignment between two scans can be cast as a non-linear optimization problem, addressed by various solvers including Levenberg-Marquardt (e.g. [6]), iterative gradient descent (e.g. [8], [12], [14]), optimization tools such as Ceres [15], least-squares solvers (e.g. [16]), or in a sliding-window, pose-graph structure (e.g. [17], [18]) with GTSAM [19].

Scan-to-map alignment: To enable global consistency across the history of scans, the computed pose is refined by aligning the current scan and the existing map. For various map representations, ranging from feature-based (e.g. [6], [11], [17], [20]), to grid-based maps (e.g. [21]) and point-based maps (e.g. [12]), one can adopt different scan-to-map alignment methods. This includes point-based alignment methods (e.g. [12]), Normalized Distributions Transform (e.g. [13], [22]) or smoothing function alignment (e.g. [21]).

Sensor fusion: While pure lidar-based methods are powerful, their performance can significantly degrade when it comes to perceptually-challenging conditions, including environments with geometrically self-similar patterns or agile robots with high-rate motions. To address these challenges, it is important to fuse lidar with additional sensing modalities, such as an inertial measurement unit (IMU) or visual-inertial odometry (VIO). The IMU can provide rotational estimates that are tightly integrated (e.g. [9], [18], [20], [21]) or loosely integrated (e.g. [6], [8], [10], [17]) with the scan matching process. When the drift is translational (referred to as *lidar-slip* in this paper), VIO, wheel-inertial odometry (WIO) and kinematic-inertial odometry (KIO) can complement IMUs by providing a full 6-DOF transform estimate. Tight integration (e.g. [7]) and loose integration (e.g. [17], [21]) of these methods with lidars have shown significant improvements over individual use of either one of these modalities. LOCUS follows the loosely-coupled model, where in addition to improving accuracy, these sensor fusion methods can reduce

Fig. 2. Architecture of the proposed lidar odometry system.

computation by providing a near-optimal prior to the lidar scan-matching optimization in the GICP.

B. Method Highlights and Contributions

The highlights and contributions of LOCUS are:

1) Architecture: The LOCUS architecture (see Fig. 2) enables accurate, robust, and real-time odometry in perceptually-stressing settings and alleviates sensor failure challenges. The architecture can be adapted to heterogeneous robotic platforms with diverse sensor inputs and computational capabilities.

2) Resilience: The system is fail-safe to drops or loss of one or more sensor channels by relying on a loosely-coupled switching scheme between sensing modalities.

3) Environment adaptability: The system further enables incorporation of domain knowledge (if available), such as flat grounds in human-made structures.

4) Field demonstration: We present an extensive field demonstration of LOCUS. In particular, we provide results and insights from deploying LOCUS as part of the CoSTAR team’s solution that won the Urban Circuit of the SubT Challenge. We present an ablation study on LOCUS and then compare the performance with six state-of-the-art methods using the data acquired in the field tests.

II. SYSTEM DESCRIPTION

In this section, we describe the system architecture reported in Fig. 2 and provide details of each submodule.

A. Point Cloud Preprocessor

Motion Distortion Correction: We assume information from one or more 360 degree lidar sensors, such as the Velodyne Puck or Ouster lidars. The raw information coming from the lidar is fed into a motion distortion correction (MDC) unit which corrects the Cartesian position of each point to account for the motion of the robot while a single scan of the lidar is completed¹. This correction is particularly important for points at large range, when high-rate rotations are experienced by the robot, and is a commonly employed step [17], [20]. The correction is informed by either an IMU, or an odometry source (e.g VIO, WIO, KIO) where

¹The Velodyne Puck lidar requires 0.1 s to complete one scan

the chosen source depends on what is reliably available and calibrated on a given robot.

Point Cloud Merger: For robots with multiple lidars, to enlarge the overall robot field-of-view, the point cloud merger (Fig. 2, top and middle) combines each motion-corrected point cloud into a single one, using the known rigid body transformation between sensors. This step is implemented so the rest of the pipeline is consistent for robots with one or multiple lidar sensors.

Point Cloud Filter: The resulting point cloud is then processed in the point cloud filter to remove noise and out-of-range points, manage the volume of data, and reduce computational load. The point cloud filter is composed by a sequential combination of a 3D voxel grid filter and a random downsampling filter, which can be individually tuned, activated and deactivated. The voxel grid filter takes the average of the points in each 3D volume (a voxel) to decrease the data size while still capturing the dominating structure of the environment. We use a voxel size of 0.1m in the tests presented in this paper. For the random downsampling filter, we use an implementation of [23] with a downsampling percentage of 90%. For both filters we use the implementation in the Point Cloud Library [24].

