AGenSE Blog: How MOSAICO and ASTIR are shaping the future of Agentic Software Engineering

Insights from the AGenSE Alliance on orchestrating AI Agents for reliable and safe software lifecycles

Authors: Anh Tuan LE[1], Jean-Marie MOTTU[2], Sebastian Scholze[1]
[1] ASTIR | [2] MOSAICO

Bridging safety-critical robotics and multi-agent orchestration

As part of the AGenSE Alliance, alongside sister projects AI4SWEng, HIVEMIND, and FAITH, members of the Horizon Europe projects MOSAICO and ASTIR recently gathered in Munich during the AgenticDev 2026 workshop they organised as part of ASE 2026. Their joint goal is to explore how agentic software engineering can elevate quality, reliability, and safety across complex development lifecycles.

While ASTIR focuses on AI-augmented engineering workbenches tailored to intelligent robotic systems, spanning requirements, architecture, code, testing, security, and runtime monitoring, MOSAICO delivers a management framework for coordinating and supervising communities of AI agents across the software engineering lifecycle. Together, their synergies highlight key pillars for next-generation software development.

Key synergies and technical integration

1. Graph-based agent collaboration across lifecycles

ASTIR is developing an agent-based engineering workbench that integrates specialised AI capabilities across the robotics software development lifecycle, including requirements engineering, system design, implementation, testing, security analysis, and runtime

monitoring. Its architecture already supports agent orchestration and communication between agents and external tools (*Figure 1*).

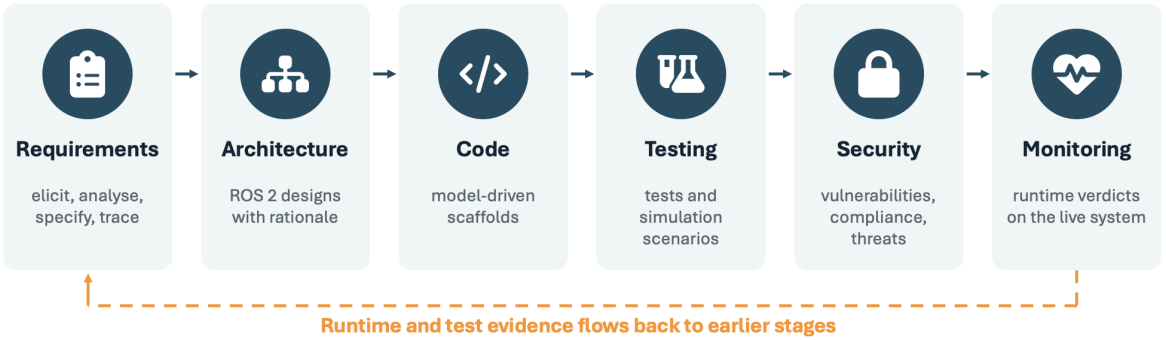

Figure 1: Overview of the ASTIR lifecycle phases and change propagation feedback loop from runtime monitoring to requirements.

MOSAICO explores graph-based collaboration mechanisms in which coordination agents distribute tasks among specialised solution agents, while consensus and supervision mechanisms assess the consistency of their outputs (*Figure 2*). These approaches could complement ASTIR's existing agent orchestration architecture by supporting more coordinated execution and verification across software engineering activities. A promising research direction is to investigate how MOSAICO's coordination mechanisms could interact with ASTIR's requirements engineering, architecture design, code generation, testing, security analysis, and runtime monitoring capabilities.

2. Trust checks for agent outputs

ASTIR investigates knowledge-graph-based grounding and systematic testing techniques to improve the trustworthiness of generative AI. These approaches could complement MOSAICO's supervision and consensus mechanisms by providing additional evidence for assessing generated outputs, identifying inconsistencies, and evaluating robustness. A worthwhile joint research direction is to investigate how such verification evidence can be incorporated directly into multi-agent decision-making processes.

3. Hybrid LLM deployment strategy

Robotics software engineering often involves proprietary development artefacts, sensitive operational data, and strict security requirements, constraints that can make local LLM deployment preferable or necessary in specific applications.
Both ASTIR and MOSAICO consider flexible deployment of AI capabilities. Their complementary approaches offer opportunities to investigate coordination between locally and remotely deployed agents while maintaining appropriate data access controls, traceability, and execution policies (*Figure 2*).

To address this, ASTIR provides AI-augmented engineering workbenches tailored specifically to intelligent robotic systems, covering the full software engineering lifecycle across requirements specification, architecture design, code synthesis, testing, security verification, and runtime monitoring. Within its framework, MOSAICO supports flexible hybrid

setups, letting teams run local LLMs whenever proprietary rules require on-premise data control or leverage remote LLM agents when appropriate.

