

Resilient Risk-Aware Autonomy

Mitch Ingham

Jet Propulsion Laboratory, California Institute of Technology

3/14/2018

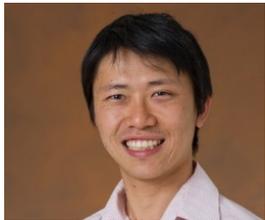
Team (1)



Dr. Michel Ingham (313D)
R&TD PI



Dr. Tara Estlin (398E)



Dr. Masahiro Ono (347E)



Dr. Leslie Tamppari (3220)



Prof. Richard Murray, KISS PI



Prof. Brian Williams



Dr. Richard Camilli

Team (2)



Ravi Lanka
398



Trevor Reed
347



Ellie Simonson
Intern in 347



Tiago Vaquero
MIT



Cat McGhan
Caltech/313

Maged Elaasar	313
Klaus Havelund	349
Ian Baldwin	347
Edoardo Bezzeccheri	Intern in 347
Simon Fang	Intern in 347; MIT

Romain Serra	Intern in 347; LAAS-CNRS
Erez Karpas	MIT
Pedro Santana	MIT
Claudio Toledo	MIT/University of São Paulo

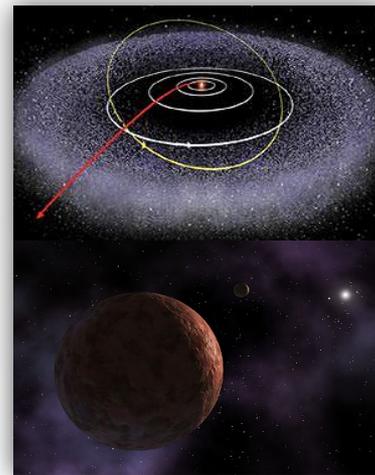
- Destinations are becoming more challenging, science questions more sophisticated, and the most accessible targets are visited
- To advance the knowledge frontier to more interesting, harsh and inaccessible destinations, such as Icy Moons, Venus, Kuiper Belt Objects, and interstellar space, next-gen spacecraft will need to reason about their own state and the state of the environment in order to:
 - predict and avoid hazardous conditions,
 - recover from internal failures, and
 - meet science objectives in the presence of substantial uncertainties



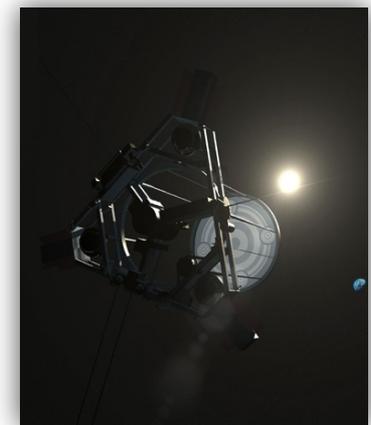
Venus Lander Concept



Europa Lander Concept



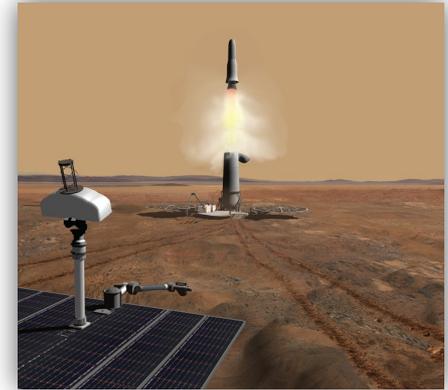
KBO Exploration Concept



**Interstellar Probe
Concept**

Overall Goal of Project

- Develop innovative software architecture, the Resilient Spacecraft Executive (RSE), to endow spacecraft with unprecedented levels of *resilience**
- Demonstrate it on two very different platforms:
 - Surface rover testbed (and high-fidelity simulation)
 - Autonomous underwater vehicle (AUV) testbed
- Impact
 - Enable greater autonomy for robotic exploration of harsh, remote, and inaccessible destinations, e.g., Venus, Outer Planet Icy Moons, and KBOs
 - Reduce operational risk and associated cost for increasingly ambitious missions



**Concept for Mars
Sample Return**



**Europa Submarine
Concept**

* During the KISS Phase 1 study, several definitions of resilience were discussed. The common conceptual core that emerged: ***adaptation in the presence of changing circumstances.***

- Introduction to Risk-Aware Autonomy
- Resilient Spacecraft Executive Architecture
- Algorithm Descriptions
- Demo (RSE Running in Simulation)
- Deployments (RSE Running in AUV System)
- Summary of Benefits

Overview of Risk-aware Autonomy

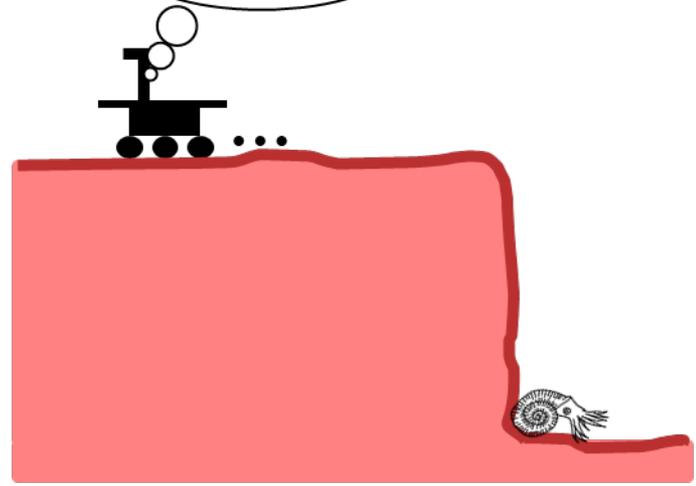


Keep the risk low.

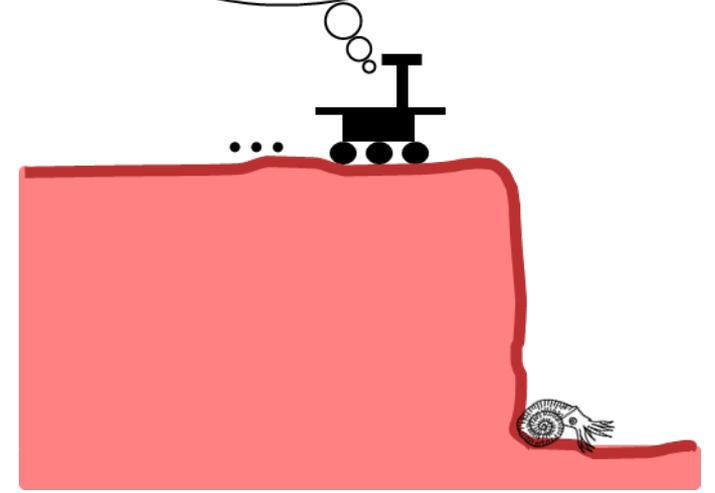


You can take more risk.

Let's stay away from the cliff... it is too dangerous.



Let's check out what's down there...