B. Scan Matching Unit

The scan matching unit (light blue box in Fig. 2) performs a GICP-based scan-to-scan and scan-to-submap matching operation to estimate the 6-DOF motion of the robot between consecutive lidar acquisitions.

Notation: We denote with $\mathcal{R}$ the coordinate system of the robot and with $\mathcal{W}$ the coordinate system of the world, that coincides with $\mathcal{R}$ at the start of the test. We therefore address the problem of determining the poses of $\mathcal{R}$ with respect to $\mathcal{W}$ by means of consecutive lidar acquisitions to incrementally reconstruct a trajectory and a map of the explored environment. We denote with L_k the lidar scan collected at time k and with L_{k-1} the lidar scan collected at time $k-1$. All lidar scans are expressed in the robot frame. We denote with $\mathbf{X}_k \in SE(3)$ the robot pose in $\mathcal{W}$ at time k and with $\mathbf{X}_{k-1}$ the robot pose in $\mathcal{W}$ at time $k-1$. We denote with $\mathbf{T}_k^{k-1} = \mathbf{X}_{k-1}^{-1} \mathbf{X}_k$ the rigid body transformation between two consecutive robot poses, where $\mathbf{T}_k^{k-1} \in SE(3)$ is the transform between $\mathbf{X}_{k-1}$ and $\mathbf{X}_k$.

1) *Sensor integration module:* In robots with multi-modal sensing, if available, we use an initial transform estimate from a non-lidar source in the scan-to-scan matching stage to improve accuracy and reduce computation.

Health monitoring: Multiple sources of odometry (e.g VIO, KIO, WIO) and raw IMU measurements are first transformed into $\mathcal{R}$, and then fed into a health monitor which selects an output from a priority queue of inputs that are deemed healthy (see bottom left of Fig. 2). The system is designed to take in a variety of sources of health metrics to evaluate the health of input sources. For example, ongoing work is looking to integrate with the Heterogeneous Robust

Odometry (HeRO) system [25] that employs custom health analysis (such as feature counts and observability analysis) on each odometry source, as well as rate and covariance checks. For our implementation presented below, we use a simple rate-check: if input messages are at a sufficient rate ($> 1\text{Hz}$), then the source is healthy.

Priority queue: The priority queue is intended to always select the highest accuracy source, based on previous testing for a given robotic system in similar environments. If the robot enters an area where the highest priority source is degraded, it is intended for this to be reflected in the health metric, that would trigger a transition to the next highest, healthy input. With this health-metric-driven dynamic switching, the priority queue is static. The priority queue for our legged robot is: VIO, KIO, IMU, no input, and for our wheeled robot is: VIO (if present), WIO, IMU, no input.

We define the pose estimate (with respect to $\mathcal{W}$) of the highest priority source that is found to be healthy as $\mathbf{Y}$. To reduce operations, we buffer only $\mathbf{Y}$ and interpolate the buffered data at lidar timestamps, t_{k-1} , and t_k to get $\mathbf{Y}_{k-1}$ and $\mathbf{Y}_k$: the pose of highest priority healthy source at times t_{k-1} and t_k , respectively. We denote with $\mathbf{E}_k^{k-1} = \mathbf{Y}_{k-1}^{-1} \mathbf{Y}_k$ the rigid body transformation of the sensor integration module output between $\mathbf{Y}_{k-1}$ and $\mathbf{Y}_k$ in the $[t_{k-1}, t_k]$ time interval where $\mathbf{E}_k^{k-1} \in SE(3)$. Each odometry source provides a rotation and translation, whereas for the IMU we only use the rotation measurements.

2) *Scan-to-scan:* In the scan-to-scan matching stage, GICP computes the optimal transformation $\hat{\mathbf{T}}_k^{k-1}$ that minimizes the residual error $\mathcal{E}$ between corresponding points in L_{k-1} and L_k .

$$\hat{\mathbf{T}}_k^{k-1} = \arg \min_{\mathbf{T}_k^{k-1}} \mathcal{E}(\mathbf{T}_k^{k-1} L_k, L_{k-1}) \quad (1)$$

When the sensor integration module is successful, we initialize the GICP with $\mathbf{T}_k^{k-1} = \mathbf{E}_k^{k-1}$. If all sensors fail, the GICP is initialized with identity rotation and zero translation and the system reverts to pure lidar odometry.

