

From SLAM Drift to Goal Drift: Measuring Coordination Sensitivity on Lunar Analogue Terrain

Andreja Malus, Tena Žužek, Jure Dvoršak, Rok Vrabič*
University of Ljubljana, Faculty of Mechanical Engineering
*rok.vrabic@fs.uni-lj.si

Abstract—Multi-robot exploration of lunar terrain requires coordination algorithms that distribute work across the fleet without duplication. Decentralised frontier assignment achieves this by having each robot independently compute the same goal allocation from shared state. However, the approach assumes accurate, consistent localisation, an assumption that breaks in featureless planetary environments where SLAM drift is large and loop closures are scarce. Within the ESA-funded TerraScout project, we deploy a distance-rank frontier deconfliction scheme on 3–5 rovers and test it outdoors and at the LUNA lunar analogue facility. By replaying every coordination decision against graph-optimised poses, we quantify how often localisation error would have changed the robot’s assigned goal. The scheme proves structurally tolerant: because only the single closest uncontested cell determines each assignment, most ranking perturbations are absorbed silently. Yet tolerance has limits: when drift reaches 2–3× the occupancy cell size, up to 41% of goals diverge from the reference solution. Two independent conditions emerge for reliable distributed coordination: position accuracy must remain within cell resolution, and occupancy grids must be exchanged in a common reference frame.

Index Terms—Multi-robot exploration, frontier-based coordination, SLAM, lunar analogue

I. INTRODUCTION

Sustained lunar presence will require autonomous exploration of large, unknown areas for resource prospecting, geological mapping, and site characterisation for future infrastructure [1]. Multi-robot teams offer faster coverage and redundancy over single-rover missions, making them a focal point of current lunar exploration programmes such as NASA’s CADRE [2], [3] and ESA-funded activities targeting long-term multi-rover operations. However, the benefit of a fleet only materialises if robots coordinate effectively to avoid duplicating work.

Frontier-based exploration [4] is the dominant paradigm for autonomous coverage. In decentralised variants, each robot independently computes the same assignment from shared state, and no negotiation or central authority is required [5]. The key assumption is that all robots observe the same, accurate world state. When this holds, a deterministic assignment rule produces identical decisions on every robot.

In practice, the assumption breaks in planetary environments. SLAM localisation accumulates drift in featureless terrain, where loop closures are scarce and visual features are repetitive [6], [7]. Position errors grow with distance travelled, and rotation errors compound further, as even small heading biases produce increasing positional displacements at

distant frontier cells [8]. Meanwhile, map sharing between robots is non-trivial: reference frames must be aligned, and exchange must be live and cumulative rather than a one-time snapshot [6], [9]. What is less often quantified is *how much* these failures degrade coordination decisions, and what accuracy threshold a given coordination mechanism actually requires.

How much error can a coordination mechanism absorb before assignments degrade? A distance-rank rule depends on the single closest uncontested cell, not the full ranking, so many rank perturbations may never propagate to an actual goal change. Quantifying this structural tolerance, and identifying where it runs out, is necessary to specify what a SLAM backend must actually deliver for reliable coordination.

This paper measures how localisation error propagates through a distance-rank frontier deconfliction scheme deployed on a fleet of 3–5 rovers within the ESA-funded TerraScout project (Fig. 1), tested in an outdoor field environment and at the LUNA lunar analogue facility. A post-hoc decision audit replays each coordination decision using graph-optimised SLAM poses, measuring the gap between what the algorithm decided and what it would have decided with accurate inputs. We make three contributions: (1) we quantify how rank perturbations propagate, or fail to propagate, to actual goal changes, revealing structural tolerance in the mechanism; (2) we show that the required localisation accuracy is set by the coordination cell resolution, not by the SLAM specification alone; and (3) we show that pose accuracy and map consistency are independent requirements that must both hold for reliable distributed coordination.



Fig. 1: TerraScout rovers during a coordinated exploration run at the LUNA lunar analogue facility (ESA/DLR, Cologne).

II. SYSTEM AND COORDINATION MECHANISM

A. Platform and Test Environments

The fleet comprises five rovers (Amos, Bobbie, Camina, James, and Naomi) based on the Leo Rover platform (approx. 8.5 kg each), each running an NVIDIA Jetson Orin Nano under ROS2. The primary sensor is an Intel RealSense D456 depth camera. RTAB-Map [9] provides visual-inertial SLAM using RGB-D and IMU data, producing a 6DOF pose and a 3D point cloud. The point cloud is projected into a 2D occupancy grid used for frontier extraction; vox_nav [10] handles path planning. Robots communicate over WiFi via ROS2 DDS, broadcasting at approximately 1 Hz their current pose and selected frontier cell.

