Program: Resilient Risk-Aware Autonomy for the Exploration of Uncertain and Extreme Environments

Campus PI: Richard M. Murray
JPL PI: Michel D. Ingham

Co-Investigators: Richard Camilli, WHOI
Tara Estlin, JPL
Masahiro Ono, JPL
Brian Williams, MIT

Brief summary of progress to date

The long term goal of this project is to develop a resilient, risk-aware software architecture for onboard, real-time autonomous operations that is intended to robustly handle uncertainty in spacecraft behavior within hazardous and unconstrained environments, without unnecessarily increasing complexity. This architecture, the *Resilient Spacecraft Executive* (RSE), serves three main functions: (1) adapting to component failures to allow graceful degradation, (2) accommodating environments, science observations, and spacecraft capabilities that are not fully known in advance, and (3) making risk-aware decisions without waiting for slow ground-based reactions. This architecture is made up of four main parts: deliberative, habitual, and reflexive control modules, as well as a state estimator that interfaces with all three control modules (Figure 1). We use a risk-aware goal-directed executive within the deliberative module to perform risk-informed planning, to satisfy the mission goals (specified by mission control) within the specified priorities and constraints. Other state-of-the-art algorithms to be integrated into the RSE include correct-by-construction control synthesis and model-based estimation and diagnosis.

Our two-year effort has taken critical initial steps in the development of highly-resilient spacecraft. This effort has developed an innovative software architecture, the Resilient Spacecraft Executive (RSE), that will endow spacecraft with unprecedented levels of resilience, leveraging planning, execution and control technologies being developed at MIT and Caltech. In 2016, our objectives were to (1) develop additional mission scenarios to exercise RSE capabilities and demonstrate applicability to problems of interest at JPL; (2) further develop the core algorithms that provide the capabilities of the RSE; (3) develop a modeling framework and tools to specify the formal behavior models of the system, which are needed by these algorithms for online reasoning and used offline for formal verification and validation (V&V) of the RSE capabilities; (4) integrate the algorithms and modeling framework into the RSE architecture; and (5) deploy and demonstrate RSE on high-fidelity simulation and hardware test platforms for surface rover and autonomous underwater vehicle (AUV) scenarios.

As shown in Figure 1, we define decision/control modules analogous to “reflexive” behavior hardwired in the nervous system; “habitual” behaviors that are performed by rote once learned through repetition and muscle-memory training; and finally, “deliberative” reasoning behavior that is used to make decisions, handle novel situations, learn from mistakes, etc. By comparison, the control paradigm for current spacecraft is biased heavily towards hard-coded reflexive behavior, with some pre-validated higher-level behaviors akin to habitual behaviors in humans, but little to no deliberative reasoning aside from that performed by ground operators.

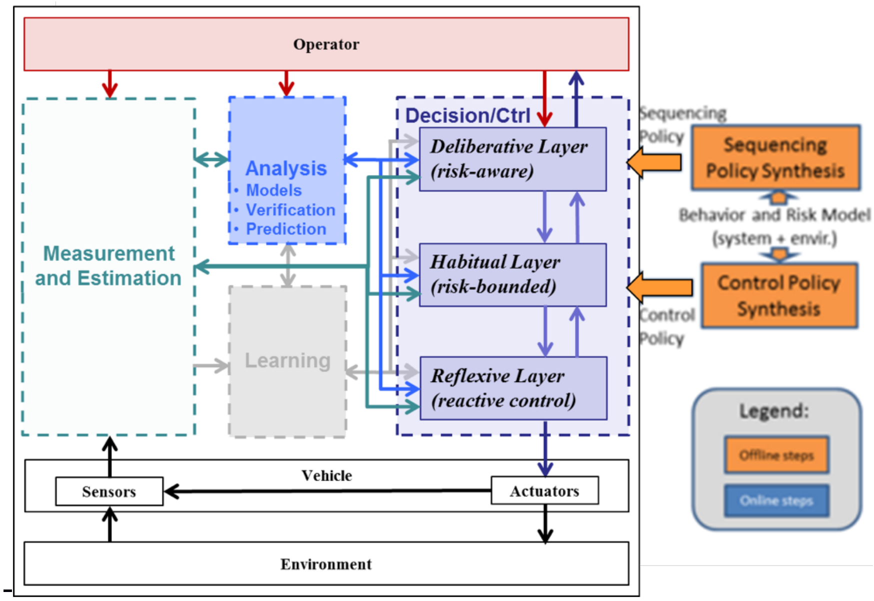

Figure 1. Conceptual view of the RSE architecture.

