SafeMove: Multi-Model Navigation for Fail-Resistant Autonomous Nuclear Material Transport – 25556

David Marquez-Gamez[1], Matt Harrison[1], Christian Pilon[1], Robert Marwood[1], Samuel Stephens[1], Michal Staniaszek[2], Tobit Flatscher[2], Chris Prahacs[2], Nick Hawes[2], and Maurice Fallon[2]
[1]AtkinsRéalis
[2]University of Oxford

ABSTRACT

The nuclear industry requires safe and efficient transportation of radioactive materials, especially in decommissioning and waste management. As waste package movements increase from monthly to daily frequencies, traditional manual methods have become unsustainable, exposing operators to hazardous environments and creating potential operational bottlenecks. This paper introduces SafeMove, an advanced autonomous transportation system developed by AtkinsRéalis and the Oxford Robotics Institute, designed to address these critical issues.

SafeMove integrates a vehicle-agnostic hardware and software autonomy stack with a sophisticated multi-sensor fusion system that combines LiDAR, cameras, and inertial sensors. This fusion enables robust localization, detailed 3D mapping, and real-time mission planning, allowing the system to navigate without reliance on GPS or extensive infrastructure modifications. Key functionalities include dynamic obstacle detection, point-to-point mission execution, and risk-aware replanning, ensuring adaptability to complex and evolving conditions within nuclear facilities. Tested in environments analogous to nuclear sites, the system demonstrated exceptional reliability, safety, and operational efficiency. To ensure readiness for deployment in active nuclear facilities, a comprehensive risk assessment based on ISO 26262 standards was conducted.

This paper details the development process, testing outcomes, and deployment considerations of SafeMove, highlighting its potential to revolutionize nuclear waste management. Its implementation offers transformative benefits for the nuclear sector, including significant reductions in human exposure to hazardous conditions, improved operational throughput, and enhanced cost efficiency. Additionally, the system’s adaptability makes it suitable for a wide range of applications, from waste transportation to autonomous inspection and monitoring tasks. By addressing the industry's need for safe, sustainable, and technologically advanced solutions, SafeMove exemplifies the role of robotics and AI in shaping the future of nuclear waste management and decommissioning activities worldwide.

INTRODUCTION

Transporting radioactive materials safely and efficiently is a critical challenge for the nuclear industry, especially in decommissioning and waste management activities. Increasing demands for nuclear waste transportation, coupled with the hazardous nature of these environments, require innovative solutions that prioritize safety, operational efficiency, and scalability. Traditional manual handling methods are increasingly unsustainable, exposing operators to significant risks and creating operational bottlenecks. This urgency drives the need for autonomous systems capable of revolutionizing operations in nuclear facilities. SafeMove addresses these challenges with a scalable and adaptable autonomous vehicle solution designed to enhance safety, reduce human exposure, and improve cost efficiency in nuclear waste transportation. By integrating advanced robotics and artificial intelligence, SafeMove offers capabilities that align with the nuclear industry’s goals of safe, environmentally responsible, and innovative waste management practices.

Autonomous vehicles have emerged as a transformative technology with the potential to address these challenges. By reducing human involvement in hazardous environments, these systems enhance safety while increasing operational throughput and cost efficiency. However, the deployment of autonomous vehicles in nuclear facilities presents unique challenges, including:

- Navigating GPS-Denied Environments: Nuclear facilities often lack GPS coverage due to shielded structures and indoor settings.
- Infrastructure Compatibility: Autonomous systems must adapt to existing infrastructure without extensive modifications.
- Dynamic Conditions: The ability to detect and respond to moving objects, such as pedestrians and vehicles, is critical for safe operations.
- Regulatory and Safety Standards: Compliance with stringent safety protocols and regulations is paramount in nuclear environments.

The use of autonomous vehicles in nuclear facilities remains in its early stages, with applications primarily focused on inspection and monitoring. For instance, Boston Dynamics' Spot robot has been employed for remote inspections, leveraging cameras and sensors to collect real-time data in hazardous areas [1]. Similarly, autonomous drones are being used for mapping radiation levels in inaccessible zones [2]. While these technologies have demonstrated value, they are limited in scope and primarily focus on data collection rather than material handling or transportation. Existing autonomous transportation systems outside the nuclear sector, such as warehouse robots and autonomous forklifts, often rely on GPS or predefined markers for navigation [3]. These solutions are not directly transferable to nuclear facilities due to the lack of GPS coverage and the need for robust obstacle detection and dynamic planning. Additionally, many current systems require significant infrastructure modifications, such as the installation of fiducial markers or specialized pathways, which are impractical in nuclear sites. This gap highlights the need for a robust and adaptable system specifically designed to address the safety, operational, and scalability demands of nuclear environments.

SafeMove bridges this gap by integrating an advanced hardware and software autonomy stack capable of transforming any electrical vehicle into a fail-resistant autonomous platform. At the core of SafeMove’s functionality is the ORI-AutoNav-System, which delivers robust localization, mapping, and mission planning through a vehicle-agnostic approach. This includes GPS-free navigation, real-time graph-based optimization, and topological mission planning that enables navigation along pre-planned routes, with real-time obstacle avoidance and risk-aware replanning. A scientific publication overviewing this work was published in a research robotics conference [4]. That publication focused more specifically on integration of the system on a walking robot operating in an indoor industrial inspection context – further demonstrating the versatility of the autonomy system.

The system uses a modular architecture that combines dynamic perception, adaptive planning, precise localization, and real-time sensing, enabling seamless operation in nuclear environments. Its perception capabilities allow it to detect, track, and predict the movements of pedestrians, vehicles, and other dynamic objects, maintaining situational awareness for informed decision-making. The planning module dynamically manages mission execution, handling tasks such as driving and parking while adapting to changing environmental conditions. Localization provides continuous feedback by integrating data from high-resolution sensors, including vision cameras, LiDAR, inertial measurement units (IMU), and digital data sharing (DDS) implementation, which supports precise positioning and robust navigation. Together, these capabilities enable SafeMove to operate without reliance on GPS or extensive infrastructure modifications. Additionally, SafeMove can be integrated into COTS (Commercial Off-The-Shelf) platforms based on operational requirements such as load capacity, speed, and terrain. This flexibility allows the system to adapt to diverse operational conditions, transforming standard electrical vehicles into autonomous agents capable of executing complex missions in hazardous environments. The system has been demonstrated across various scenarios, validating its ability to navigate large, dynamic sites efficiently while maintaining fail-resistant performance.

