2025 I-RIM Conference
October 17-19, Rome, Italy
ISBN: 9788894580570
10.5281/zenodo.17629902

Mixed-Reality Interface for Robot Teams

Omotoye Shamsudeen Adekoya
DIBRIS, University of Genoa

Antonio Sgorbissa
DIBRIS, University of Genoa

Carmine Recchiuto
DIBRIS, University of Genoa

***Abstract*—We present an update on *HORUS* (Holistic Operational Reality for Unified Systems), a mixed-reality (MR) interface for managing heterogeneous robot teams from a Meta Quest 3 headset. HORUS anchors robot state, plans, sensor data, and controls in a shared mini-map ("Ground Station") with per-robot *Robot Manager* panels and a semi-immersive view for precise teleoperation. Relative to the prior wheeled-robot study, we (i) move to ROS 2 for local/remote operation, (ii) replace post-hoc map merging with multi-robot SLAM maintaining a single shared occupancy map, (iii) add a usable *legged* integration (Boston Dynamics Spot) with a body/stance panel while preserving the wheeled workflow, (iv) implement an *aerial* pathway flight-tested indoors, and (v) define an *ongoing* LLM planner copilot that communicates plan *intent* via concise text and lightweight 3D highlights. We restate key numbers from the wheeled phase.**

***Index Terms*—Mixed Reality, multi-robot systems, human–robot interaction, supervisory control**

I. Background and Scope

Heterogeneous teams improve coverage and robustness in cluttered, partially known environments, but conventional 2D dashboards fragment situational awareness across maps, camera feeds, and status widgets. MR shifts that burden by *spatially anchoring* data and controls in the operator's physical context so the "where" and "what" co-locate [1]. Prior MR tools demonstrate benefits for single-robot visualization/control (e.g., ARviz and iViz) [2], [3] and explore multi-robot tasking [4]. Multi-robot mapping is trending toward shared-map formulations [5].

HORUS at a glance. HORUS renders a shared mini-map of all robots (Fig. 1). Selecting a robot opens its *Robot Manager* with four tabs: *Status* (battery, pose, task), *Data Viz* (LiDAR with proximity cues; projected/overhead video), *Tasks* (goals, waypoints, labeled poses, sketch-to-path), and *Teleop* (mini-map driving and a semi-immersive camera view).

Device, networking, and stack. The front end runs on Meta Quest 3 (Unity, Meta XR SDK). Communication uses ROS 2 locally and across sites via a mesh VPN, preserving identical workflows. The *HORUS Bridge* (ROS↔Unity) translates UI actions (*set goal*, *add waypoint*, *draw path*) into ROS services/actions and rate-limits overlays. Multi-robot SLAM maintains a *single* shared occupancy map with per-robot localization (replacing earlier map merging), simplifying planning overlays and operator reasoning [5]. Navigation employs TEB for reactive trajectory optimization [6]. *Availability:* compiled application releases are open-source[1]; an SDK for integrating additional robot backends is under active development[2].

[1]https://github.com/RICE-unige/horus
[2]https://github.com/RICE-unige/horus_sdk

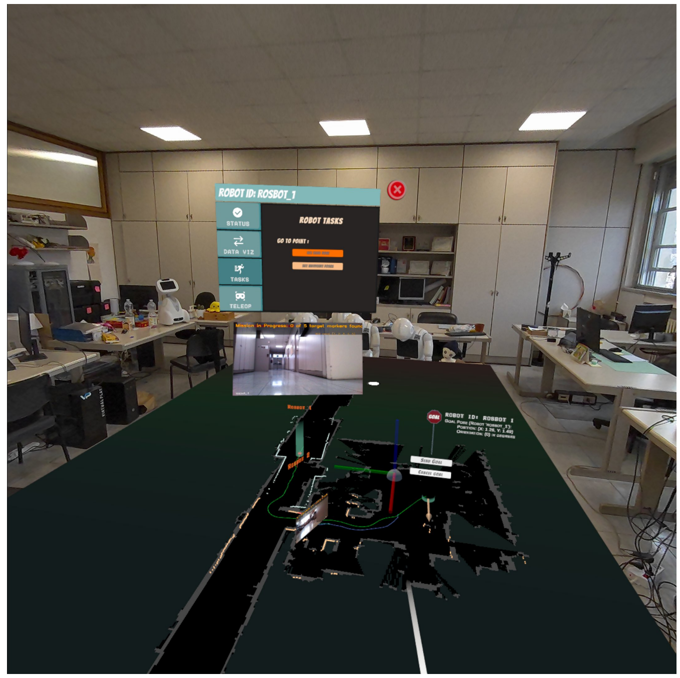

Fig. 1. **Interface overview (Fig. 1).** Shared mini-map with per-robot panels and camera modes.