Experiments were conducted in two environments. The **outdoor tests** took place in an outdoor field with rich visual features providing favourable conditions for visual SLAM. The full coordination pipeline was operational, with all robots maintaining peer visibility throughout. The **facility tests** were conducted at the LUNA facility in Cologne, Germany (ESA/DLR). LUNA provides 700 m² of basalt-derived volcanic terrain processed to approximate lunar regolith, with craters, boulders, slopes, and simulated lunar illumination. The monochromatic, low-texture surface presents a challenging environment for visual SLAM, as distinctive features are scarce and repetitive.

B. Distance-Rank Frontier Deconfliction

The coordination mechanism distributes frontier cells across the fleet without negotiation. All robots share peer positions and a frontier set. The assignment rule is a greedy distance rank: iteratively, the robot closest to its highest-scoring available cell claims that cell and its 8-cell neighbourhood until every robot has a goal. Because the rule is deterministic and operates on shared inputs, every robot independently computes the same assignment [5]. Correctness therefore requires that all robots observe the same peer positions and the same frontier set.

Grid maps are shared via cell-by-cell synchronisation: a robot updates only cells still unknown in its own map, marking them as traversable or obstacle based on the peer’s data, with own observations taking priority. No coordinate transform is applied before merging, so the mechanism assumes all grids share the same origin. This holds at the start of a run, but each robot’s SLAM operates in its own local frame. As drift accumulates, local frames diverge, and cell indices progressively refer to different physical locations. The merge then places information at wrong grid positions, corrupting the frontier set (Fig. 2).

III. LOCALISATION ERROR AND ITS IMPACT ON COORDINATION

A. Post-Hoc Decision Audit

RTAB-Map stores graph-optimised poses in its SQLite database after pose-graph optimisation and loop closure detection, representing the best achievable trajectory estimate from




Fig. 2: Distance-rank frontier deconfliction. (a) With accurate state, both robots agree on frontier cell c and Robot A, being closer, claims it. (b) With SLAM drift, robot positions shift (dashed arrows) and each robot’s occupancy grid drifts independently, so they observe the frontier at different grid cells (c_A , c_B) and compute independent, incompatible assignments.

the recorded sensor data. We extract these poses, apply SE(2) alignment transforms to bring all robots into a common global frame, and interpolate to each decision tick timestamp. Self-localisation error is the 2D Euclidean distance between the live pose used at decision time and the corresponding optimised pose.

At each tick we replay the deconfliction rule using optimised poses for the robot and all peers. We record two metrics: *rank changed*, indicating that any frontier cell’s distance ordering changed between the live and optimised reconstructions, and *goal changed*, indicating that the robot’s claimed goal cell differed from what it would have been with optimised poses. The first measures sensitivity of the distance computation; the second measures whether that sensitivity propagated to an actual assignment change.

B. Results

Table I summarises the coordination pipeline across all scenarios. The pipeline proceeds left to right: frontier cells are first reduced by *cluster selection* (choosing the best-scoring connected frontier region), then by *peer-based deconfliction* (removing cells claimed by closer peers). The remaining columns report rank and goal change rates from the post-hoc audit. The gap between the two reveals the mechanism’s structural tolerance to perturbation.

Outdoor tests (OR1, OR2). The full coordination pipeline was operational: all robots maintained peer visibility throughout. Mean self-localisation errors were 0.11 m to 0.29 m, well below the cell size. Even so, 27–52% of distance rankings changed. But only 0–13% of goal assignments differed from what optimised poses would have produced. Most rank perturbations did not propagate to actual goal changes, because only the single closest uncontested cell matters.

Facility tests (FR1–FR3). At the LUNA facility, the peer-based deconfliction stage was not operational, so the analysis covers only pose-error sensitivity. Pose errors were an order of magnitude larger than outdoors (Fig. 3), ranging from mean 0.8 m to 1.3 m in FR1 to 2.5 m to 3.3 m in FR3. The tolerance observed outdoors degrades with error magnitude but does not disappear uniformly. In FR1, one robot with mean error 0.84 m

TABLE I: Coordination pipeline and sensitivity across all scenarios. All values except N and Err. are percentages of frontier cells or decision ticks affected.