SafeMove's development directly addresses the critical challenges associated with deploying autonomous vehicles in nuclear facilities. These include:

- Dynamic Obstacle Management: Real-time perception and planning ensure the detection and avoidance of moving objects.
- GPS-Independent Navigation: Integration of LiDAR, graph-based mapping, and inertial sensors enables SafeMove to navigate effectively in GPS-denied environments.
- Risk-Aware Replanning: Adaptive algorithms manage battery performance, obstacles, and safety risks during mission execution.
- Infrastructure Compatibility: The modular design enables seamless integration with existing vehicles and site configurations.

This paper explores SafeMove's development and deployment, focusing on how its innovations overcome the challenges of autonomous vehicle integration in nuclear facilities. By addressing critical operational and safety needs, SafeMove represents a significant step forward in enhancing the efficiency, safety, and sustainability of nuclear waste management operations.

SYSTEM ARCHITECTURE

To achieve the goal of autonomous material transportation, SafeMove leverages a modular system architecture that integrates key functional components. These components include mapping, localization, perception, planning, and sensing capabilities, which collectively enable robust autonomous operation in dynamic and complex environments. Figure 1 illustrates the overall system architecture, highlighting the modular design that facilitates seamless integration of SafeMove's components (shown in blue) with a wide range of control systems and vehicle interfaces. This architecture supports vehicle-agnostic deployment, ensuring compatibility with diverse platforms and operational conditions. The core modules handle mapping and Localization, object detection and tracking and mission planning, while interfacing with external vehicle-specific systems for trajectory execution and control.

SAFEMOVE PAYLOAD – HARDWARE

The core hardware component of the SafeMove system is the Frontier device (Figure 2). The Frontier integrates a NUC mini-PC with an Intel i7-1165G7 processor running at 2.8 GHz and 32 GB of RAM, a Hesai XT-32 LiDAR, three time-synchronized Sevensense Alphasense fisheye cameras, and a Bosch BMI085 inertial measurement unit (IMU). These components are housed in a lightweight, compact, 3D-printed enclosure with a total weight of 1.5 kg.

The Frontier is integral to SafeMove's functionality, as it runs the AutoNav system, the autonomous navigation software stack that powers SafeMove's perception, localization, and mission execution capabilities. The sensors on the Frontier are precisely calibrated to ensure accuracy and consistency. Camera intrinsics and extrinsics are calibrated using Kalibr [5], and the camera-to-LiDAR extrinsics are calibrated following established methodologies [6]. Since all sensors are rigidly mounted, the device can be transferred between vehicles without requiring recalibration, providing flexibility for deployment across different platforms. The AutoNav system operates entirely on the Frontier, eliminating the need for continuous network connectivity except when receiving operator commands. Operators monitor the system remotely using an external computer, ensuring oversight and control during operations.

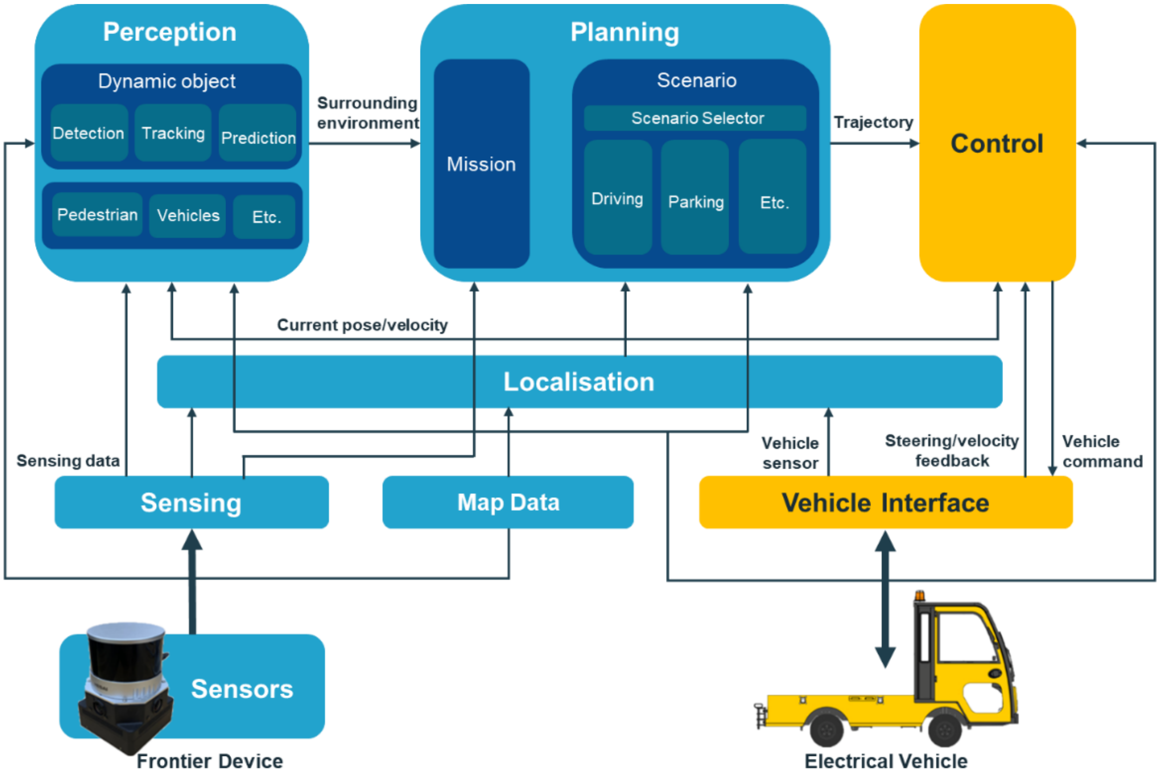

Figure 1. Autonomous Vehicle System Architecture. The architecture includes SafeMove's core modules (shown in blue), such as mapping, localization, perception, and planning, which integrate seamlessly with various control systems and vehicle interfaces. This modular design ensures vehicle-agnostic deployment and adaptability to diverse environments.