II. System Architecture and Experiment Setup

A. Architecture used in the wheeled-robot study

Fig. 2 summarizes the evaluated stack. Each ROSbot 2.0 publishes odometry and 2D LiDAR; a multi-robot SLAM setup maintains a *single* shared occupancy map consumed by each robot's navigation. The *HORUS Bridge* translates Unity UI actions to ROS 2, throttles video/scan overlays, and maintains per-robot registration. The Quest 3 runs untethered; a mesh VPN connects remote operators and robots with unchanged MR workflows.

B. Task and measures

Participants (**n=20**; two groups of 10) controlled two robots to find five ArUco tags in offices along a corridor. With HORUS, users combined high-level tasking (incremental goals, waypoints) and targeted mini-map teleop near doorways; a teleop-only baseline offered joystick velocity control and a large camera view. We measured completion time, SUS, NASA-TLX, and pre/post SSQ. HORUS achieved **8:42** vs. **11:17** (p<0.001, d=1.93); SUS averaged **82.3** ("excellent") vs. **68.5** ("good"), and TLX *frustration* was lower (**1.6** vs.

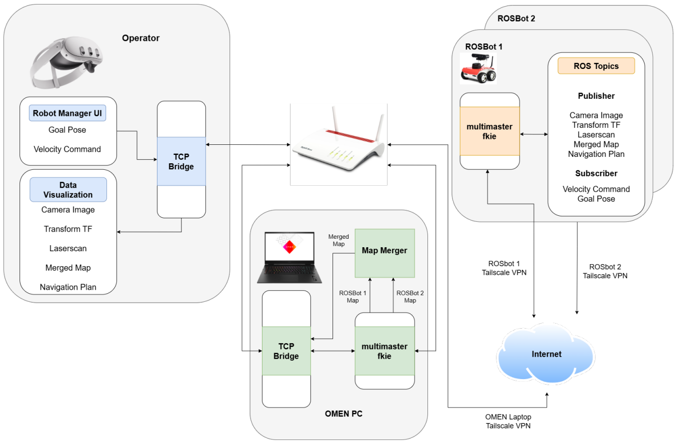

Fig. 2. **Experiment architecture.** Multi-robot SLAM (shared occupancy map; per-robot localization) feeds each robot's TEB-based navigation [6]. HORUS Bridge (ROS 2↔Unity) synchronizes goals/waypoints and overlays.

TABLE I
RESULTS SNAPSHOT (TWO-ROBOT CORRIDOR SEARCH; n=10+10)

Metric	HORUS	Teleop-only
Completion time	8:42 (SD 62.3s)	11:17 (SD 87.5s)
SUS score	82.3 (SD 8.2)	68.5 (SD 11.7)
TLX Frustration	1.6 (SD 0.70)	2.5 (SD 1.43)
Training time	∼18 min	∼7 min

Time: $p<0.001$, Cohen's $d=1.93$.

2.5). Training took ∼18 min for HORUS vs. 7 min for teleop-only, yet participants rated HORUS more usable.

III. Updates for IRIM 2025

A. Legged robot support (Boston Dynamics Spot)

Motivation. Extend from level floors to stairs/uneven terrain without changing operator habits.
Minimal deltas. A *legged adapter* in the bridge (i) registers the agent with stance-aware limits and safe stops; (ii) standardizes status/camera topics so the same four tabs work unchanged; (iii) maps goal/waypoint actions to the legged navigation API.
Current capability. Spot appears as a first-class teammate: pose on the mini-map, overhead video in *Data Viz*, point-to-point goal dispatch, velocity teleop, and a *body/stance* panel for posture/stability (Fig. 3). **Level-floor operation is fully implemented; with 3D map visualization (in progress), we will run participant studies on multi-floor operation.**

B. Aerial robot pathway (status and testing)

A UAV pathway with 3D goal widgets, headset-controller teleop, and video streaming is implemented and has been flight-tested indoors; Remote operation across sites via VPN/ROS 2 is the next milestone.