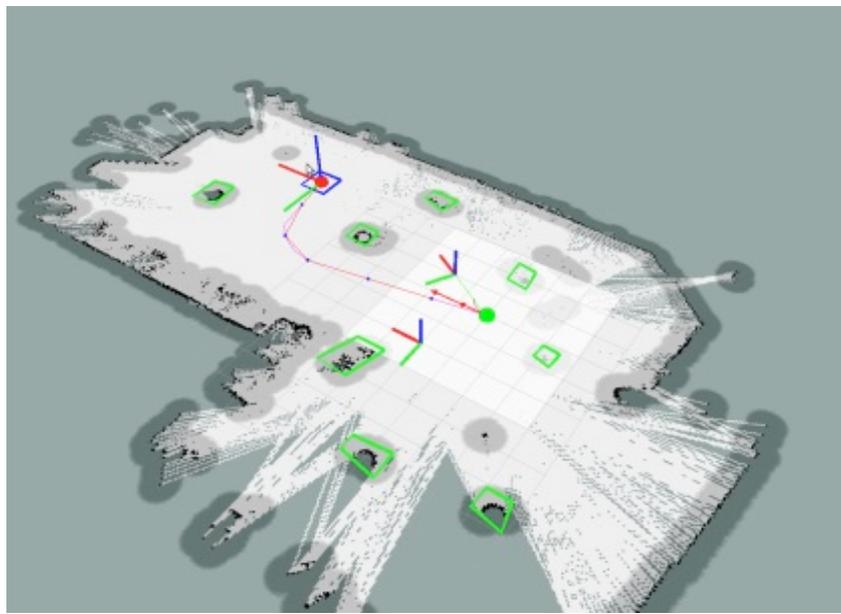
Overview of Risk-aware Autonomy



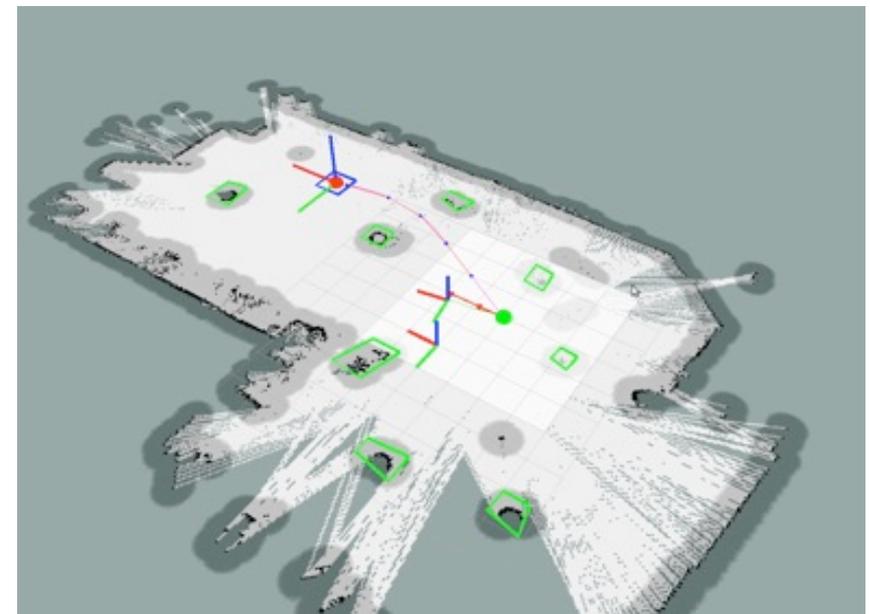
Keep the risk low.



You can take more risk.



Low Risk



High Risk

Overview of Risk-aware Autonomy



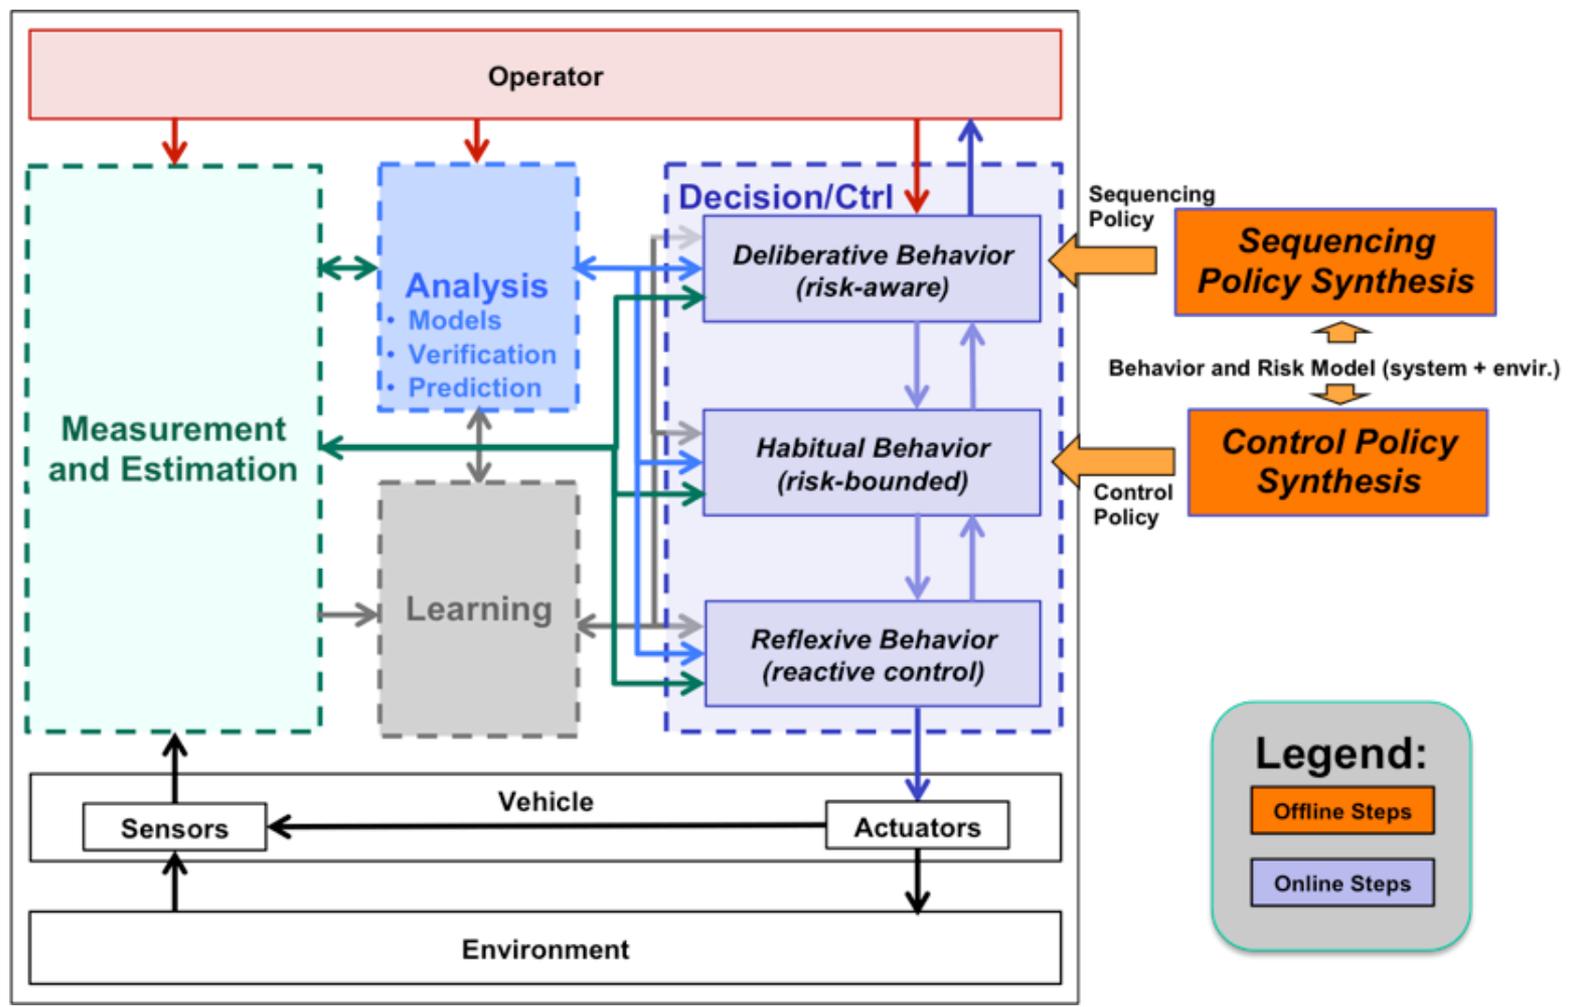
Keep the risk low.