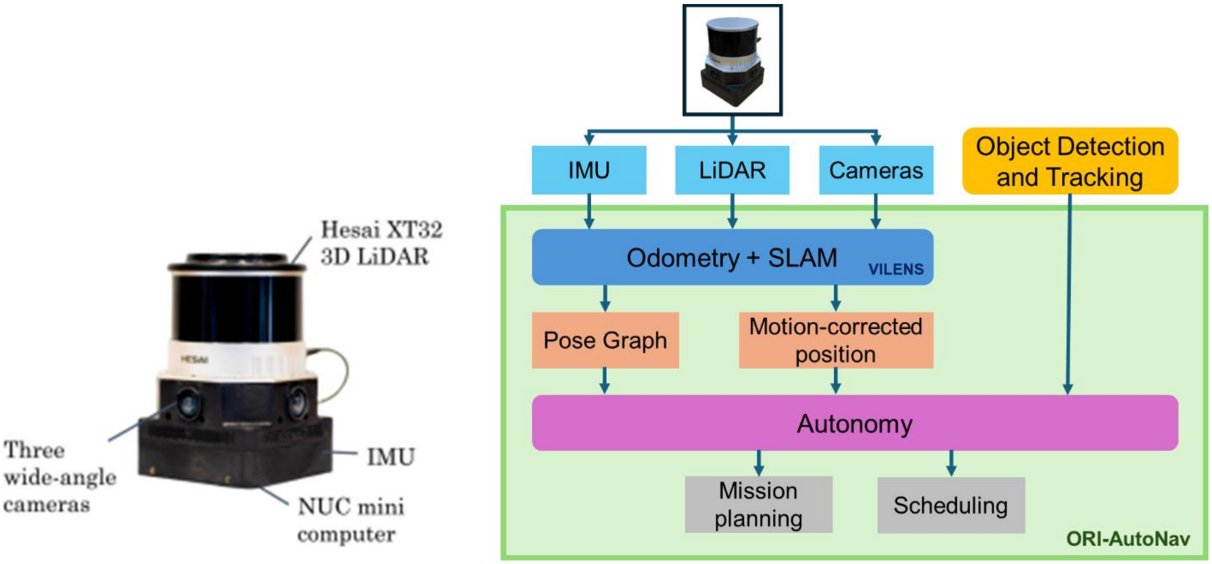

Figure 2. (Left) SafeMove payload, the Frontier device, consisting of a Hesai LiDAR, three Sevensense Alphasense cameras, and an IMU. (Right) Overview of the ORI AutoNav system, using the Frontier device: IMU measurements, visual and LiDAR data from Frontier, are used by VILENS to generate a continuous pose estimate of the vehicle and a SLAM pose graph. These inputs are used by the topological autonomy system to provide mission planning, and scheduling capabilities to the operator.

MAPPING AND LOCALIZATION

The mapping and localization system uses the VILENS odometry [7] and SLAM systems [8], which fuse data from the IMU, and LiDAR to provide robust and accurate odometry. To ensure mapping fidelity, the system compensates for LiDAR motion during each scanning sweep. This correction is critical to counteract the effects of dynamic vehicle motion, enabling the generation of high-resolution, accurate maps required for safe and efficient autonomous operation. The VILENS system has been successfully demonstrated in various challenging environments, including construction [9], forestry [10], and aerial inspections with drones [11].

Initial SLAM Mapping

To enable autonomous operation, the system localizes within a pre-built map of the environment. This map can be generated from existing 3D LiDAR scans acquired with a terrestrial LiDAR scanner or created from scratch using our SLAM system. When generating a map from scratch, odometry is provided by VILENS [7], which can be configured to utilize various odometry sources, including IMU, visual feature tracking, LiDAR ICP registration and produces motion-corrected LiDAR scans. The motion-corrected LiDAR scans are processed by VILENS SLAM [8], which employs the iSAM2 solver [12] to perform pose graph optimization. Loop closure proposals are generated based on geometric constraints and place recognition. Place recognition utilizes the ScanContext descriptor [13] to identify previously visited locations. An overview of the system is presented in Figure 3. The output of the mapping step is a pose graph with associated individual pointclouds, as well as a global map in which all individual pointclouds have been registered in a global reference frame.

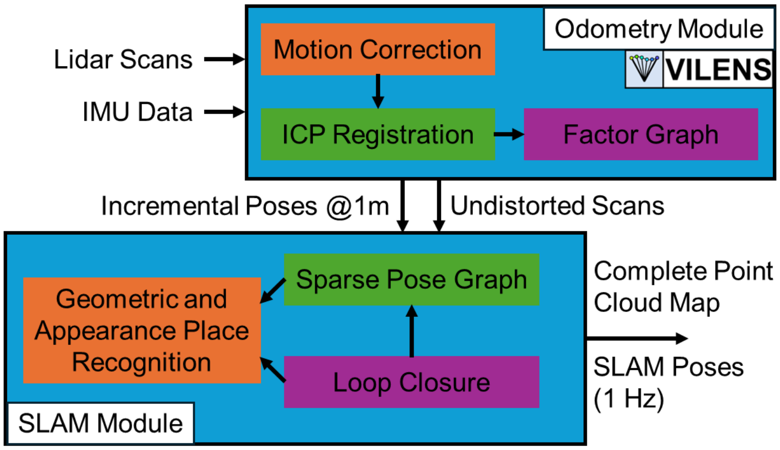

Figure 3. Overview of the VILENS odometry and VILENS SLAM system.