C. LLM planner copilot (ongoing; no experiments yet)

Goal: a copilot that accepts natural-language missions, decomposes them into tasks, proposes allocations over heterogeneous robots, and surfaces *intent* in MR for confirmation/editing. Current status: architecture defined; HORUS hooks render short rationales and highlight affected rooms/paths; operator edits feed back as constraints. Planned evaluation will measure plan-validation time, intervention frequency, and perceived clarity.

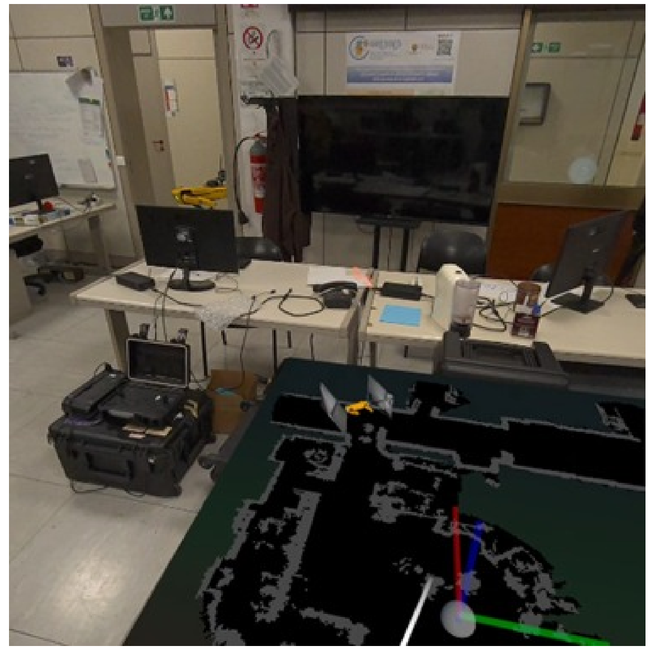

Fig. 3. **Spot in the UI.** Legged agent as a peer in the mini-map with the same Robot Manager tabs and video streaming camera frame.

IV. Discussion and Future Work

Planned work centers on headset-native **3D map visualization** for both legged and aerial agents, multi-level context for Spot and volumetric situational awareness for the UAV. **Remote management experiments** across sites (ROS 2 over VPN) with heterogeneous teams assess scalability, latency tolerance, and operator workload. **LLM copilot**: integrate planning, concise intent summaries, and editable constraints with uncertainty cues; validate with a focused user study on decision speed and trust.

References

[1] J. J. Roldán, E. Peña Tapia, A. Martín-Barrio, M. A. Olivares-Méndez, J. Del Cerro, and A. Barrientos, "Multi-robot interfaces and operator situational awareness: Study of the impact of immersion and prediction," *Sensors*, vol. 17, p. 1720, 2017.
[2] K. C. Hoang, W. P. Chan, S. Lay, A. Cosgun, and E. A. Croft, "ARviz: An augmented reality-enabled visualization platform for ROS applications," *IEEE Robot. Autom. Mag.*, vol. 29, pp. 58–67, 2022.
[3] A. Zea and U. D. Hanebeck, "iviz: A ROS visualization app for mobile devices," *Softw. Impacts*, vol. 8, p. 100057, 2021.
[4] F. Kennel-Maushart, R. Poranne, and S. Coros, "Interacting with multi-robot systems via mixed reality," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2023.
[5] J. Zhang *et al.*, "Overview of multi-robot collaborative slam from the perspective of data fusion," *Machines*, vol. 11, no. 6, p. 653, 2023.
[6] C. Rösmann, F. Hoffmann, and T. Bertram, "Integrated online trajectory planning and optimization in distinctive topologies," *Robotics and Autonomous Systems*, vol. 88, pp. 142–153, 2017.