3) *Scan-to-submap:* The motion estimated in the scan-to-scan matching stage is further refined by a scan-to-submap matching step. Here L_k is matched against a local submap S_k which is taken from the local region of the global map M_k given the current estimate of the robot pose in $\mathcal{W}$.

$$\tilde{\mathbf{T}}_k^{k-1} = \arg \min_{\mathbf{T}_k^{k-1}} \mathcal{E}(\mathbf{T}_k^{k-1} L_k, S_k) \quad (2)$$

In this optimization, $\mathbf{T}_k^{k-1}$ is initialized with the $\hat{\mathbf{T}}_k^{k-1}$ result from Eqn. 1. The global map is a point cloud stored in an octree format that is an accumulation of point clouds after every t meters of translation, or r degrees of rotation: for our results, we use $t = 1$, $r = 30^\circ$. We use an octree with a minimum resolution of 0.001m to store the map, which usually maintains all points in an easily searchable format.

Output: After scan-to-scan and scan-to-submap matching, the final estimated motion $\tilde{\mathbf{T}}_k^{k-1}$ between consecutive lidar acquisitions is used to update the robot pose in $\mathcal{W}$: the generated odometry is therefore the integration of all computed incremental transforms.

Both accuracy and computational speed are improved by the incremental estimation from input odometry to scan-to-scan and finally scan-to-map (shown, for example, in [12]). A good initial estimate in both Eqn. (1) and Eqn. (2) reduces the chances of converging in a sub-optimal local minima, and reduces the number of iterations needed to converge, lowering computation time.

4) *Notes on multi-threading:* The computational speed of the scan-to-scan and scan-to-submap matching has been greatly increased through the development of a multi-threaded GICP approach, modified from the PCL implementation [24]. The multi-threading utilizes a user-specified number of cores for the normal computation stage in GICP, which represents over 70% of the computation time in the process. For the evaluations performed in this paper, we use 4 threads unless otherwise stated.

C. Environment Adaptation: Flat Ground Assumption

In human-made environments there are many areas with flat grounds, which if known prior, could be utilized to aid odometry algorithms. When detected or known, the flat ground assumption (FGA) can be activated to limit drift in Z and error in roll and pitch (lower-right blue box in Fig. 2). FGA operates on the output of both scan-to-scan and scan-to-submap alignment, by zeroing any Z movement, roll or pitch, all in a global, gravity aligned reference frame.

FGA activation modalities: The system provides two ways to detect a flat ground and activate FGA: context-based, and sensor-based. The first approach relies on prior knowledge of the environment that can be acquired by a human supervisor, for instance in single floor exploration of urban environments. For stair-climbing robots, the initiation of a stair mission can be used to deactivate FGA, and then reactivate it when the stair mission is complete through the input of a human operator (see [26] for an example). In the second approach, an IMU monitor can be used to detect periods when the robot has near-zero roll and pitch over a sufficient time period to activate FGA, and then deactivate it when this condition is no longer met.

D. Adaptation for Different Platforms

The system includes adjustable components to adapt to heterogeneous robotics platforms with different computing power and sensors. These adaptations are primarily in: the number of lidars, the filtering, and number of threads for GICP and the measurements used for the initial transform estimate. Sec. III-C, demonstrates the flexibility of LOCUS through application to two different robotic platforms.

III. FIELD EXPERIMENTS

In this section we present the experimental results obtained from tests in the Tunnel and Urban circuits of the SubT Challenge. We first use three datasets from a Clearpath Husky ground rover (Fig. 1.a-b) to perform an ablation study on LOCUS, and compare it with state-of-the-art lidar odometry solutions. We then showcase results achieved during live operations in field tests across heterogeneous robotic platforms. See https://youtu.be/5QQkkQ_YrbU for visualization of the results.

Dataset description: Each dataset comprises 3D lidar scans coming from 2 on-board VLP16 lidars (one flat, one pitched forward 30°), along with IMU (Vector Nav 100) and WIO measurements for a 60-minute run. Each dataset is selected to contain components that are challenging for lidar odometry. The Urban datasets (Alpha Course, Fig. 1.f and Beta Course, Fig. 1.e) are collected in a dismissed power plant located in Elma, WA that presents many challenges for robot perception such as long feature-poor corridors and large open spaces (the test area dimensions are 100x100x20 m). The Tunnel dataset (Safety Research Course, Fig. 1.c) is recorded in the Bruceton Research Mine in Pittsburgh, PA that is characterized by self-similar and self-repetitive geometries (the test area dimension are 200x200x10 m). All datasets have substantial vibrations and large, sudden accelerations as is characteristic of a skid-steer wheeled robot traversing rough terrain and rubble. See Fig. 1 for sample images of the environments.

Lidar scans are recorded at 10Hz. WIO and IMU are recorded at 50Hz in the Urban datasets, while a higher-rate IMU recording (100Hz) is available for the Tunnel dataset. Both motion corrected and raw points are available for the Urban datasets, whereas the Tunnel dataset only has raw points available. We use LOCUS to do scan matching on the ground-truth map provided by DARPA to estimate the ground-truth reference of the robot trajectory.