Subsequent Localization in Prior Map

To re-localize within a global prior map, the system utilizes ICP alignment, which requires an accurate initial pose estimate. This estimate can be provided by the operator or achieved by initializing the robot in a known location. The pose estimate is then iteratively refined at a rate of 2 Hz, aligning LiDAR data to the prior map through ICP. Alternatively, localization can be performed using a prior map composed of individual pose-graph pointclouds. In this approach, place recognition is employed to determine the initial pose estimate for each iteration. This method does not rely on an explicit initial pose and is particularly advantageous in large-scale environments when performing ICP directly on a single large pointcloud would be computationally prohibitive.

OBJECT DETECTION AND TRACKING

Dynamic or moving objects within the operational environment of an autonomous vehicle pose significant hazards. While the 3D mapping and localization subsystem ensures accurate positioning to prevent collisions, it does not classify objects or track their movements. To address this limitation, SafeMove integrates an object detection and tracking system designed to identify, classify, and monitor objects in real time.

Object Detection

Several object detection methods were evaluated during the development of SafeMove, including both 2D and 3D detection approaches. A 2D object detection framework was selected for its optimal balance between computational efficiency, detection accuracy, and real-time performance. The D-FINE network was chosen due to its superior latency and reliability. Tested on standard PC hardware, the system achieved detection latencies as low as 30 milliseconds, making it suitable for real-time deployment. The framework is also compatible with edge computing platforms, providing versatility for a wide range of hardware configurations. While 3D object detection offers the potential for richer spatial information, it remains a relatively nascent methodology with less mature solutions than 2D detection. As 3D technologies advance, they are expected to surpass 2D methods in terms of convenience and performance, presenting opportunities for future enhancements.

Object Tracking

The system employs the ByteTracker algorithm to track objects across sequential frames. This algorithm effectively links detections over time, assigning unique identifiers to individual objects and monitoring their trajectories. ByteTracker was selected for its ability to balance computational speed and tracking accuracy, ensuring reliable performance in dynamic environments.

3D Projection

To enhance situational awareness, 2D object detections are re-projected into 3D space using data from SafeMove’s multi-sensor array. This capability allows for precise localization of objects in the global environment, facilitating more informed decision-making and safer navigation in complex scenarios

AUTONOMY

A key component of the SafeMove system is its autonomy subsystem, which integrates topological mapping, topological navigation, mission planning and scheduling, and a user interface, forming the foundation for autonomous navigation and task execution (Figure 4).

At the core of the autonomy subsystem is a topological map representation, which provides a modular and efficient system for navigation and mission execution. Combined with the mapping and localization subsystem, this representation enables the vehicle to operate in its environment without the need for continuous monitoring or input from operators. Topological maps are particularly effective for representing large physical spaces [14] and have been extensively used as a navigation abstraction for mobile robots since their introduction by Brooks [15]. These maps are well-suited for incorporating domain knowledge during deployments and adapting the system’s behavior based on specific operational needs. Additionally, they support advanced planning and resource allocation algorithms, which are critical for mission scheduling and optimization. Topological maps have demonstrated their utility in both field deployments, such as in office environments [16] and agricultural applications [17], and in theoretical research [18].

Beyond their technical advantages, topological maps offer significant usability benefits. They provide an intuitive visual representation of the vehicle operational area, allowing end users to easily understand and interact with the system. Features such as naming locations within the map enhance communication about missions, facilitating efficient coordination. Topological maps can be constructed using various methods, including aerial imagery [19] and 2D or 3D maps [20]. SafeMove employs a hybrid approach, automatically generating an initial map from the SLAM pose graph during global map construction and then tailoring it manually using graphical user interface (GUI) tools.

Integrating a new platform with the autonomy subsystem involves implementing a localizer for the graph and an edge traversal interface. This modular structure supports the core tasks of mission planning and scheduling while accommodating different localization and navigation methods. This adaptability ensures that the system can integrate with COTS platforms, transforming standard electric vehicles into autonomous platforms capable of operating in complex environments.

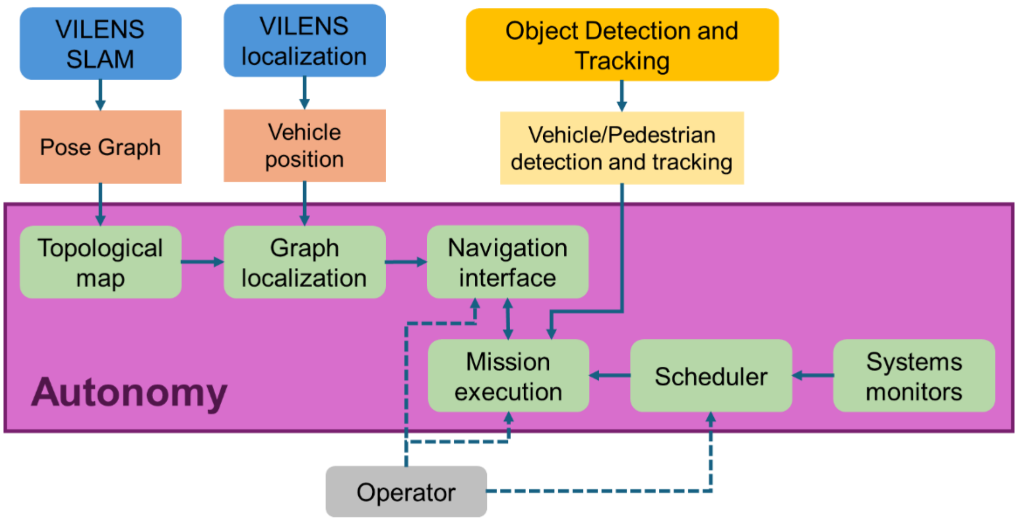

Figure 4. Overview of the autonomy subsystem. The topological map is created using data from the VILENS SLAM pose graph, and the vehicle determines its position on this map using 3D pose data from the localization system. Navigation is handled through the navigation interface. Mission execution performs tasks based on mission specifications. The scheduler organizes and runs missions on a set schedule, allowing interruptions from system monitors. Operators typically interact with the system by scheduling or initiating missions directly.