JPL, Caltech and MIT have collaborated on the integration of the individual algorithmic components into the RSE, including specification of the various interfaces, both internal and external to the architecture, and assembly of the overall Executive framework. The open-source Robot Operating System (ROS) middleware software provides a flexible software substrate for the architecture.

In particular, the team has been developing two main risk-aware planning components of the risk-aware goal-directed executive: (1) an activity planning component and (2) a path planning component. The *risk-aware activity planning component* accounts for temporal disturbances while generating plans that satisfy the given mission goals and constraints on operational risk (called chance constraints). The current planner incorporates probabilistic models of activity delay, which it employs within a risk-aware temporal consistency checker to allow the planner to explicitly evaluate the uncertainty in action duration. The *risk-aware path planning component* consists of a global path planner and a kino-dynamic path planner that together allow a vehicle to traverse an environment with bounded risk on obstacle collision. The activity and path planning components have been integrated to allow the activity planner to invoke the path planner to check and compute safe paths for a vehicle traversal actions.

The team has also focused on detecting and handling disruptions and off-nominal cases during plan execution (e.g., camera failure; drilling failure, the inability to detect a scientifically-interesting rock to sample in a target location; and the inability to complete a traverse activity within the specified deadline). A hybrid state estimation capability has been implemented to allow diagnosis of subtle degraded and failure modes of behavior and components (e.g., health of the onboard camera dependent on its current temperature). The team has also extended the capability of the aforementioned activity planning component to allow the generation of pre-computed risk-bounded control policies for managing instrument and component reconfiguration (e.g., computing conditional plans for drilling process that bound the risk of breaking drill bits and using spare ones). An implementation of the risk-aware goal-directed executive, consisting of the two integrated planning components, the hybrid state estimator and an activity dispatcher have been incorporated into the deliberative layer and integrated with the habitual and reflexive layers.

The integrated system has been demonstrated in a set of simulated Mars rover scenarios and in three Autonomous Underwater Vehicle (AUV) hardware deployments: 1) at Scott Reef (Timor Sea), from March 24 to April 6, 2015; 2) the Santa Barbara Basin (Eastern Pacific) from September 5 to 16, 2016; and 3) a 1-day, multiple glider deployment in Cape Cod Bay (North Atlantic) on November 17, 2016, in which the mission included three gliders.

In the simulated rover deployment, we demonstrated resilience by injecting failures/degradations at various stages and showing how the integrated RSE enabled recovery and graceful degradation capabilities. In particular, we considered (1) Camera failure; (2) drilling failure; (3) the inability to detect a scientifically-interesting rock to sample in a target location; and (4) the inability to complete a traversal within a specified deadline. In these cases, the anomalous events are detected and managed by the activity planner, which replans the mission in light of these unexpected events, taking into account the modeled risks and the mission risk posture specified by mission operators. For example, it handled the failure of the Mastcam by replanning all remaining imaging activities to use the Hazcam (i.e., leveraging functional redundancy), and it handled the absence of a scientifically-interesting rock at one site by replanning the activities at a later site to include the sampling activity that was missed. We performed successful demonstrations of resilient risk-aware autonomy capabilities in the Gazebo simulation environment at the mid-year review in April 2016, and in October 2016.

In the AUV hardware deployment, MIT and WHOI team members deployed a risk-aware goal-directed planning capability as a decision support system for AUV gliders. The operators used the executive to plan a series of observations of target regions in between surfacing for data communication, and to plan underwater paths for the observations. The executive then generated mission scripts that were directly executable by the gliders. The demonstrations included replanning of trajectories to mitigate for spatial and temporal changes in ocean currents. These deployments showed promising impact of the technology towards endowing vehicles with higher levels of resilience and autonomy. In the first deployment, the system allowed the glider to navigate safely through a reef. To the best of our knowledge, an AUV glider has never been used inside a reef before, due to the challenges present in that environment. The Santa Barbara Basin deployment (Fig. 2) successfully demonstrated autonomous survey of seafloor hydrocarbon seeps using an experimental payload mass spectrometer operating on the AUV glider. This field site served as an analog for other planetary bodies such as Titan, Enceladus, and Europa, which are thought to contain hydrothermal vents and/or hydrocarbon seeps. For this demonstration, the mission planner optimized science goal activities and vehicle trajectory, accounting for navigation uncertainty, available time windows, vehicle drift (due to water currents). During these demonstrations the AUV operated to a maximum depth of 1,880 ft (570m), visiting 72% of its science goal locations (13 of 18 locations) within the specified temporal window, and arriving within 15 meters of the 3D goal point, despite operating in water column currents that were up to 80% of the vehicle speed (Fig. 3 and Fig. 4). In the third demonstration deployment, the planner successfully demonstrated simultaneous allocation of multiple vehicles to multiple target regions within time windows and deadlines. This demonstration is a stepping stone to the recently awarded project through the PSTAR program (Grant #NNX16AL08G) entitled “Cooperative Exploration With Under-actuated Autonomous Vehicles in Hazardous Environments.”