A. Ablation Study

To investigate the impact of each component of LOCUS on the overall pipeline accuracy, we evaluate the Absolute Position Error (APE) of the robot-trajectory in the Urban Alpha dataset². The results are summarized in Fig. 3. The study confirms that the use of motion-corrected points is essential, and highlights the effectiveness of the filtering approaches that limit the data volume and reduce computational load. Feature-based filtering (e.g. LOAM-type features of edges and planes) can result in greater accuracy, yet with a higher CPU load (25% more than the baseline configuration), leading to non-real-time performance on our system. Environmental knowledge of a flat ground is not essential, but can help to improve accuracy for exploration of human-made buildings. The loose sensor integration of WIO or IMU results in minor improvements, however, we use this approach as baseline to robustly operate in scenarios with high-rate motions or low lidar observability.

²While these results are for one dataset, we observed similar trends from tests on the Urban Beta dataset.

Fig. 3. Evolution of the Absolute Position Error (APE) of the proposed method for different processing configurations in the Alpha course of the SubT Challenge. The inset gives more detail on the four best configurations. **baseline**: all features in Sec. II. **imu_int**: no WIO integration, only IMU integration, **no_int**: neither WIO or IMU integration, **loam_feat**: using LOAM feature extraction instead of filtering, **fga_off**: no FGA, **rdf_off**: no random downsample filter, **vgf_off**: no voxel-grid filter, **mdc_off**: no MDC.

TABLE I

SUMMARY OF STATE-OF-THE-ART, OPEN-SOURCE ALGORITHMS

Algorithm	Align.	Opt.	IMU	Odom.	MDC*
LOCUS	Dense	GICP	Loose	Loose	Yes
BLAM [27] ³	Dense	GICP	None	None	No
ALOAM [6] ⁴	Features	Ceres	None	None	No
FLOAM [6] ⁵	Features	Ceres	None	None	No
Cartog. [21] ⁶	Grid	Ceres	Tight	Loose	Yes
LIO-Map. [20] ⁷	Features	Ceres	Tight	None	Yes
LIO-SAM [18] ⁸	Features	GTSAM	Tight	None	Yes

* See section IIIB for more details.

B. Evaluation Against the State-of-the-Art

We compare the proposed algorithm against a variety of the state-of-the-art open-source lidar odometry systems, selected to cover the range of alignment methods and sensor integration methods, as summarized in Table I. While we would like to compare against systems integrating with visual odometry (e.g. [7], [17]), these algorithms do not have open source implementations to readily test. FLOAM and ALOAM are two modern implementations of LOAM aimed to simplify the code structure and increase computational speed, respectively. Cartographer is an LIO algorithm distinguished by its use of grid-based matching. LIO-Mapping is a more recent LIO algorithm that combines the IMU pre-integration approach from VINS-Mono [28] with LOAM-type feature alignment. LIO-SAM is yet more recent, and builds from LeGo-LOAM [11], adding in IMU pre-integration in a smoothing-and-mapping approach.

Comparison criteria: We aim to compare the lidar odometry systems holistically, hence we use three criteria: i) Accuracy, ii) Robustness and iii) Efficiency.

Comparison setup: Each algorithm is setup for the best performance, yet with minimal parameter tuning, only input adjustment (number of lidars, motion corrected points). WIO is the odometry input, if used. Loop closures are disabled to

focus on the lidar odometry performance. Variations in the input for each algorithm are summarized below.

LOCUS, BLAM, and ALOAM each use two lidars in all datasets, with motion-corrected points in the Urban datasets. FLOAM only succeeded with one lidar in the Urban Alpha dataset and otherwise uses two lidars, in each case (with motion correction in Urban datasets). Cartographer can do internal motion correction, but only with the input of points as individual UDP packets from a single lidar. We elected to instead feed pre-corrected scans from two lidars to Cartographer. LIO-Mapping and LIO-SAM are both set up to do internal motion-correction on the point clouds, leveraging the integrated IMU data, hence uncorrected scans are used as inputs. For LIO-Mapping, 2 lidars were used, except for Urban Alpha, where only a 1 lidar test was successful.

LIO-SAM was only able to run with a single lidar input, as there are internal assumptions of a structured point cloud of rings for motion correction and feature extraction. We were not able to get LIO-SAM working on the Urban datasets, likely due to the IMU rate, at 50Hz being lower than the recommended 200Hz for LIO-SAM.

1) *Accuracy Evaluation:* We evaluate the accuracy with two metrics: position error and map error.