Topological Map

The topological map forms the core of the autonomy subsystem, serving as a structured representation of the environment. An example of this map is illustrated in Figure 5. The map is represented as a graph consisting of nodes and edges, where an edge connecting two nodes indicates a navigable path between those locations. This representation is both simple and versatile, enabling annotations of nodes and edges with domain-specific information relevant to deployment scenarios. Additionally, it provides a structured abstraction of the environment, facilitating input to advanced planning algorithms.

The topological map primarily captures spatial relationships implicitly through the connections between nodes and edges. For most deployments, nodes are augmented with 3D positions that specify their location within the global reference frame, enhancing spatial precision and compatibility with localization systems. A key feature of the topological map is its adaptability; the user interface allows real-time editing of the map without interrupting the system’s autonomous operation. This flexibility ensures that the autonomy subsystem can accommodate dynamic changes in the environment while maintaining consistent performance.

Figure 5. Top-down view of pointcloud with overlaid topological map.

Topological Navigation

The topological map serves as the foundation for navigation within the autonomy subsystem. To direct the vehicle to a specific node within the graph, the system computes the shortest path by evaluating the costs associated with each edge. The vehicle follows this path by traversing each edge sequentially. If an edge cannot be traversed—for example, due to an obstacle—the system autonomously reroutes by temporarily deactivating the obstructed edge and recalculating the shortest path. This ensures that the vehicle can adapt to dynamic changes in the environment without requiring operator intervention. The navigation subsystem interfaces with the vehicle's navigation system to execute edge traversal. These interfaces translate the planned edge transitions into precise commands for the vehicle, enabling execution of the navigation plan.

Mission Execution and Scheduling

The mission execution subsystem integrates topological navigation with task execution, enabling the autonomous vehicle to complete transportation missions. A task is defined as performing a specific action, such as transporting material from Point A to Point B. A mission consists of one or more tasks, which are typically specified by human operators but can also be generated programmatically. Missions are executed sequentially, with the vehicle navigating to the location of each task, performing the associated action, and proceeding to the next task until all tasks are successfully completed. Minimizing travel distance when manually constructing missions can be challenging for operators. To address this, the system can optimize the order of tasks by solving a standard traveling salesperson problem, reducing overall path costs and enhancing efficiency. Task execution may fail due to an inability to reach the task location or an error during the action itself. When such failures occur, the system can be configured to either continue executing the remaining tasks in the mission or abort the mission entirely, depending on the operational requirements.

The final layer of the autonomy subsystem is the scheduler, which extends mission execution by enabling missions to be repeated on predefined schedules. Missions can be scheduled to run periodically (e.g., hourly, daily at 09:00, or weekly on Tuesdays at 16:00) or as one-time executions at a specific time. The scheduler also incorporates system monitors that can conditionally trigger or prevent mission execution. For example, a battery monitor prevents the initiation of scheduled or user-requested missions when the vehicle’s battery level is low and instead triggers a mission directing the vehicle to the nearest charging station.

RESULTS

Demonstration

The SafeMove system was demonstrated at Oxford's Begbroke Business Park, a controlled environment selected for its suitability to replicate conditions found in nuclear facilities. The demonstration involved equipping an electric vehicle with the SafeMove Payload and testing its core capabilities, including mapping, localization, mission planning, and real-time obstacle detection. The vehicle was manually driven during the entire demonstration to collect data and validate these functionalities.

The demonstration route consisted of an 800-meter loop (approximately 2 kilometers in total), as shown in Figure 6. The area featured a mix of normal roads (red), car parks (orange), and pedestrian paths (green). The route included waypoints labeled A through L to simulate a series of point-to-point transportation missions. These waypoints represented key locations connected by the system's graph-based map, enabling the planning and demonstration of point-to-point missions, mimicking material transport in nuclear facilities. The route was specifically chosen to include varied environments accessible to small vehicles while adhering to a speed limit of 5–10 mph. These conditions closely resembled operational constraints in nuclear facilities, such as limited space, diverse terrain, and strict safety requirements.

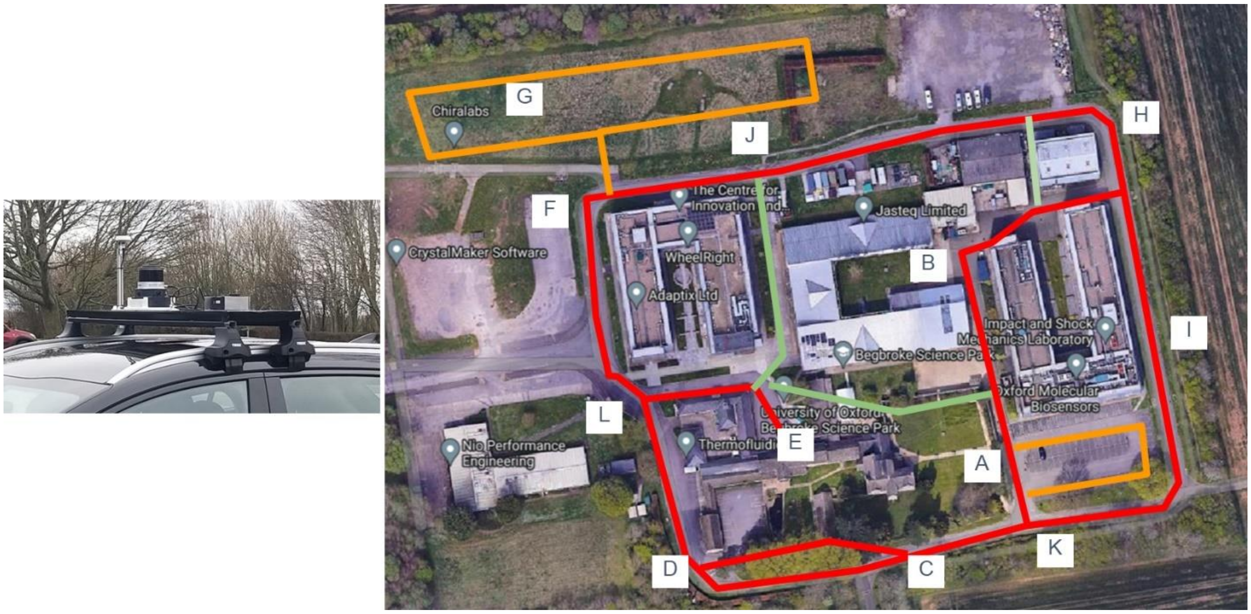

Figure 6. (Left) SafeMove Payload mounted on an electric vehicle using a standard roof rack. (Right) Overview of Oxford's Begbroke Business Park. The red lines represent normal roads, the orange sections indicate car parks, and the green paths correspond to pedestrian paths. Waypoints labeled A through L represent starting and destination points for point-to-point missions. The route spans 800 meters per loop, totaling approximately 2 kilometers.