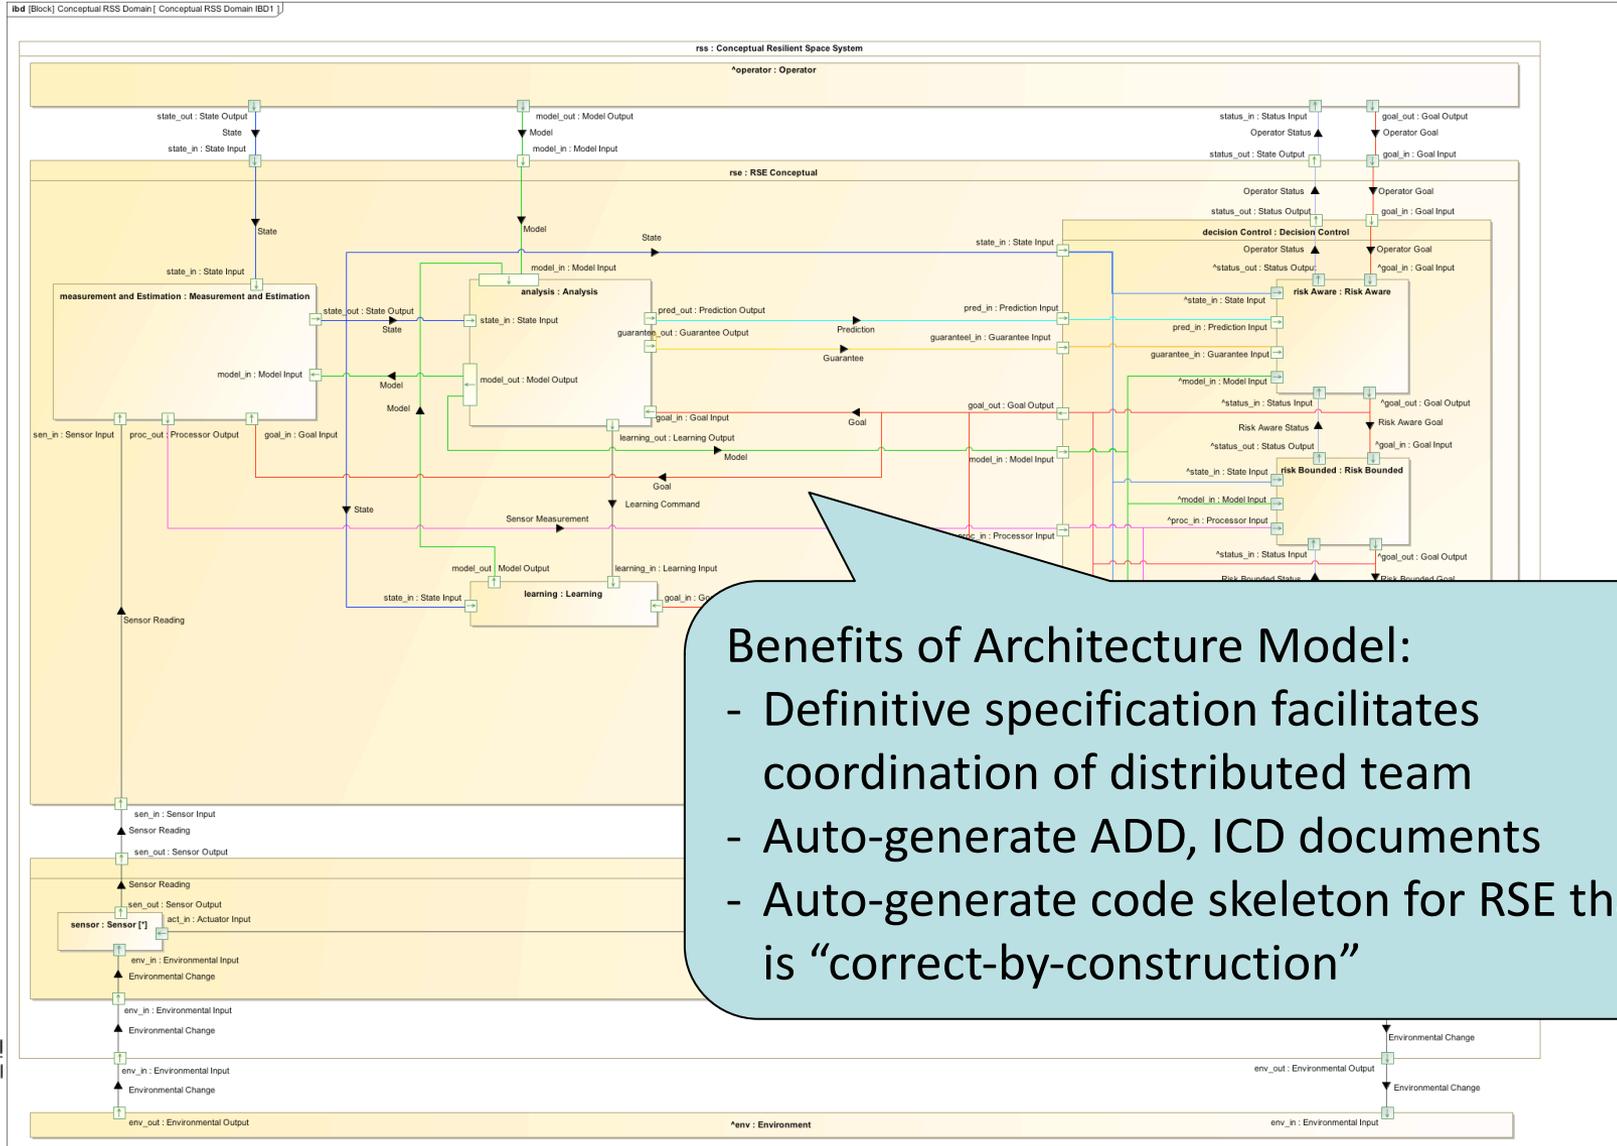


You can take more risk.



RSE Architecture





Benefits of Architecture Model:

- Definitive specification facilitates coordination of distributed team
- Auto-generate ADD, ICD documents
- Auto-generate code skeleton for RSE that is “correct-by-construction”