Position error: To evaluate the localization accuracy, we use evo [29] to compute the absolute position error (APE) of the trajectories estimated by the different methods against the ground-truth reference. We report in Fig. 4 a boxplot visualization of the APE results for the Urban and Tunnel datasets and summarize these results in Table II.

The results show that LOCUS is equal to or better than the state-of-the-art in all datasets. FGA does help to improve LOCUS performance, but is not essential for LOCUS to perform well. LIO-Mapping has similarly low error, as expected with tight integration of IMU and lidar, yet with a large delay from a mean processing time of 1s per scan. The larger optimization being performed with pre-integration, scan alignment and extrinsic estimation all together likely leads to the longer computation times.

LIO-SAM performs the best in the Tunnel dataset, and with feasible computational speeds, yet the strict requirements on appropriate input data limit the range of platforms and datasets it can be applied to. Cartographer and ALOAM perform relatively well in Urban, showing the effectiveness of edge and plane features as well as grid methods in human-made environments.

For dense alignment methods, BLAM performs adequately in the Urban datasets, but poorly in the Safety Research dataset, suggesting the WIO integration employed by LOCUS is an important component in the self-similar and low lidar observability conditions seen in that dataset.

Map error: As second quantitative evaluation, we compare the maps obtained with each algorithm to the DARPA provided ground-truth map to compute the overall cloud-to-cloud error. To account for potential calibration misalignments, we run Iterative Closest Point (ICP) between the reconstructed map and ground-truth map before performing error analysis. We report in Table II, a numerical summary of

³github.com/erik-nelson/blam_slam

⁴github.com/HKUST-Aerial-Robotics/A-LOAM

⁵github.com/wh200720041/floam

⁶github.com/cartographer-project/cartographer

⁷github.com/hyye/lio-mapping

⁸github.com/TixiaoShan/LIO-SAM

the RMSE values of the map error (ME) computed for each algorithm on each relevant dataset. The results show similar trends to the position error, with some negligible differences in the order of the algorithms

2) *Robustness Evaluation:* The previous section highlighted the robustness to low lidar observability, substantial vibrations, large accelerations and self-similar environments through the accuracy results on the datasets. In this section we focus on another aspect of robustness: the ability to handle a sudden failure of an input source. Specifically, we test the following scenarios: i) failure of WIO/IMU, ii) failure of WIO, iii) failure of lidar. Each of these failure scenarios have been experienced in real field tests in preparation for the SubT Challenge. We artificially create these failures in our datasets to have a controlled way of isolating the source of the failure, and the resulting impact on the algorithms. The results are summarized in Table III, with example maps resulting from different failure modes shown in Fig 5.

WIO/IMU failure: We simulate sensor failure after 1200s, by shutting down WIO and IMU streams for the rest of the run. This failure only affects those algorithms that use IMU, where the algorithms cease to run, either relying on a synced callback with WIO and IMU (e.g. Cartographer) or relying on pre-integrated IMU to provide odometry updates between scans as well as initial scan to map alignment estimates (LIO-Mapping, LIO-SAM). In contrast, LOCUS processes the input data separately, and hence can automatically switch from lidar odometry with WIO integration, to lidar odometry with IMU integration, to pure lidar odometry, demonstrating robust handling of sensor failures in a cascaded fashion. This behavior is highly desirable to accommodate the unforeseen challenges posed by rough terrains in real-world applications where hardware failures are likely to happen, or sensors sources can become unreliable (e.g. dark areas with no visual texture for VIO).

WIO failure: We simulate a loss of WIO after 1200s. Cartographer and LOCUS are the only algorithms affected, with the same result as the WIO/IMU case.

Lidar failure: We stress the systems by subtracting the most fundamental data source: lidar. More specifically, we simulate a 10s gap in lidar data while the robot is in motion.

There are three responses to this failure. The first response is that the algorithm stops running until the lidar returns, resulting in large map errors (BLAM, ALOAM, FLOAM, and Cartographer due to the synced callback). The second response is that the algorithm runs purely on IMU integration, leading to an accumulation of drift before the lidar returns (LIO-Mapping, LIO-SAM). The final response is only seen by LOCUS, where the loose coupling allows WIO to accumulated over the 10s of no lidar data to produce an accurate initial transform when the lidar returns.

3) *Efficiency Evaluation:* We profile the time needed from the different algorithms to process a single lidar scan when running the algorithms on an Intel Hades Canyon NUC8i7HVKVA (4x1.9 GHz, 32 GB RAM) running Ubuntu 18.04 LTS. Fig. 6 shows the resulting times per scan with scans at 10Hz (LIO-Mapping is omitted as the processing

time, 1s per scan, is too large). Additionally, Table II shows the CPU loads for each algorithm. All values are from the Urban Beta dataset, except for LIO-SAM, which is on the Tunnel Safety Research dataset.