Mapping and Localization

The first stage of the demonstration involved using the SafeMove Payload to generate a detailed 3D map of the campus while the vehicle was manually driven along the designated route. This 3D map formed the basis for the system's topological navigation and mission planning by creating a graph network that connected key physical locations. The mapping process was entirely independent of GPS, leveraging LiDAR-based place recognition to localize the vehicle within the generated map. This capability demonstrated SafeMove's suitability for GPS-denied environments, a critical requirement for nuclear facilities.

The mapping process supports progressive updates to extend operational areas by merging data from multiple mapping sessions. As shown in Figure 7, the mapping system progressively integrates new sessions, with the single-session map expanding to include data from three and six sessions, resulting in a comprehensive and detailed representation of the environment."

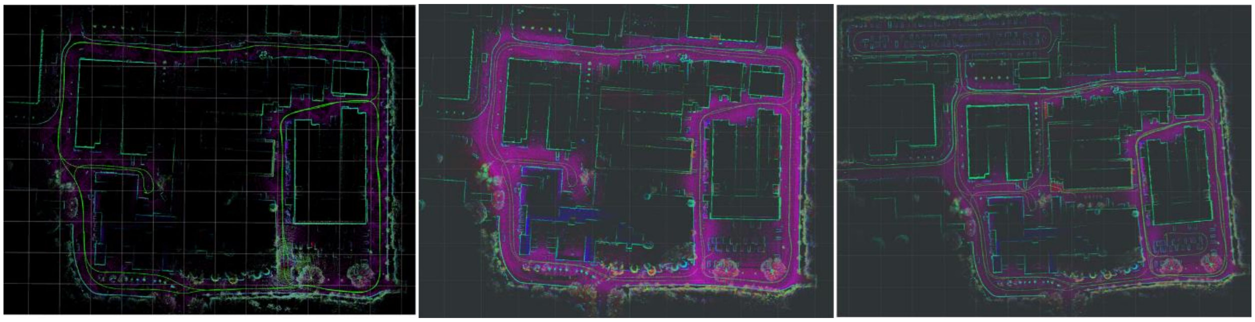

Figure 7. Results of Multi-Sensor Mapping. Left image shows a single-session map, center image demonstrates the combined map after three sessions, and right image presents the extended map after six sessions. This sequence demonstrates the system's ability to integrate data from multiple mapping sessions, progressively extending the operational area while ensuring detailed and accurate representations.

Figure 8 illustrates the LiDAR place recognition system, which operates in two distinct phases: map building and online operation. During the map-building phase, a 3D map of the test site is constructed, consisting of nodes spaced 1–2 meters apart. The connectivity between these nodes forms the autonomy "highway," where traveling along the nodes is assumed to be safe unless obstructed. Additionally, a database of descriptors is created for each node to support future localization. During the online operation phase, the system identifies matching nodes from the map to estimate the vehicle's position within the local point cloud. Once multiple matches are accumulated, the vehicle is declared "locked into the map," ensuring robust and precise localization.

Figure 8. Results of LiDAR Place Recognition. (Left) A section of the 3D map showing node connectivity during the map-building phase. (Right) Online operation illustrating localization within the map. The system identifies matching nodes in real-time to estimate the vehicle's position within the environment, enabling precise and reliable localization. A video demonstrating these results is available at https://youtu.be/zSbx-vCsm2Q

Point-to-Point Mission Planning

The demonstration showcased point-to-point mission planning using the constructed topological map. With the map as a foundation, the system was able to plan efficient routes between waypoints, simulating material transport missions such as moving from a home base at Point A to destinations including Points B, C, D, E, and others. Although navigation was not autonomous during this demonstration, the system successfully computed routes, optimized path planning, and demonstrated the process of mission creation and execution within the mapped environment.

As illustrated in Figure 9, the topological map serves as a navigation graph, representing a structured "highway" for the vehicle to follow. Nodes signify key waypoints, and edges represent traversable paths connecting these waypoints. The left image shows the navigation graph overlaid on the 3D map, with the green dot indicating the starting point of the mission and the yellow dot representing the endpoint. The right image highlights the system’s local path planning, where the yellow line depicts the specific edges to be traversed based on the mission planning policy. The point-to-point mission planning workflow transitions from the 3D map to the autonomy graph. The system dynamically generates a mission planning policy to ensure safe and efficient navigation, specifying which edges should be traversed when the vehicle reaches a node. This capability validates the robustness of the SafeMove system in mission creation and execution.