Key Distinguishing Features & Innovations



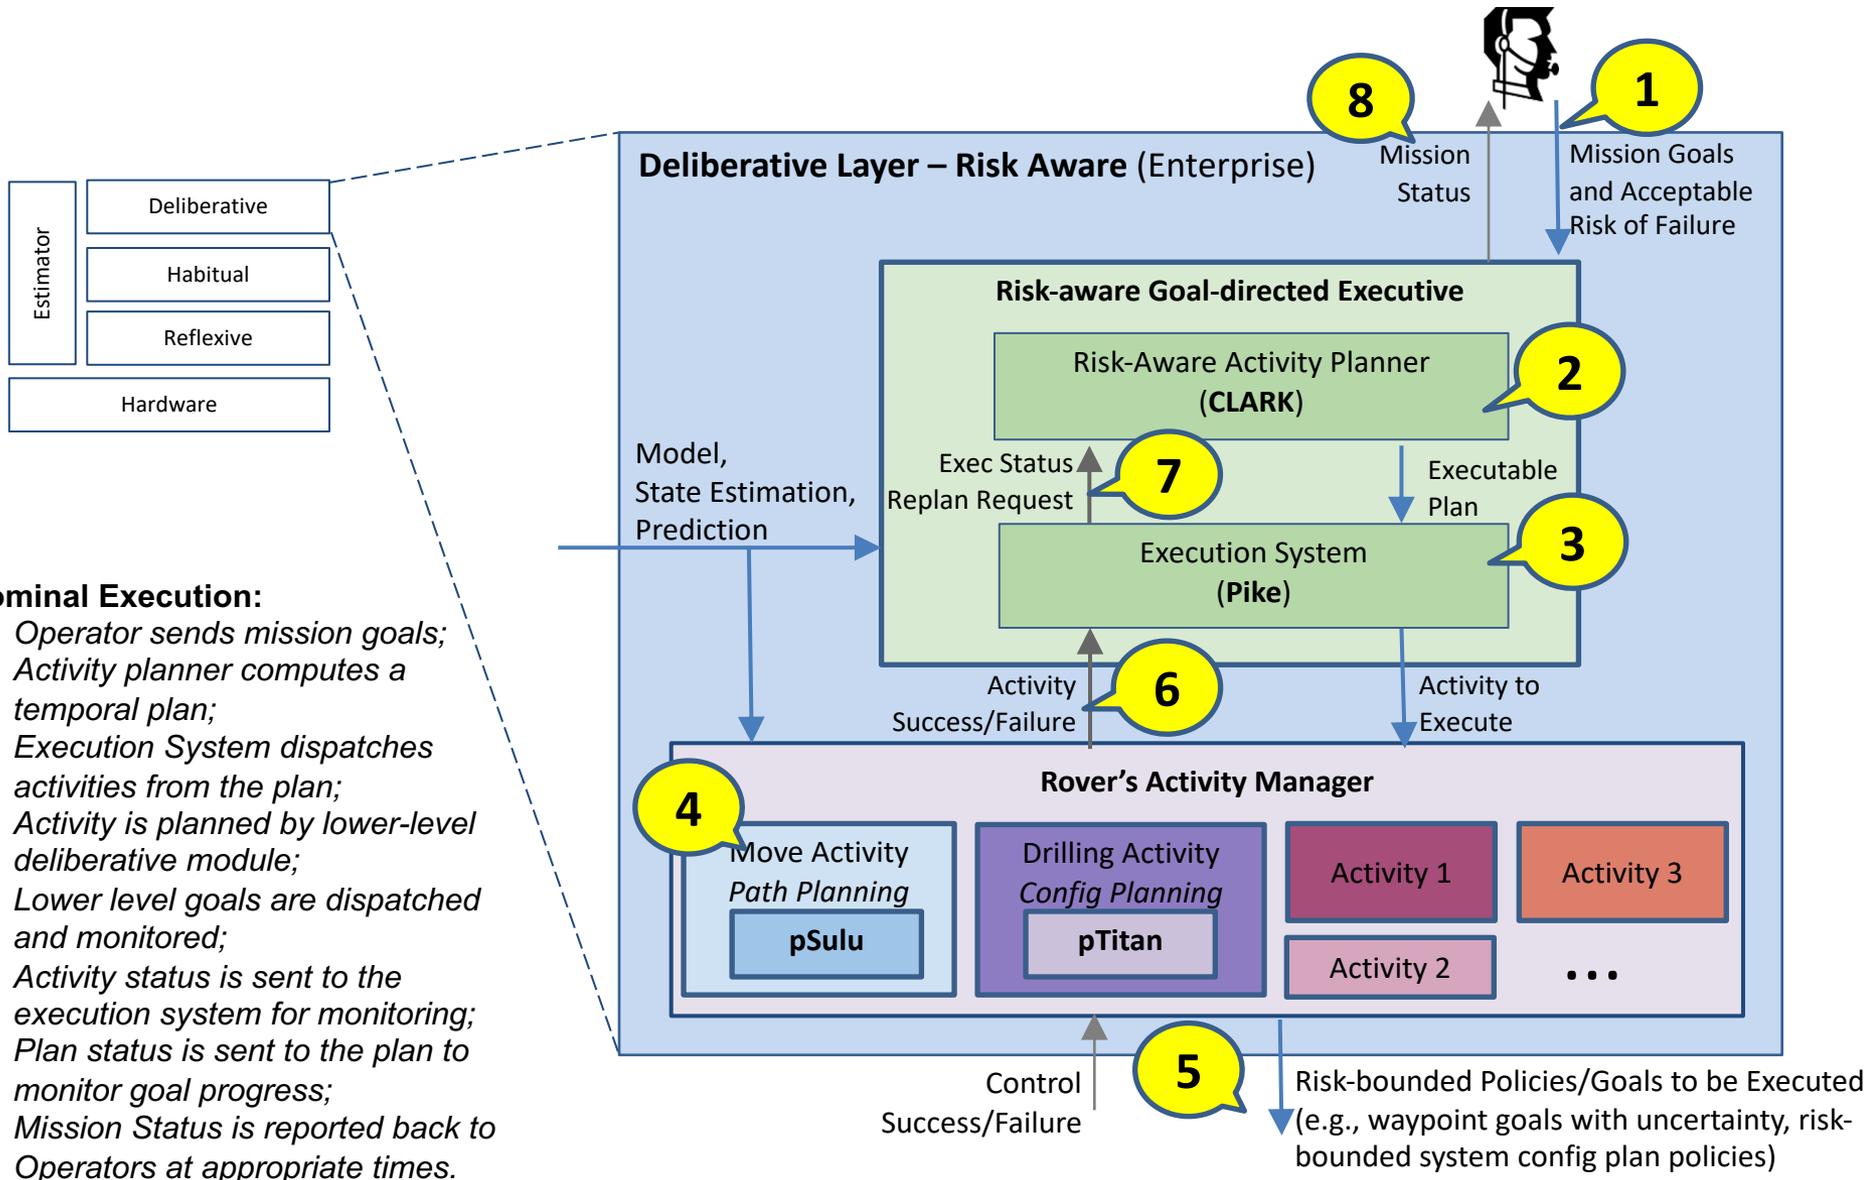
- (i) Sequencing and control policies that are “correct by construction”
 - Use of model-based policy synthesis addresses challenge of ***assuring correctness*** of system behavior in the face of growing complexity.

- (ii) Risk-aware onboard deliberative reasoning
 - Critical to ***managing unprecedented uncertainty*** in environments to be explored in future missions, and ***managing space of possible executions*** far too large to be completely covered by design-time control policies.

- (iii) Formal architectural analysis to perform tradeoffs and inform appropriate allocation of capabilities to architectural modules
 - Result in systems with flexibility to adapt to uncertain environments and potential mission changes. [Analysis capability still in work.]