LOCUS, ALOAM and BLAM can all maintain real-time processing, whereas ALOAM and LIO-SAM can only stay real-time with a lower rate of lidar scans. LIO-SAM can use the IMU pre-integration to cope with a lower IMU rate, and by using features, ALOAM can also handle a lower rate for certain datasets. Both FLOAM and LIO-Mapping do not appear to be feasible for real-time operation. Cartographer has both the quickest processing time and the lowest CPU load, yet the accuracy is not as strong as the other algorithms. LOCUS produces the best accuracy, with real-time performance, yet requires the largest CPU load.

C. Real-Time Operation Across Different Platforms

In this section, we demonstrate real-time field operation of LOCUS on different robotic platforms during the Urban Circuit of the SubT Challenge and provide statistics from logged online operation. Results come from the four competition runs, two in Alpha course and two in Beta course.

1) *Hardware and Tuning:* During the competition, we deployed LOCUS on two very different robots: i) the Husky from the datasets used above (see Fig. 1.b), and ii) Spot from Boston Dynamics (see Fig. 1.d).

Husky:

In addition to the sensors described in Sec. III, Husky carries an AMD RYZEN 9 3900X 12-Core 3.8 GHz for computation.

Spot: A legged robot that is equipped with 1 VLP16 and an Intel NUC7i7DN 4-Core 1.9 GHz for computation. Both VIO and KIO are available from the Boston Dynamics API, and can be used for integration into LOCUS. We choose VIO as it was shown to be more accurate than KIO in our tests.

Adaptation: The parameters of LOCUS are tuned to achieve accurate and robust real-time operation on both platforms while accounting for differences in computational capabilities and hardware configurations. Table IV summarizes the configurations used during the competition.

2) *Performance:* We report in Table V the average value of the number of lidar scans dropped per second by each robot in each course of the competition during real-time operation. Point clouds are subscribed to at 10hz and we do not buffer any lidar scans, to minimize the delay of the computed odometry. Therefore, the number of dropped scans per second represent how frequently the lidar processing time exceeded 0.1s. Spot drops 2 scans a second, due to the less powerful computer onboard. However, Spot has a more accurate additional odometry source than Husky, with VIO. The reliable initial transform from VIO enables Spot to process fewer scans per second, and still perform well.

Real-time accuracy profiling: LOCUS performed accurately for both Husky and Spot in the competition, as evident in the overall team's performance, winning first

Fig. 4. Boxplot visualization of the Absolute Position Error (APE) computed for the different methods on the test datasets. For clarity, only the best six algorithms in each dataset are shown.

TABLE II
SUMMARY OF ACCURACY ANALYSIS RESULTS ON ALPHA, BETA AND SAFETY RESEARCH DATASETS

Algorithm	Alpha Course				Beta Course				Safety Research Course				CPU ^β	
	APE [m]			ME [m]	APE [m]			ME [m]	APE [m]			ME [m]	max	mean
LOCUS	1.69	0.62	0.57	0.29	1.51	0.88	0.51	0.69	3.39	1.67	0.76	0.63	3.39	2.72
LOCUS.FGA	0.63	0.26	0.18	0.28	1.20	0.58	0.39	0.48	-	-	-	-	3.39	2.72
BLAM	3.44	1.01	0.94	0.43	3.89	2.27	0.89	1.27	171.34	35.45	51.91	5.37	1.14	0.93
ALOAM	4.33	1.38	1.19	0.60	2.58	2.11	0.44	0.99	18.61	10.01	6.01	6.11	1.65	1.41
FLOAM	29.49	9.19	8.96	1.73*	40.64	3.94	8.42	3.73*	85.31	32.49	25.73	20.16	1.76	1.44
Cartographer	5.84	2.91	1.60	1.05	2.64	1.37	0.67	0.31	50.05	14.31	13.45	14.25	1.75	0.88
LIO-Mapping	2.12	0.99	0.51	0.45	1.60	1.18	0.22	0.61	3.31	1.99	0.55	0.76	1.80	1.53
LIO-SAM	-	-	-	-	-	-	-	-	2.45	1.26	0.58	0.52	2.75 ⁺	2.00 ⁺

* failure leads to a low map error. ^β CPU loads are computed from the Urban Beta dataset, and Tunnel dataset for LIO-SAM (+)

TABLE III
SUMMARY OF ROBUSTNESS TEST RESULTS

Algorithm	Robustness Test Result		
	a) WIO/IMU Fail	b) WIO Fail	c) Lidar Drop
LOCUS	OK	OK	OK
BLAM	NA	NA	Errors
ALOAM	NA	NA	Errors
FLOAM	NA	NA	Errors
Cartographer	Stops	Stops	Errors
LIO-Mapping	Stops	NA	Errors
LIO-SAM	Stops	NA	Errors

OK; negligible degradation in accuracy. NA: Not Applicable - the algorithm does not use that sensor source. Err: Error

Fig. 5. Robustness test in Beta course: a) results on WIO/IMU failure, b) results on WIO failure, c) results on Lidar failure. The failure locations are circled in all cases.