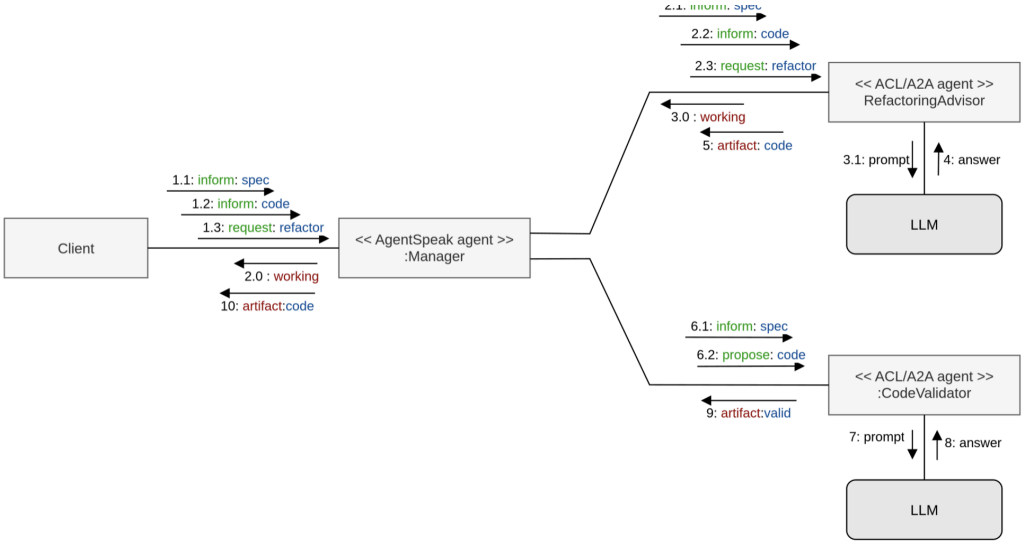

Figure 2: Hybrid LLM architecture enabling coordination between manager and specialized agents across local and remote environments.

4.Model-driven artefacts:

ASTIR investigates model-based representations, including UML/SysML models and Behaviour Trees, to support selected robotics software engineering activities. Such artefacts could provide structured interfaces between specialised agents in MOSAICO, complementing natural-language communication. A promising research direction is to examine how machine-readable models, explicit metamodels, and validation constraints can support more consistent information exchange and traceability across multi-agent workflows.

5. Redundant verification for strict safety requirements

Safety-critical robotics applications require systematic verification throughout the software engineering lifecycle. MOSAICO's multi-agent architecture offers opportunities to coordinate complementary verification activities, including independent analyses of generated artefacts, where discrepancies between agent outputs are identified through dedicated supervision mechanisms and subjected to further assessment. Combining these approaches with ASTIR's testing, security analysis, and dependability assurance technologies could support more robust verification workflows, though the effectiveness of such approaches, including their ability to detect correlated errors, will need to be evaluated.

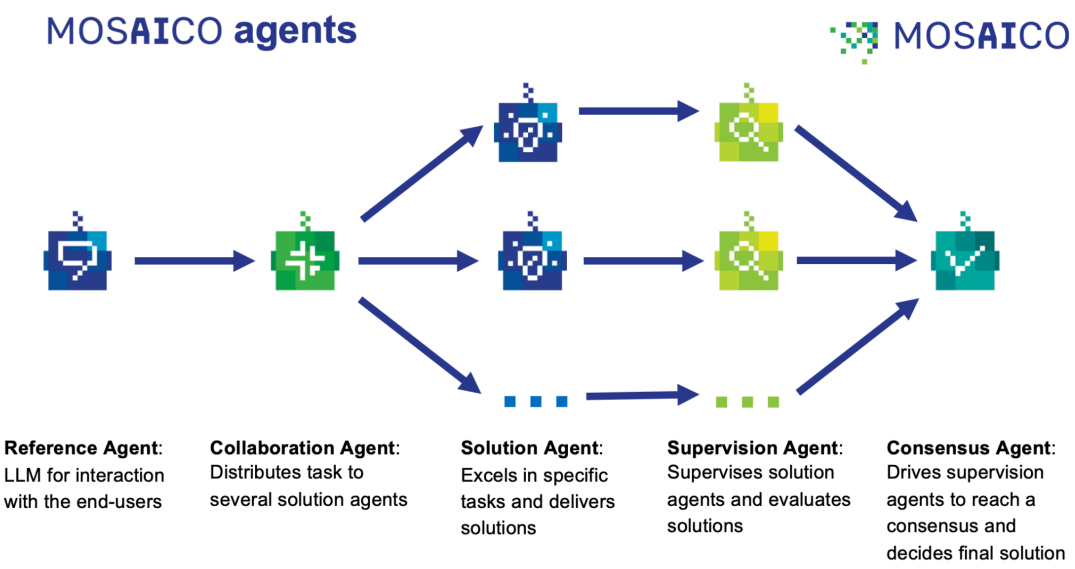

Figure 3: MOSAICO multi-agent collaboration model showing task distribution, solution execution, supervision, and consensus decision-making.

6. Runtime monitor generation for continuous supervision

ASTIR investigates AI-assisted generation of formal specifications and runtime monitors to support the detection of deviations from expected system behaviour. These capabilities could complement MOSAICO's supervision mechanisms by providing live information about system behaviour during operation. A joint collaboration would investigate how monitoring observations can be fed back into agent-supported engineering workflows, enabling developers to identify affected requirements, designs, or implementation artefacts and initiate controlled revisions.

Looking ahead: Next steps for AGenSE

The discussions at AgenticDev 2026 identified several promising directions for further collaboration between ASTIR and MOSAICO, including investigating interoperability between agent orchestration approaches, exploring the integration of verification evidence into multi-agent supervision, and examining opportunities for shared evaluation scenarios. ASTIR's robotics software engineering use cases and benchmark framework could provide a useful basis for assessing selected multi-agent engineering workflows. Moving forward, the projects will explore which of these opportunities can be developed into concrete joint research activities within the AGenSE Alliance.