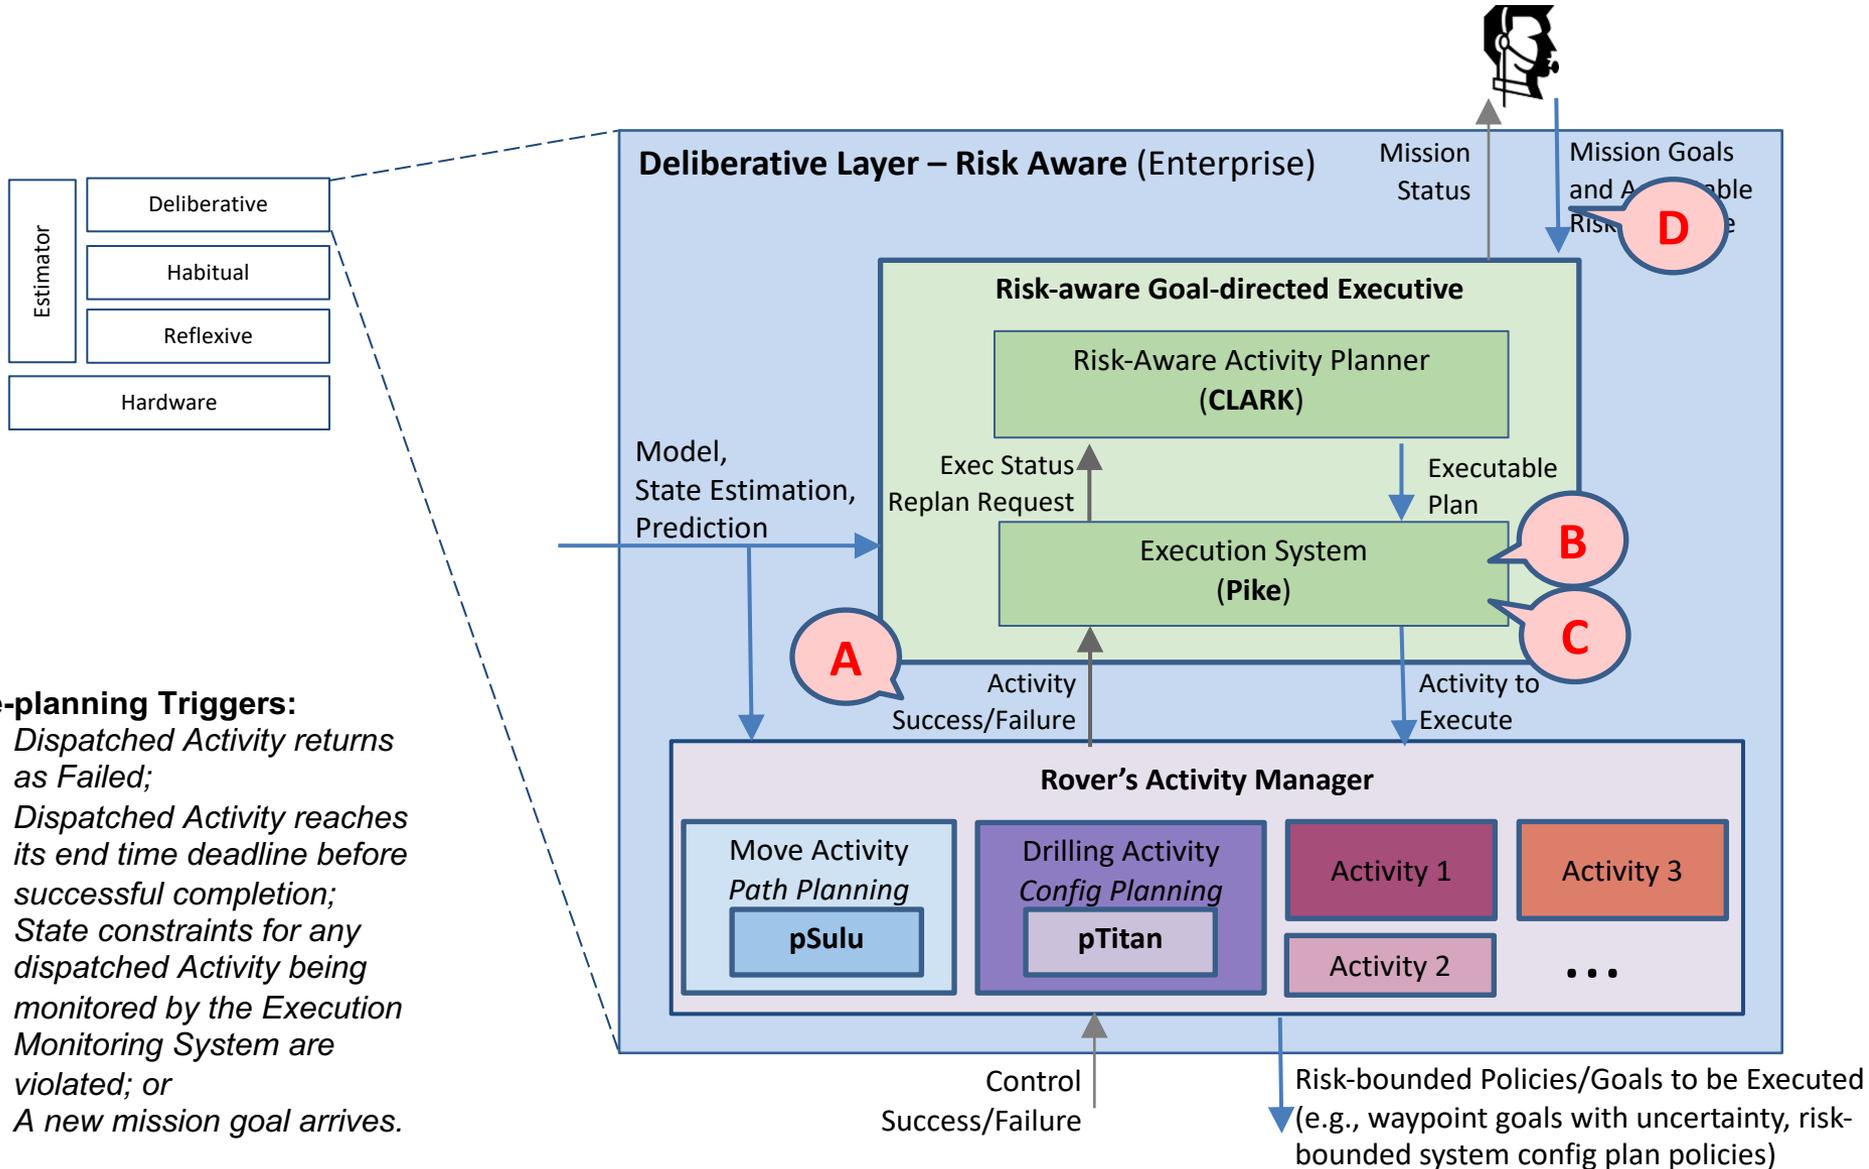
Resilient Spacecraft Executive

Risk-aware Goal-directed Executive

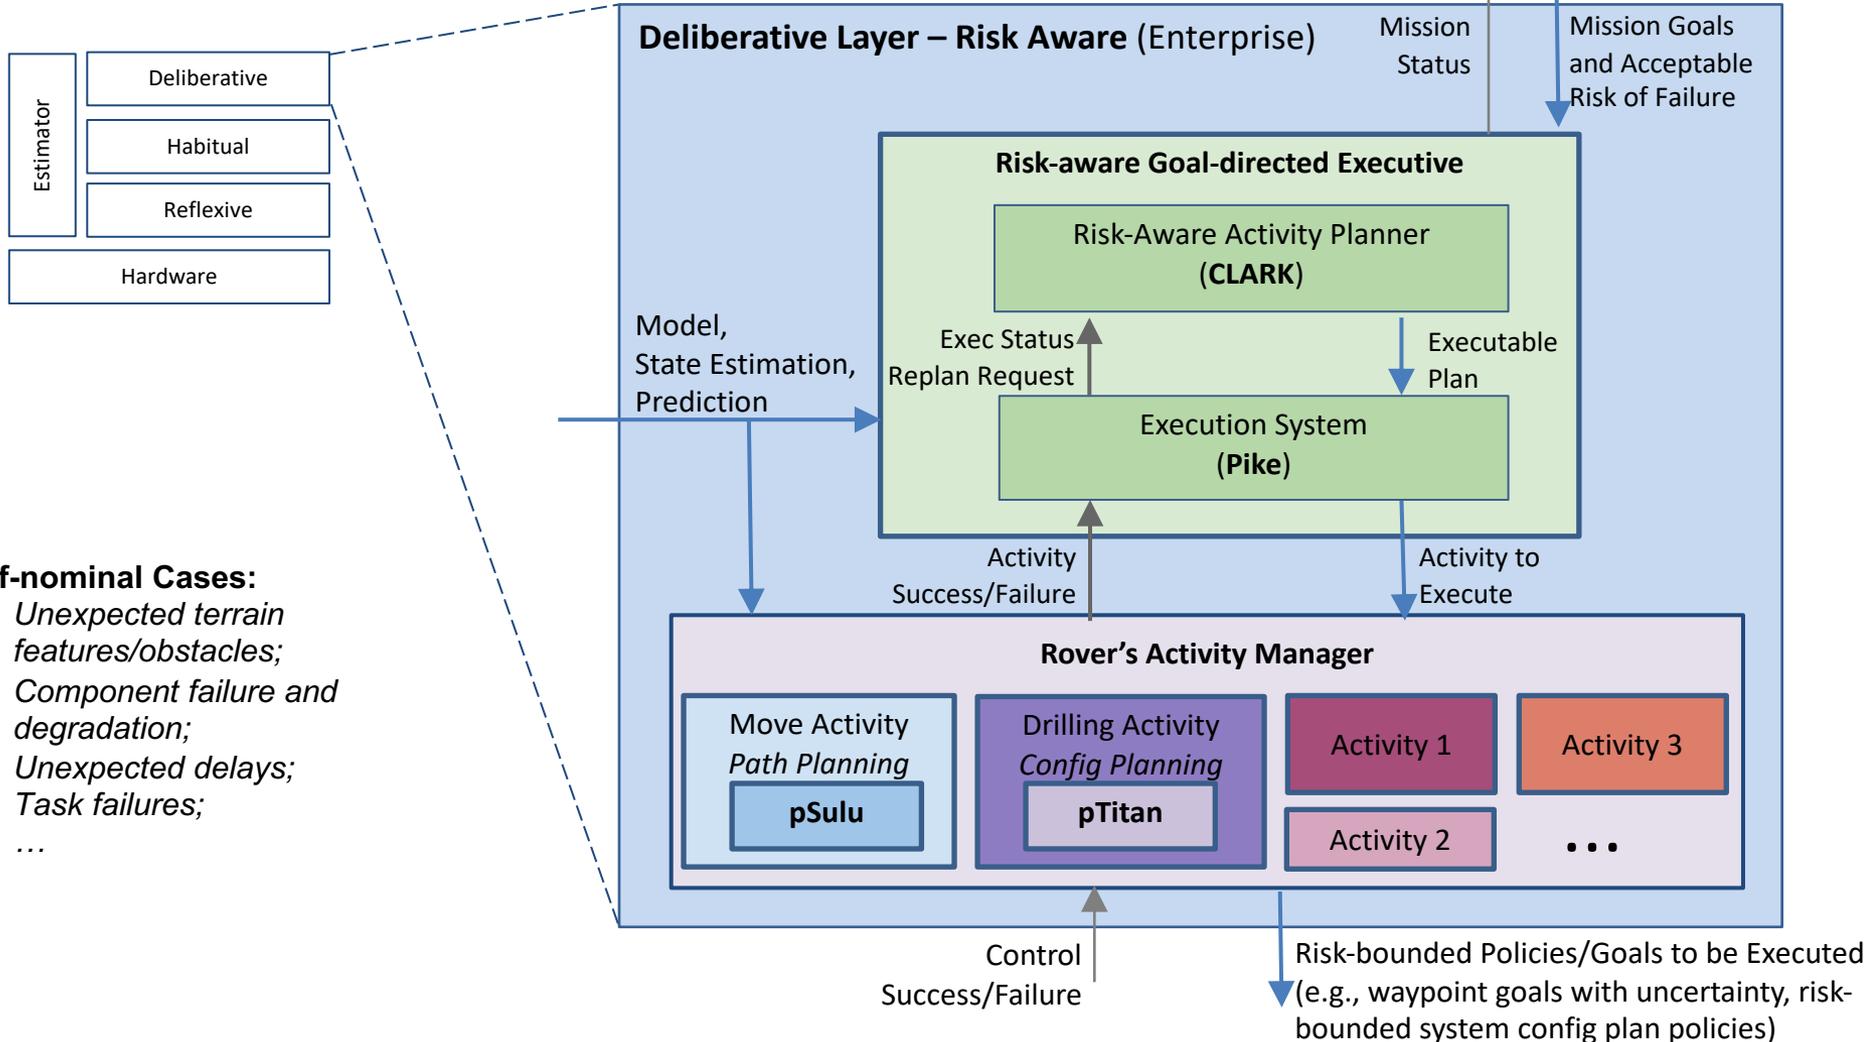


Resilient Spacecraft Executive