Fig. 6. Comparison of lidar processing time across the different lidar odometry algorithms. The times are the duration for processing a single scan. Top - Urban Beta dataset, Bottom - Tunnel Safety Research dataset. A processing time of 0.1 s indicates real-time performance (10 Hz scans).

SUMMARY OF LOCUS SETTINGS ON DIFFERENT ROBOTS

Parameter	Husky	Spot
Number of lidars	2	1
Voxel Grid Filter leaf size (m)	0.1	None
GICP iterations in scan-to-submap	20	25
GICP number of cores	4	1
Sensor Integration	WIO	VIO

place⁹. Fig. 7, shows the live performance of LOCUS with FGA on the Husky in Beta 2, with a slightly larger error than the post-processed results (on a different computer), yet still highly competitive.

For Spot, LOCUS was run live in a multi-level exploration

⁹The LOCUS output was integrated with a robust odometry aggregator [25], and then fed to a back-end SLAM algorithm [12]

TABLE V
DROPPED LIDAR SCANS FROM REAL-TIME ON-ROBOT TESTS

Robot	Number of dropped scans / s				
	Alpha 1	Alpha 2	Beta 1	Beta 2	Average
Husky	0	0	0	0	0
Spot	2.082	2.205	1.833	2.016	2.034

Fig. 7. Absolute Position Error (APE) of the trajectories estimated by the different methods against ground-truth in Beta course, including the performance when running live in the SubT Challenge (LOCUS_REALTIME).

in the Urban Alpha 2 course. In this run, the mean APE was 0.586 m, and the maximum APE was 2.599 m, which was sufficiently small for scoring in the competition.

D. Discussion

Overall, fusion of additional sensing modalities is crucial to enable accurate operation in such extreme explorations: by relying on a loosely-coupled mechanism, LOCUS is robust to potential failures of sensors, and can achieve improved performance with respect to tightly-coupled approaches in settings where the extrinsic sensors calibration is not ideal. Additionally, by not making assumptions on the environment type, LOCUS can use a larger number of points with respect to feature-based methods during scan registration, and process a greater amount of information at reasonable computational cost by taking advantage of the priors informed by the additional sensing modalities to seed the GICP.

IV. CONCLUSIONS

Achieving accurate lidar odometry in perceptually-challenging conditions can be difficult due to the lack of reliable perceptual features, presence of noisy sensor measurements, and high-rate motions. While integrating additional sensing modalities can help address these challenges, potential sensor failures can have dramatic impacts on the mission outcome if not robustly handled. In this paper, we present a lidar odometry system to enable accurate and resilient ego-motion estimation in challenging real-world scenarios. The proposed system, LOCUS, provides an accurate multi-stage scan matching unit equipped with an health-aware sensor integration module for seamless loose integration of additional sensing modalities. The proposed architecture is adaptable to heterogeneous robotic platforms and is optimized for real-time operation.

We compare LOCUS against state-of-the-art open-source algorithms and demonstrate top-class accuracy in perceptually challenging real-world datasets, top-class computation time and superior robustness to sensor failures, yet with

greater CPU load. Finally, we demonstrate field-proven real-time operation of LOCUS on two different robots involved in fully autonomous exploration of Satsop power plant during the Urban Circuit of the DARPA Subterranean Challenge, where the proposed system was a key component of CoSTAR team's solution that achieved first place.