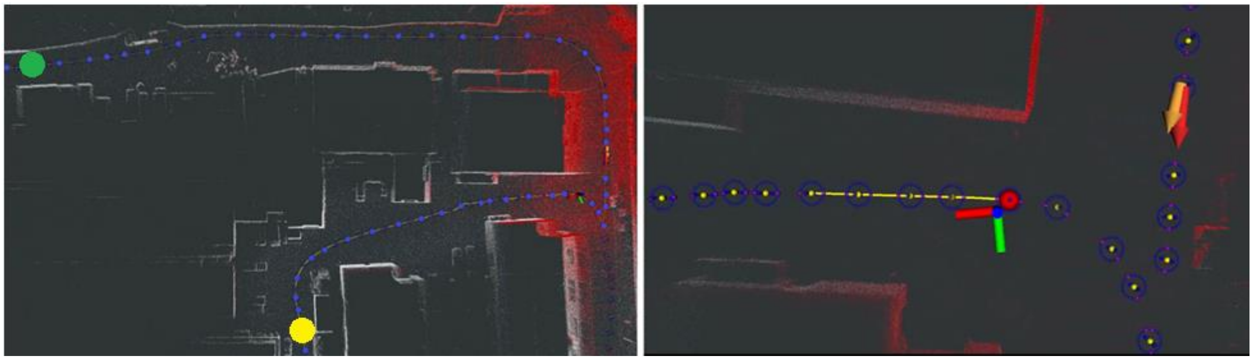

Figure 9. Results of Point-to-Point Mission Planning. (Left) Navigation graph overlaid on the 3D map, with the green dot representing the starting point and the yellow dot representing the endpoint. (Right) Local path planning is illustrated by the yellow line, indicating the edges to be traversed within the navigation graph. The system transitions from a 3D map to an autonomy graph, enabling precise route planning for autonomous missions. A video demonstrating these results is available at https://youtu.be/iXn4U9DBWtQ

Object Detection and Tracking

The final stage of the demonstration focused on validating the real-time object detection and tracking capabilities of the SafeMove system. The system effectively identified and tracked dynamic objects, including pedestrians and vehicles, providing the necessary situational awareness for safe navigation in complex environments. These capabilities are critical for ensuring safety in operational deployments at nuclear facilities.

Object Detection

As illustrated in Figure 10, the system demonstrates its object detection capabilities. The left image shows vehicles being detected, while the right image highlights the detection of pedestrians. Each detected object is assigned a classification label and confidence score, showcasing the system's ability to accurately identify different object types in real time. This detection functionality serves as the foundation for subsequent tracking processes, enabling precise monitoring of object trajectories during navigation.

Figure 10. Results of Object Detection. (Left) Detection of vehicles, including classification labels and confidence scores. (Right) Detection of pedestrians.

The object detection process is powered by the D-FINE network, which currently holds the highest mAP (mean Average Precision) for the COCO dataset. The network delivers high detection performance with a mean latency of less than 30 milliseconds on a standard laptop PC, enabling real-time operation without compromising classification accuracy. Comparisons with other detection models, such as SSDLite + MobileNet and FasterRCNN, confirmed D-FINE's superior balance of speed and reliability. The performance metrics for these models are summarized in Table 1.

Table 1: Performance Comparison of Object Detection Models in Terms of Accuracy and Latency

Model	Mean Accuracy (COCO)	Latency (laptop GPU)
SSDLite + MobileNet V3	21.3%	14ms
FasterRCNN + MobileNet v3	22.8% / 32.8%	25ms / 31ms
D-FINE Nano	42.8%	27ms

Object Tracking

The ByteTracker algorithm consistently tracked objects across frames, assigning unique IDs and generating reliable object trajectories. Testing in environments representative of nuclear facilities confirmed its effectiveness in detecting and tracking dynamic obstacles, such as vehicles and pedestrians, with high accuracy and efficiency. As illustrated in Figure 10, ID 73 is assigned to a person and consistently maintained throughout the sequence. Additionally, another person is identified as ID 75 for part of the sequence, demonstrating the algorithm's capability to manage multiple objects simultaneously. These results highlight the robustness of the tracking system, which is crucial for ensuring safe navigation in dynamic and complex environments.

Figure 11. Results of Object Tracking. The sequence demonstrates robust tracking performance, with ID 73 consistently assigned to a person throughout the scenario and ID 75 assigned to another individual for part of the sequence. A video demonstrating these results is available at https://youtu.be/r9t8vxUKLP0

Risk Assessment for Deploying Autonomous Vehicles in Nuclear Facilities

The deployment of autonomous vehicles in nuclear facilities requires a detailed hazard analysis to ensure safety, compliance, and reliability. Guided by ISO 26262 standards, a Hazard Analysis and Risk Assessment (HARA) was conducted to identify and mitigate potential risks. This process evaluated hazards based on severity, exposure, and controllability, producing actionable mitigation strategies summarized in Table 2. This structured risk assessment serves as a framework for deploying autonomous vehicles (including systems like SafeMove) in nuclear facilities, addressing complex operational challenges and ensuring safe, reliable integration.

Table 2: Hazards, and Mitigation Strategies for Deploying Autonomous Vehicles in Nuclear Facilities

ID	Hazard Description	Hazardous Event	Severity	Exposure	Controllability	Mitigation Strategy
H001	Software glitch/bug in autonomous system	Autonomous vehicle fails to respond correctly	High	High	Low	• Rigorous testing and validation of autonomous systems. • Implement fail-safe mechanisms (e.g., emergency braking, manual override). • Regular maintenance and updates.
H002	Malfunction and technical/hardware issues	Autonomous vehicle malfunctions	High	Medium	Medium	• Real-time monitoring and adaptive responses to unexpected scenarios. • Backup transportation methods (e.g., manual drivers). • Plan for extreme weather conditions and road closures.
H003	Cybersecurity vulnerabilities	Unauthorized access to vehicle systems	High	High	Low	• Robust cybersecurity protocols (encryption, intrusion detection). • Isolation of critical vehicle systems. • Regular security audits. • Secure communication channels
H004	Non-compliance with safety regulations	Violation of safety regulations	High	Medium	High	• Collaborate with regulatory bodies during the system design. • Obtain necessary permits and approvals. • Maintain accurate records
H005	Accidental spills or leaks during transportation	Environmental contamination	Critical	Medium	Low	• Secure packaging and containment of waste packages.