Risk-aware Goal-directed Executive

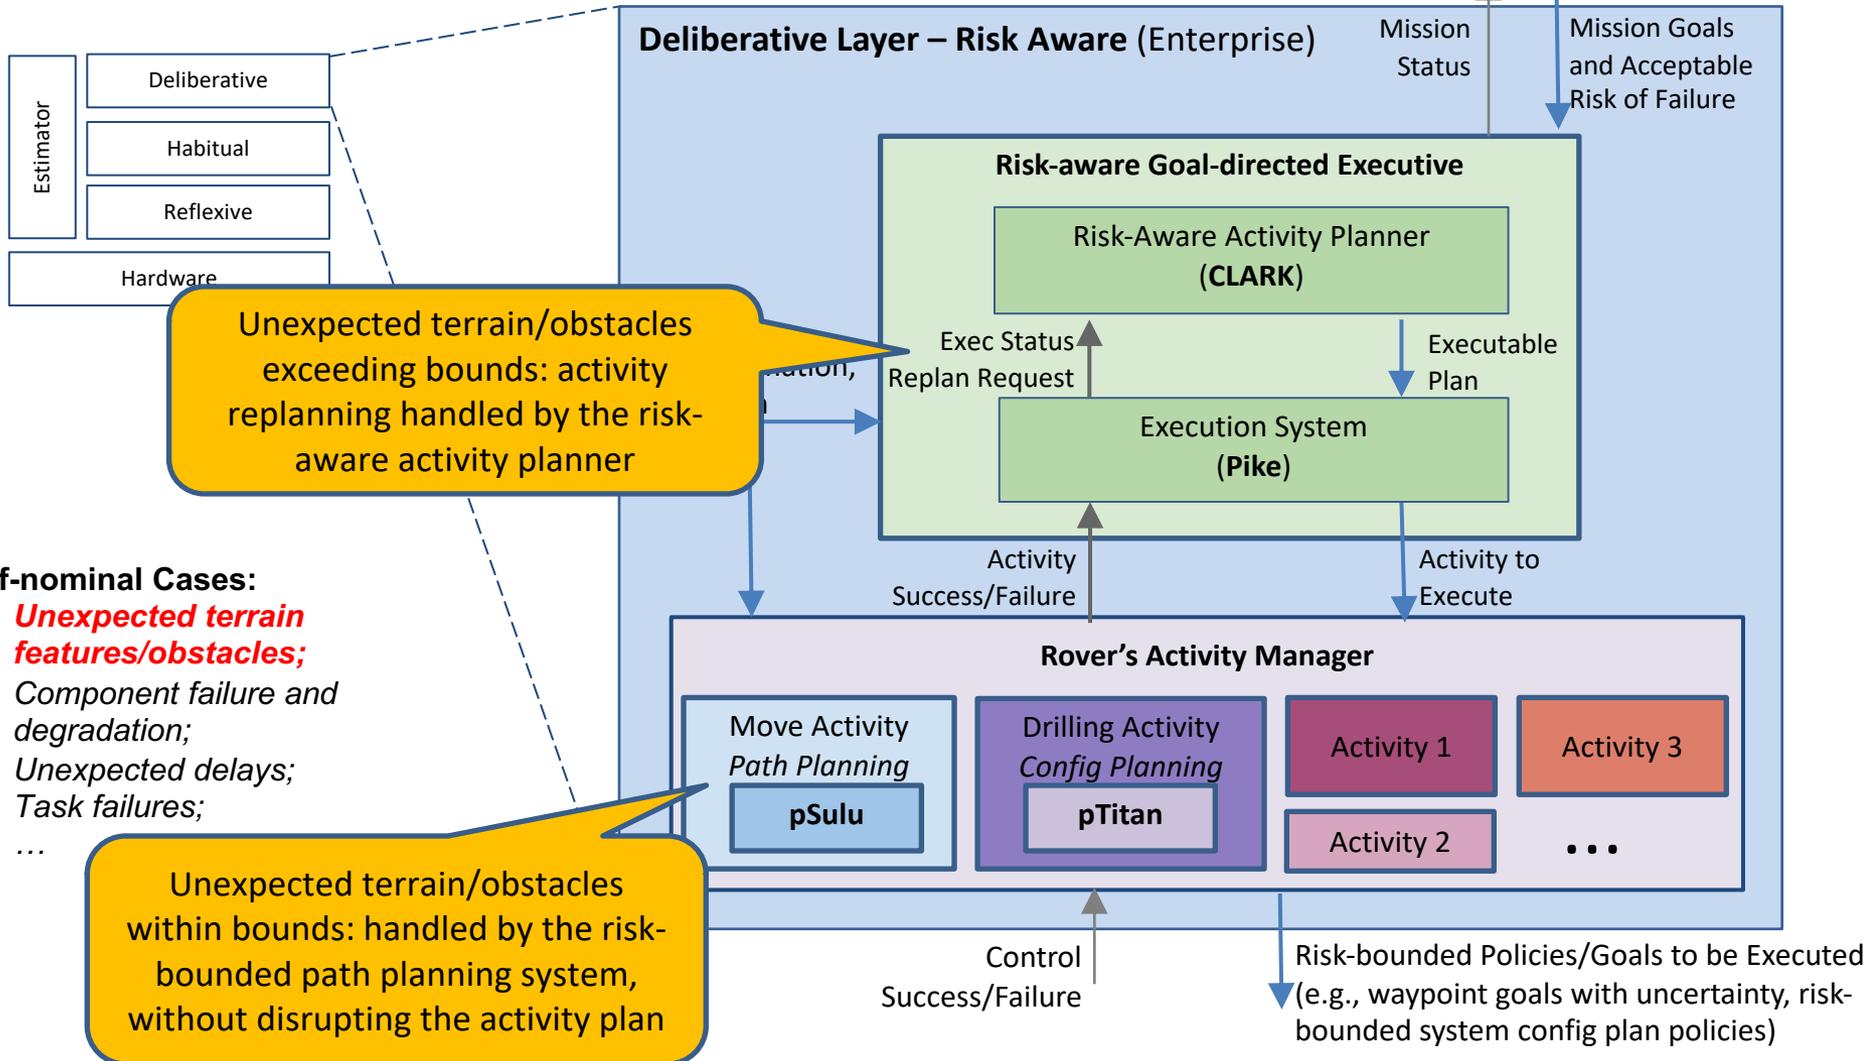


Handling Off-Nominal Cases: Robustness throughout the Architecture

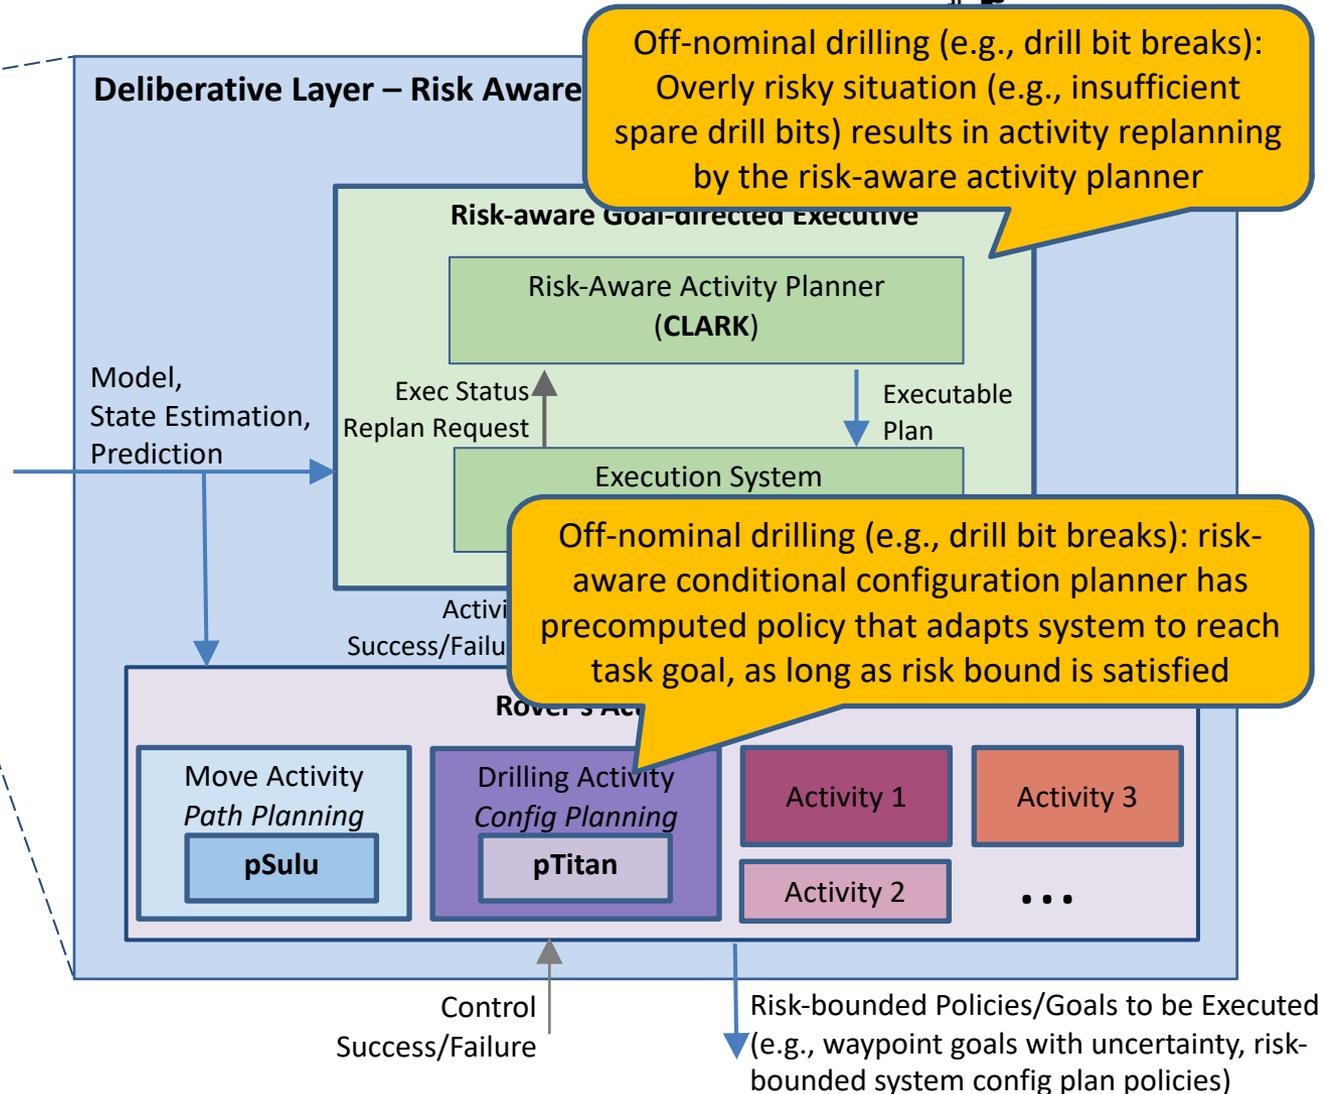
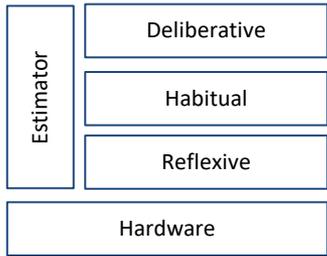


- Off-nominal Cases:**
- Unexpected terrain features/obstacles;
 - Component failure and degradation;
 - Unexpected delays;
 - Task failures;
 - ...