REFERENCES

- [1] R. Lösch, S. Grehl, M. Donner, C. Buhl, and B. Jung, "Design of an autonomous robot for mapping, navigation, and manipulation in underground mines," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1407–1412.
- [2] J. S. Jennings, G. Whelan, and W. F. Evans, "Cooperative search and rescue with a team of mobile robots," in *1997 8th International Conference on Advanced Robotics. Proceedings. ICAR'97*. IEEE, 1997, pp. 193–200.
- [3] R. Bogue, "Robots for monitoring the environment," *Industrial Robot: An International Journal*, 2011.
- [4] J. Haruyama, T. Morota, S. Kobayashi, S. Sawai, P. G. Lucey, M. Shirao, and M. N. Nishino, "Lunar holes and lava tubes as resources for lunar science and exploration," in *Moon*. Springer, 2012, pp. 139–163.
- [5] A. Agha, *CoSTAR team website*, 2020. [Online]. Available: <https://costar.jpl.nasa.gov/>
- [6] J. Zhang and S. Singh, "LOAM: Lidar odometry and mapping in real-time," in *Robotics: Science and Systems*, vol. 2. IEEE, 2014, p. 9.
- [7] S. S. J. Zhang, "Laser-visual-inertial odometry and mapping with high robustness and low drift," in *Journal of Field Robotics*, 2018, pp. pp. 1242–1264.
- [8] J.-L. Blanco-Claraco, "A modular optimization framework for localization and mapping," in *Robotics: Science and Systems*, 2019.
- [9] C. Le Gentil, T. Vidal-Calleja, and S. Huang, "In2laama: Inertial lidar localization autocalibration and mapping," *IEEE Transactions on Robotics*, 2020.
- [10] M. Bosse and R. Zlot, "Continuous 3d scan-matching with a spinning 2d laser," in *ICRA*, 2009, pp. pp. 4312–4319.
- [11] T. Shan and B. Englot, "LeGO-LOAM: Lightweight and ground-optimized lidar odometry and mapping on variable terrain," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. pp. 4758–4765.
- [12] K. Ebadi, Y. Chang, M. Palieri, A. Stephens, A. Hatteland, E. Heiden, A. Thakur, N. Funabiki, B. Morrell, S. Wood, L. Carlone, and A.-a. Agha-mohammadi, "LAMP: Large-Scale Autonomous Mapping and Positioning for exploration of perceptually-degraded subterranean environments," in *ICRA*. IEEE, 2020.
- [13] P. Biber and W. Straßer, "The normal distributions transform: A new approach to laser scan matching," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 3. IEEE, 2003, pp. 2743–2748.
- [14] A. Segal, D. Haehnel, and S. Thrun, "Generalized-icp," in *Robotics: science and systems*, vol. 2, no. 4. Seattle, WA, 2009, p. 435.
- [15] e. a. S. Agarwal, K. Mierle, "Ceres solver available online.," in <http://ceres-solver.org>.
- [16] G. Grisetti, T. Guadagnino, I. Aloise, M. Colosi, B. Della Corte, and D. Schlegel, "Least squares optimization: from theory to practice," *Robotics*, vol. 9, no. 3, p. 51, July 2020.
- [17] W. Shao, S. Vijayarangan, C. Li, and G. Kantor, "Stereo visual inertial lidar simultaneous localization and mapping," in *IROS*, 2019.
- [18] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *IEEE/RSJ IROS*. IEEE, 2020.
- [19] F. Dellaert, "Factor graphs and gtsam: A hands-on introduction," Georgia Institute of Technology, Tech. Rep., 2012.
- [20] H. Ye, Y. Chen, and M. Liu, "Tightly coupled 3d lidar inertial odometry and mapping," in *ICRA*, 2019, pp. pp. 3144–3150.
- [21] D. Kohler, H. Rapp, and D. Andor, "Real-time loop closure in 2d lidar slam," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. pp. 1271–1278.
- [22] E. Takeuchi and T. Tsubouchi, "A 3-D scan matching using improved 3-D normal distributions transform for mobile robotic mapping," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 3068–3073.
- [23] J. S. Vitter, "Faster methods for random sampling," *Communications of the ACM*, vol. 27, no. 7, pp. 703–718, 1984.

- [24] R. B. Rusu and S. Cousins, “3d is here: Point cloud library (pcl),” in *ICRA*. IEEE, 2011, pp. 1–4.
- [25] A. Santamaria-navarro, R. Thakker, D. D. Fan, B. Morrell, and A. Agha-mohammadi, “Towards resilient autonomous navigation of drones,” in *International Symposium on Robotics Research*, 2019.
- [26] A. Bouman, M. Ginting, N. Alatur, M. Palieri, D. Fan, T. Touma, T. Pailevanian, S. Kim, K. Otsu, J. Burdick, and A. Agha-Mohammadi, “Autonomous spot: Long-range autonomous exploration of extreme environments with legged locomotion,” *IROS 2020*.
- [27] E. Nelson. (2016) Berkley localization and mapping. [Online]. Available: <https://github.com/erik-nelson/blam>
- [28] T. Qin, P. Li, and S. Shen, “Vins-mono: A robust and versatile monocular visual-inertial state estimator,” *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [29] M. Grupp, “evo: Python package for the evaluation of odometry and slam.” <https://github.com/MichaelGrupp/evo>, 2017.