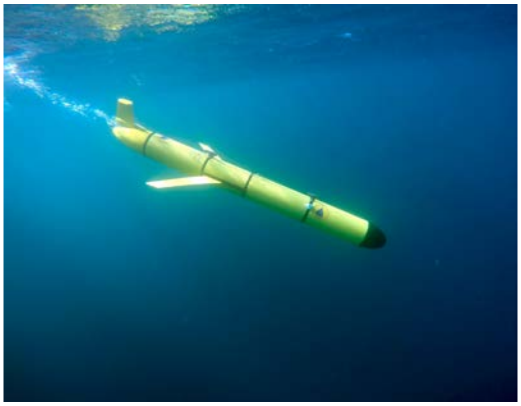

Figure 2: WHOI glider deployment (during descent).

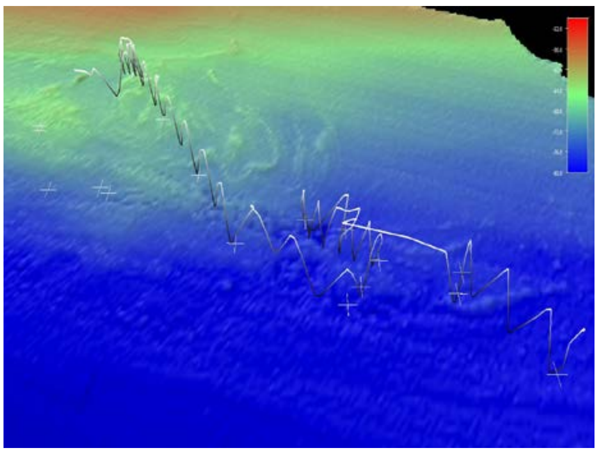

Figure 3: Perspective view of reconstructed autonomous survey results shows science goal locations (white crosshairs), and vehicle track log (gray scale line, where white indicates surface and black indicates >70 m depth). Underlay shows bathymetry with color bar at upper right indicating seafloor depth in meters.

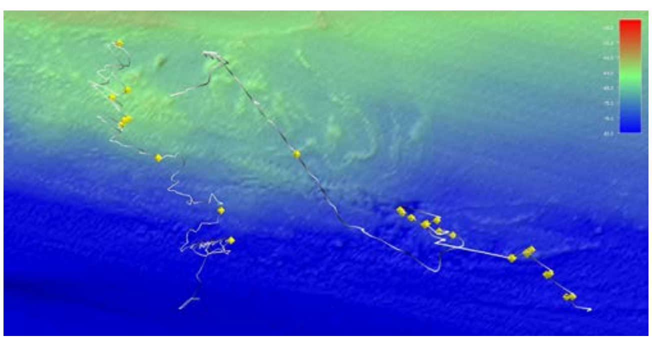

Figure 4: Plan view of reconstructed autonomous survey results showing anomalous hydrocarbon locations identified over a seafloor tectonic fault. Anomaly locations are indicated as yellow diamond icons. The gray-scale line shows vehicle track (where white indicates surface and black indicates >70 m depth). Underlay shows bathymetry with color bar at upper right indicating seafloor depth in meters.

Status of Collaborations (Campus/JPL/External)

The team is distributed across four institutions: Caltech, JPL, MIT and WHOI. Project meetings were held every other week through September 2016, alternating locations between campus and lab (MIT and WHOI researchers participate via WebEx). In addition to the bi-weekly project meetings, there are frequent meetings between Caltech and JPL researchers, as well as between MIT and WHOI researchers. The team also held two technical exchange meetings, in August 2015 and August 2016, at the Woods Hole Oceanographic Institution (WHOI) in Massachusetts.

Caltech postdoc Catharine McGhan had an office on campus and on lab, allowing her to spend her time in either location. Caltech postdoc Tiago Vaquero has an office at MIT, allowing him to collaborate closely with the MIT researchers.

Papers / Technical Reports to date

(bullet list of ALL published, accepted, in preparation, from throughout the project etc.)