	N	Err. (m)	Clust. rem.	Peer rem.	Rank chg.	Goal chg.
<i>Outdoor</i>						
OR14		0.11–0.17	44–58%	22–36%	27–40%	0–8%
OR23		0.25–0.29	51–84%	6–44%	31–52%	1–13%
<i>Facility</i>						
FR1 5		0.84–1.30	75–82%	—	34–86%	0–24%
FR2 5		1.45–2.43	52–71%	—	40–77%	12–24%
FR3 5		2.47–3.27	48–83%	—	54–95%	13–41%

OR1–OR2: outdoor field, 5–24 min. FR1–FR3: LUNA facility, 5 robots, 5–15 min. All runs use 0.80 m to 1.0 m cells.

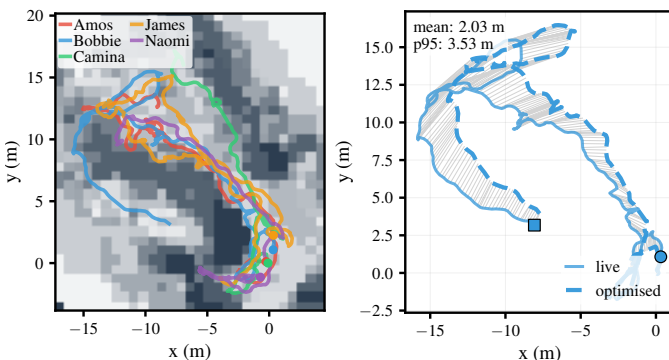


Fig. 3: Left: all-robot live trajectories overlaid on the shared occupancy grid (FR2, 5 robots). Right: live (solid) vs. graph-optimised (dashed) trajectory for *Bobbie* in FR2, with grey lines showing per-tick displacement (mean error 2.03 m, p95 3.53 m).

on a 0.80 m cell grid recorded zero goal changes, the only case where localisation accuracy was in proportion with cell resolution. In FR3, errors of nearly three cell widths produced up to 41% goal changes (Fig. 4).

Compounding of position drift and grid misalignment.

Both inputs to the deconfliction rule are independently degraded: peer positions by SLAM drift, and frontier sets by the grid synchronisation mechanism described in Section II-B. As local reference frames diverge, the merged grid accumulates misplaced cells, inflating the frontier set with cells that have already been explored by peers.

IV. CONCLUSIONS

We measured how SLAM localisation error propagates through a distance-rank frontier deconfliction scheme, using data from outdoor field tests and the LUNA lunar analogue facility. The results lead to three findings.

The mechanism has structural tolerance to localisation error. Because the assignment depends on the single closest uncontested frontier cell, most rank perturbations are absorbed without affecting the goal. In outdoor tests, 70–100% of rank

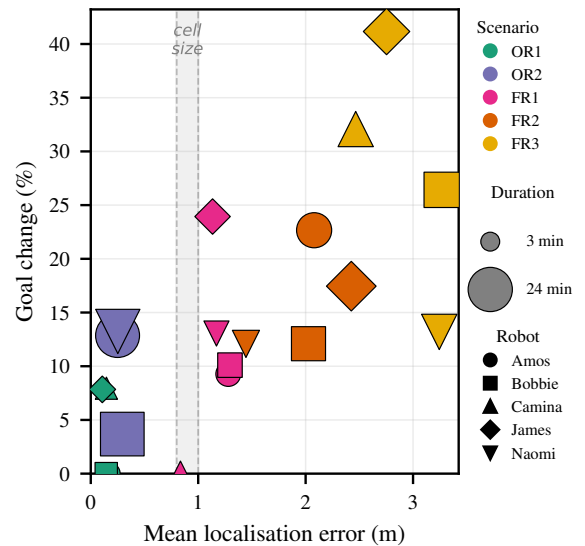


Fig. 4: Per-robot mean self-localisation error vs. goal-cell change fraction (OR1–FR3). Marker size scales with robot active duration; shape encodes robot identity, colour encodes scenario. The grey band marks the 0.80 m to 1.0 m cell size.

reorderings did not change the claimed goal. This tolerance is bounded: at $2\text{--}3\times$ cell size, up to 41% of assignments are affected.

