Towards Verifiable Safety for Autonomous Robots

Safe-ROS: An Architecture for Autonomous Robots in Safety-Critical Domains

Diana C. Benjumea[1], Louise A. Dennis[1], Marie Farrell[1]

The use of autonomous robots in safety-critical domains can improve human safety, task efficiency and cost. However, without formal evidence that these systems are free from unexpected and hazardous behaviour, deployment in such domains is still restricted in practice.

1 Our contribution

The **Safe-ROS** architecture for developing *reliable* and *verifiable* autonomous robots.

2 Case Study

Platform: AgileX Scout Mini.
Scenario: Nuclear Inspection.

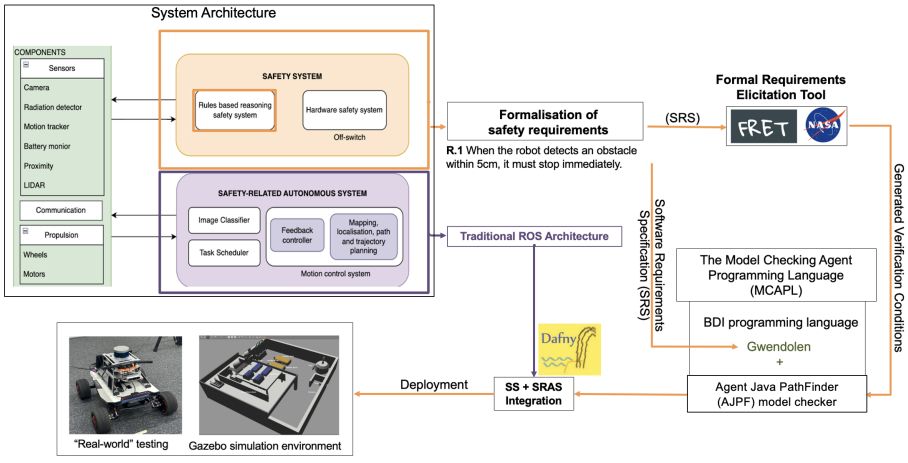

2.1 From Safety Requirement to Verifiable Property:

R1: *When the robot detects that an obstacle is within 5cm of it, then it must stop immediately.*

whenever too_close agilex_agent shall satisfy stopped

LTL expression: G (too_close → F stopped)

[] (B(agilex_agent, too_close) →<> B(agilex_agent, stopped))

Validation: Testing

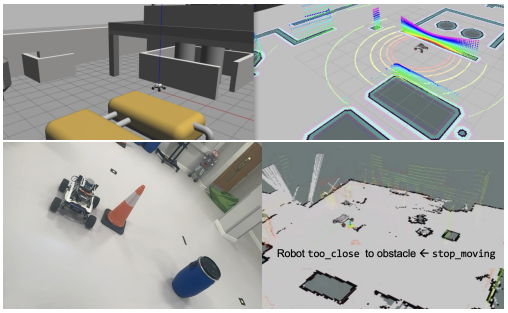

Future work

Extending Safe-ROS to incorporate richer safety properties (e.g., returning to a door)

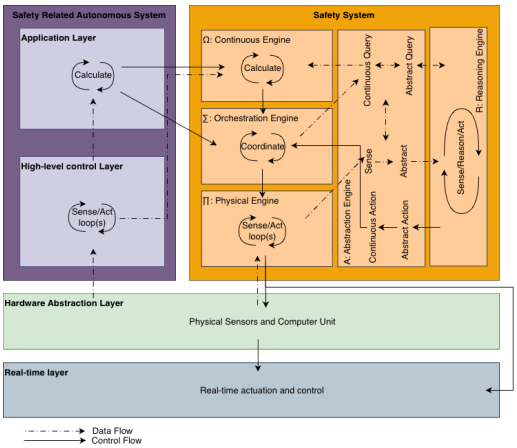

This work was funded in part by The University of Manchester, the EPSRC-funded CRADLE project (EPSRC grant EP/X02489X/1), and the Royal Academy of Engineering, and benefited from a Fellowship at RAICo (The Robotics and AI Collaboration).

1. University of Manchester

diana.benjumeahernandez@manchester.ac.uk

MANCHESTER 1824
The University of Manchester

Download the paper →