1. C. L. R. McGhan, Richard M. Murray, R. Serra, M. D. Ingham, M. Ono, T. Estlin, B. C. Williams, “A Risk-Aware Architecture for Resilient Spacecraft Operations”. Proceedings of the IEEE Aerospace Conference, 2015.
2. C. L. R. McGhan and Richard M. Murray, “Application of Correct-by-Construction Principles for a Demonstratively Resilient Risk-Aware Architecture”. IEEE Space Forum 2015.
3. C. L. R. McGhan, R. M. Murray, T. Vaquero, B. C. Williams, M. D. Ingham, M. Ono, T. Estlin, A. R. K. Lanka Subrahmanya, and M. E. Elaasar, “The Resilient Spacecraft Executive: An Architecture for Risk-Aware Operations in Uncertain Environments,” AIAA Space and Astronautics Forum and Exposition (AIAA SPACE) 2016.
4. C. L. R. McGhan, Y. Wang, R. M. Murray, T. Vaquero, B. C. Williams, and M. Colledanchise, “Towards Architecture-wide Analysis, Verification, and Validation for Total System Stability During Goal-Seeking Space Robotics Operations,” AIAA Space and Astronautics Forum and Exposition (AIAA Space) 2016.
5. P. Santana, T. Vaquero, E. Timmons, B. C. Williams, C. L. R. McGhan, R. M. Murray, C. Toledo, "Risk-aware Planning in Hybrid Domains: An Application to Autonomous Planetary Rovers," AIAA Space and Astronautics Forum and Exposition (AIAA SPACE) 2016.
6. P. Santana, T. Vaquero, C. Toledo, A. Wang, C. Fang, and B. Williams, "PARIS: a Polynomial-time, Risk-Sensitive Scheduling Algorithm for Probabilistic Simple Temporal Networks with Uncertainty," The 26th International Conference on Automated Planning and Scheduling (ICAPS), June 2016.
7. E. Timmons, T. Vaquero, B. C. Williams, R. Camilli, “Risk-aware Planning Executive for Autonomous Underwater Gliders,” International Conference on Automated Planning and Scheduling (ICAPS), Workshop on Planning and Robotics (PlanRob), 2016.
8. L. M. Zapata, A. Mallio, P. Ridao, R. Camilli, O. Pizarro, “Preliminary results on terrain-aided localization algorithm for glider AUV,” IEEE International Conference on Robotics and Automation, Workshop on Marine Robot Localization and Navigation. Stockholm Sweden, 2016

Presentations / Conferences to date

- C. L. R. McGhan, “A Risk-Aware Architecturefor Res ilient Spacecraft Operations”. IEEE Aerospace Conference, 12 March 2015.
- C. L. R. McGhan, “Application of Correct-by-Construction Principles for a Demonstratively Resilient Risk-Aware Architecture”. IEEE Space Forum 2015, 1 September 2015.
- C. L. R. McGhan, “Resilient Risk-Aware Autonomy for the Exploration of Uncertain and Extreme Environments”. University of Cincinnati invited talk, 15 January 2016.
- C. L. R. McGhan, “The Resilient Spacecraft Executive: An Architecture for Risk-Aware Operations in Uncertain Environments”. AIAA SPACE Conference, Long Beach, CA, 15 September 2016.
- C. L. R. McGhan, “Towards Architecture-wide Analysis, Verification, and Validation for Total System Stability During Goal-Seeking Space Robotics Operations”. AIAA SPACE Conference, Long Beach, CA, 15 September 2016.

- T. Vaquero, “Risk-aware Planning in Hybrid Domains: An Application to Autonomous Planetary Rovers ”. AIAA SPACE Conference, 15 September 2016.
- P. Santana, “PARIS: a Polynomial-time, Risk-Sensitive Scheduling Algorithm for Probabilistic Simple Temporal Networks with Uncertainty,” ICAPS Conference, London, UK, 16 June 2016.
- E. Timmons, “Risk-aware Planning Executive for Autonomous Underwater Gliders,” PlanRob workshop at ICAPS Conference, London, UK, 13 June 2016.
- T. Vaquero, “Risk-aware Planning for Resilient Spacecraft Systems,” Invited talk at the University of Huddersfield, UK, 9 June 2016.
- L. M. Zapata et al, “Preliminary results on terrain-aided localization algorithm for glider AUV,” IEEE International Conference on Robotics and Automation, Workshop on Marine Robot Localization and Navigation. Stockholm Sweden, 20 May 2016

Undergraduate students, graduate students and postdocs who have worked on the project

- Catharine McGhan, CMS postdoctoral fellow, Caltech (co-located at JPL)
- Tiago Vaquero, CMS postdoctoral fellow, Caltech (co-located at MIT)
- Eric Timmons, PhD research assistant, MIT (supported through the WHOI-MIT Joint Program)
- Cheng Fang, PhD Student, MIT, 2015 intern at JPL
- Riashat Islam, University College London undergraduate, 2015 SURF fellow
- Sandra Liu, Caltech undergraduate, 2015 SURF fellow and 2016 SURF fellow
- Yuh Shyang (Mickey) Wang, CDS graduate student, Caltech