Coordination resolution sets the localisation requirement. Localisation error must remain below the occupancy grid cell size for reliable assignments. This is a cross-layer requirement: it is not captured by the SLAM specification alone, and it must hold throughout the run, not just at initialisation. Rotation errors can further compound this effect, as heading biases produce growing displacements at distant frontier cells.

Pose accuracy and map consistency are independent, co-required conditions. Even with perfect localisation, occupancy grids synchronised without frame correction will diverge as SLAM drift accumulates, causing robots to deconflict over misaligned frontier sets.

Limitations. Peer deconfliction was disabled at LUNA to record a pose-only baseline first; operational constraints in the campaign window prevented re-enabling it before the runs ended. The absent grid-misalignment effect would only compound the reported rates. Loop-closure analysis indicates residual error in the graph-optimised reference of approximately 0.26 m (p95, 140 loop closures) at LUNA, well below cell size. With five scenarios we report per-robot ranges rather than aggregate confidence intervals; reported goal-change rates are therefore conservative.

Future work. Integration of a multi-robot SLAM backend [11] is in progress, addressing pose accuracy via shared loop closures and map consistency via graph-based exchange. Quantifying the downstream cost of misassignments and the scaling of structural tolerance to larger fleets remain open.

ACKNOWLEDGEMENTS

The authors thank the European Space Agency and the LUNA team for the invitation and support. This research was funded by the European Space Agency project TerraScout (grant nb. 4000143611/24/NL/MH/mp) and the Slovenian Research and Innovation Agency (grant nb. P2-0270).

REFERENCES

- [1] B. Martinez Rocamora, C. Kilic, C. Tatsch, G. A. S. Pereira, and J. N. Gross, "Multi-robot cooperation for lunar in-situ resource utilization," *Frontiers in Robotics and AI*, vol. 10, p. 1149080, 2023.
- [2] S. Nayak, G. Lim, F. Rossi, M. Otte, and J.-P. De La Croix, "Multi-robot exploration for the CADRE mission," *Autonomous Robots*, vol. 49, no. 2, p. 17, 2025.
- [3] J.-P. De La Croix, F. Rossi, R. Brockers, D. Aguilar, K. Albee, E. Boroson, A. Cauligi, J. Delaune, R. Hewitt, D. Kogan, G. Lim, B. Morrell, Y. Nakka, V. Nguyen, P. Proença, G. Rabideau, J. Russino, M. S. Da Silva, G. Zohar, and S. Comandur, "Multi-agent autonomy for space exploration on the CADRE lunar technology demonstration," in *Proceedings of the IEEE Aerospace Conference, 2024*, pp. 1–14.
- [4] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation, 1997*, pp. 146–151.
- [5] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 376–386, 2005.
- [6] P.-Y. Lajoie, K. Soma, H. M. Bong, A. Lemieux-Bourque, R. Zhang, V. S. Varadharajan, and G. Beltrame, "Multirobot decentralized collaborative SLAM in planetary analogue environments: Dataset, challenges, and lessons learned," *IEEE Transactions on Field Robotics*, vol. 2, pp. 516–528, 2025.
- [7] K. Ebadi, L. Bernreiter, H. Biggie, G. Catt, Y. Chang, A. Chatterjee, C. E. Denniston, S.-P. Deschênes, K. G  linas, P. H  naff *et al.*, "Present and future of SLAM in extreme environments: The DARPA SubT challenge," *IEEE Transactions on Robotics*, vol. 40, pp. 936–959, 2024.
- [8] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart, "Distributed multirobot exploration and mapping," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1325–1339, 2006.
- [9] M. Labb   and F. Michaud, "RTAB-Map as an open-source lidar and visual slam library for large-scale and long-term online operation," *Journal of Field Robotics*, vol. 36, no. 2, pp. 416–446, 2019.
- [10] F. Atas, G. Cielniak, and L. Grimstad, "Elevation state-space: Surfel-based navigation in uneven environments for mobile robots," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022*, pp. 5715–5721.
- [11] X. Liu, J. Lei, A. Prabhu, Y. Tao, I. Spasojevic, P. Chaudhari, N. Atanasov, and V. Kumar, "SlideSLAM: Sparse, lightweight, decentralized metric-semantic SLAM for multi-robot navigation," *IEEE Transactions on Robotics*, vol. 41, pp. 6529–6548, 2025.