						• Emergency response protocols for spills or leaks.
H006	Data privacy concerns	Unauthorized access to sensitive data	High	High	Medium	• Implement strict data access controls. • Anonymize or pseudonymize sensitive data. • Comply with privacy regulations.
H007	Cost overruns during development and deployment	Project exceeds budget	Low	Medium	High	• Detailed cost analysis during development and deployment.
H008	Ethical dilemmas	Ethical guidelines are not followed	Medium	Low	High	• Establish clear ethical guidelines for AV behaviour. • Consider societal impact in decision-making.
H009	Interoperability and standards	Non-compliance with industry standards	Medium	Medium	High	• Align with industry standards (e.g., communication protocols, safety requirements). • Collaborate with other techno providers and AV manufacturers.
H010	Infrastructure readiness	Infrastructure is not compatible with AVs	High	Medium	Medium	• Assess road infrastructure (e.g., signage, lane markings) for AV compatibility. • Invest in necessary infrastructure upgrades.

CONCLUSIONS

SafeMove addresses critical challenges in nuclear waste transportation by demonstrating capabilities in mapping and localization, point-to-point mission planning, and object detection and tracking. Its modular architecture integrates GPS-independent navigation, advanced planning, and real-time situational awareness, ensuring reliable operation in complex and hazardous environments. Field demonstrations validated the system's ability to construct detailed 3D maps, plan and simulate efficient point-to-point missions, and detect and track dynamic objects, establishing its readiness for deployment in nuclear facilities. The system’s modular design facilitates integration with a wide range of vehicle platforms, while its advanced mapping and localization capabilities ensure precision without requiring extensive infrastructure modifications. Object detection and tracking features enhance safety and compliance with nuclear industry standards, supporting robust operations in challenging environments.

Future work will focus on transitioning SafeMove from manual demonstrations to fully autonomous deployment, including the integration of control systems and vehicle interfaces. Additionally, its modular architecture supports further applications, such as autonomous inspection and monitoring in other large-scale industrial settings, ensuring its continued relevance in advancing operational safety and efficiency.

REFERENCES

1. AtkinsRéalis, "Sellafield robots leading an evolution in nuclear decommissioning," [Online]. Available: https://www.gov.uk/government/news/sellafield-robots-leading-an-evolution-in-nuclear-decommissioning [Accessed: Nov. 2024].
2. Kromek, "UAV radiation mapping drone - Kromek," [Online]. Available: https://www.kromek.com/product/autonomous-airborne-radiation-monitoring-system/ [Accessed: Nov. 2024].
3. DHL, "DHL Supply Chain Deploys Its First Autonomous Forklifts in North America," [Online]. Available: https://www.dhl.com/us-en/home/press/press-archive/2021/dhl-supply-chain-deploys-its-first-autonomous-forklifts-in-north-america.html [Accessed: Nov. 2024].
4. M. Staniaszek, T. Flatscher, J. Rowell, H. Niu, W. Liu, Y. You, R. Skilton, M. Fallon, N. Hawes, “AutoInspect: Towards Long-Term Autonomous Industrial Inspection”, IEEE Int. Conf. On Robotics and Automation (ICRA): Workshop on Field Robotics, May 2024.
5. J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, “Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes,” in IEEE Intl. Conf. on Robotics and Automation (ICRA), 2016, pp. 4304–4311.
6. L. F. T. Fu, N. Chebrolu, and M. Fallon, “Extrinsic calibration of camera to lidar using a differentiable checkerboard model,” in IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), 2023, pp. 1825–1831
7. D. Wisth, M. Camurri, and M. Fallon, “VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots,” IEEE Trans. Robotics, vol. 39, no. 1, pp. 309–326, 2023.
8. M. Ramezani, G. Tinchev, E. Iuganov, and M. Fallon, “Online LiDARSLAM for legged robots with robust registration and deep-learned loop closure,” in IEEE Intl. Conf. on Robotics and Automation (ICRA), 2020, pp. 4158–4164.
9. L. Zhang, M. Helmberger, L. Fu, D. Wisth, M. Camurri, D. Scaramuzza, and M. Fallon, “Hilti-Oxford Dataset: A millimeter-accurate benchmark for simultaneous localization and mapping,” IEEE Robot. Autom. Lett. (RA-L), vol. 8, pp. 408–415, 01 2023.
10. A. Proudman, M. Ramezani, S. T. Digumarti, N. Chebrolu, and M. Fallon, “Towards real-time forest inventory using handheld LiDAR,” IEEE Robot. Auton. Syst., vol. 157, p. 104240, 2022.
11. R. Border, N. Chebrolu, Y. Tao, J. D. Gammell, and M. F. Fallon, “Osprey: Multi-session autonomous aerial mapping with lidar-based SLAM and next best view planning,” CoRR, vol. abs/2311.03484, 2023
12. F. Dellaert and M. Kaess, “Factor Graphs for Robot Perception,”Foundations and Trends in Robotics, vol. 6, pp. 1–139, 2017..
13. G. Kim and A. Kim, “Scan Context: Egocentric spatial descriptor for place recognition within 3D point cloud map,” in IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), 2018, pp. 4802–4809.
14. B. Kuipers, “Modeling spatial knowledge,” Cognitive Science, vol. 2, no. 2, pp. 129–153, Apr. 1978.
15. R. Brooks, “Visual map making for a mobile robot,” in IEEE Int. Conf. on Robotics and Automation (ICRA), vol. 2, 1985, pp. 824–829.
16. N. Hawes et al., “The STRANDS project: Long-term autonomy in everyday environments,” IEEE Robot. Autom. Mag., vol. 24, no. 3, pp. 146–156, Sep. 2017.
17. G. Das et al., “A unified topological representation for robotic fleets in agricultural applications,” Sep. 2023.
18. B. Lacerda, F. Faruq, D. Parker, and N. Hawes, “Probabilistic planning with formal performance guarantees for mobile service robots,” Int. J. Robot. Res., vol. 38, no. 9, pp. 1098–1123, Aug. 2019.
19. Z. Li, J. D. Wegner, and A. Lucchi, “Topological map extraction from overhead images,” in Int. Conf. on Computer Vision (ICCV), IEEE, Oct. 2019, pp. 1715–1724.
20. F. Blochliger et al., “Topomap: Topological mapping and navigation based on visual SLAM maps,” in IEEE Int. Conf. on Robotics and Automation (ICRA), May 2018, pp. 3818–3825.