Handling Off-Nominal Cases: Robustness throughout the Architecture



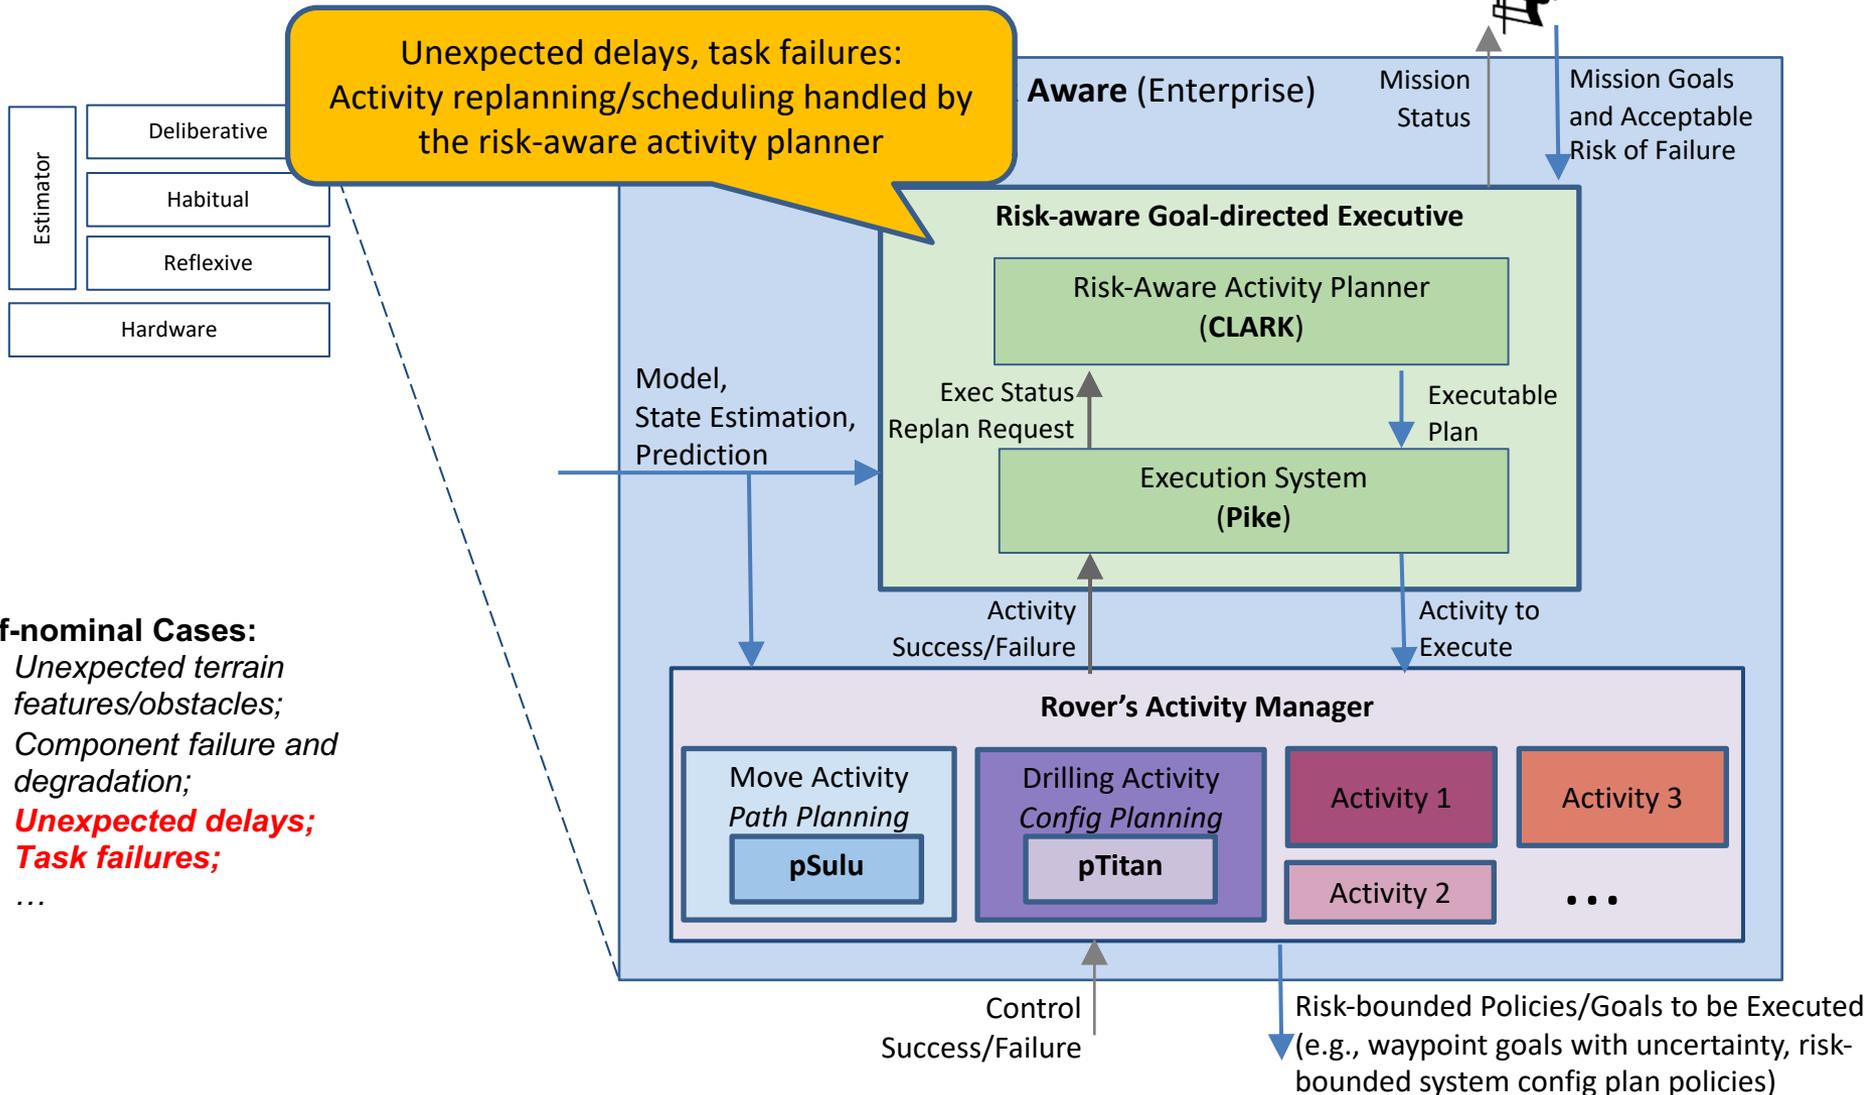
Handling Off-Nominal Cases: Robustness throughout the Architecture



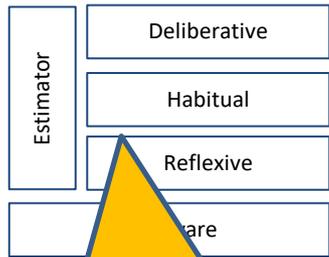
Off-nominal Cases:

- Unexpected terrain features/obstacles;
- **Component failure and degradation;**
- Unexpected delays;
- Task failures;
- ...

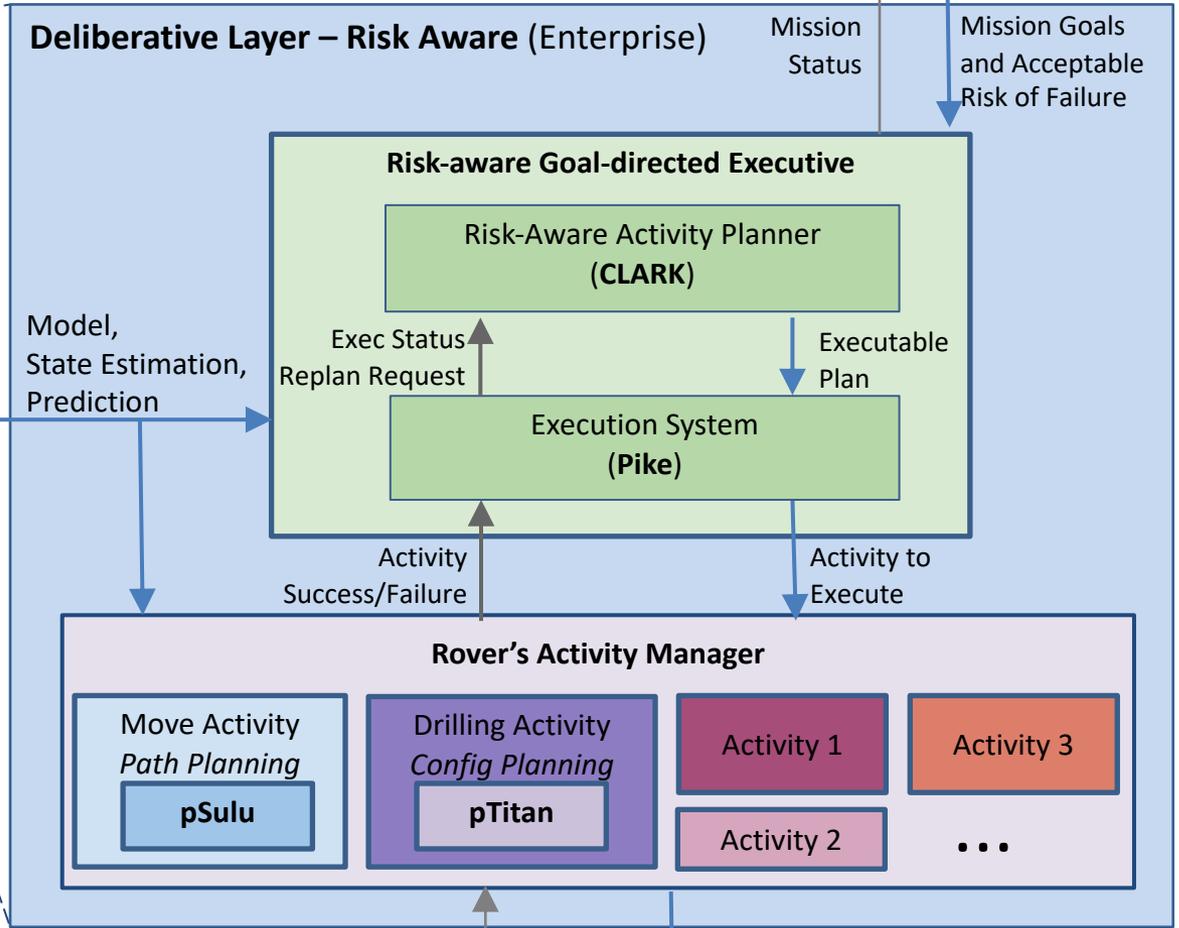
Handling Off-Nominal Cases: Robustness throughout the Architecture



Handling Off-Nominal Cases: Robustness throughout the Architecture



These are just cases handled by the Deliberative Module; many more off-nominal situations handled by the habitual and reflexive modules



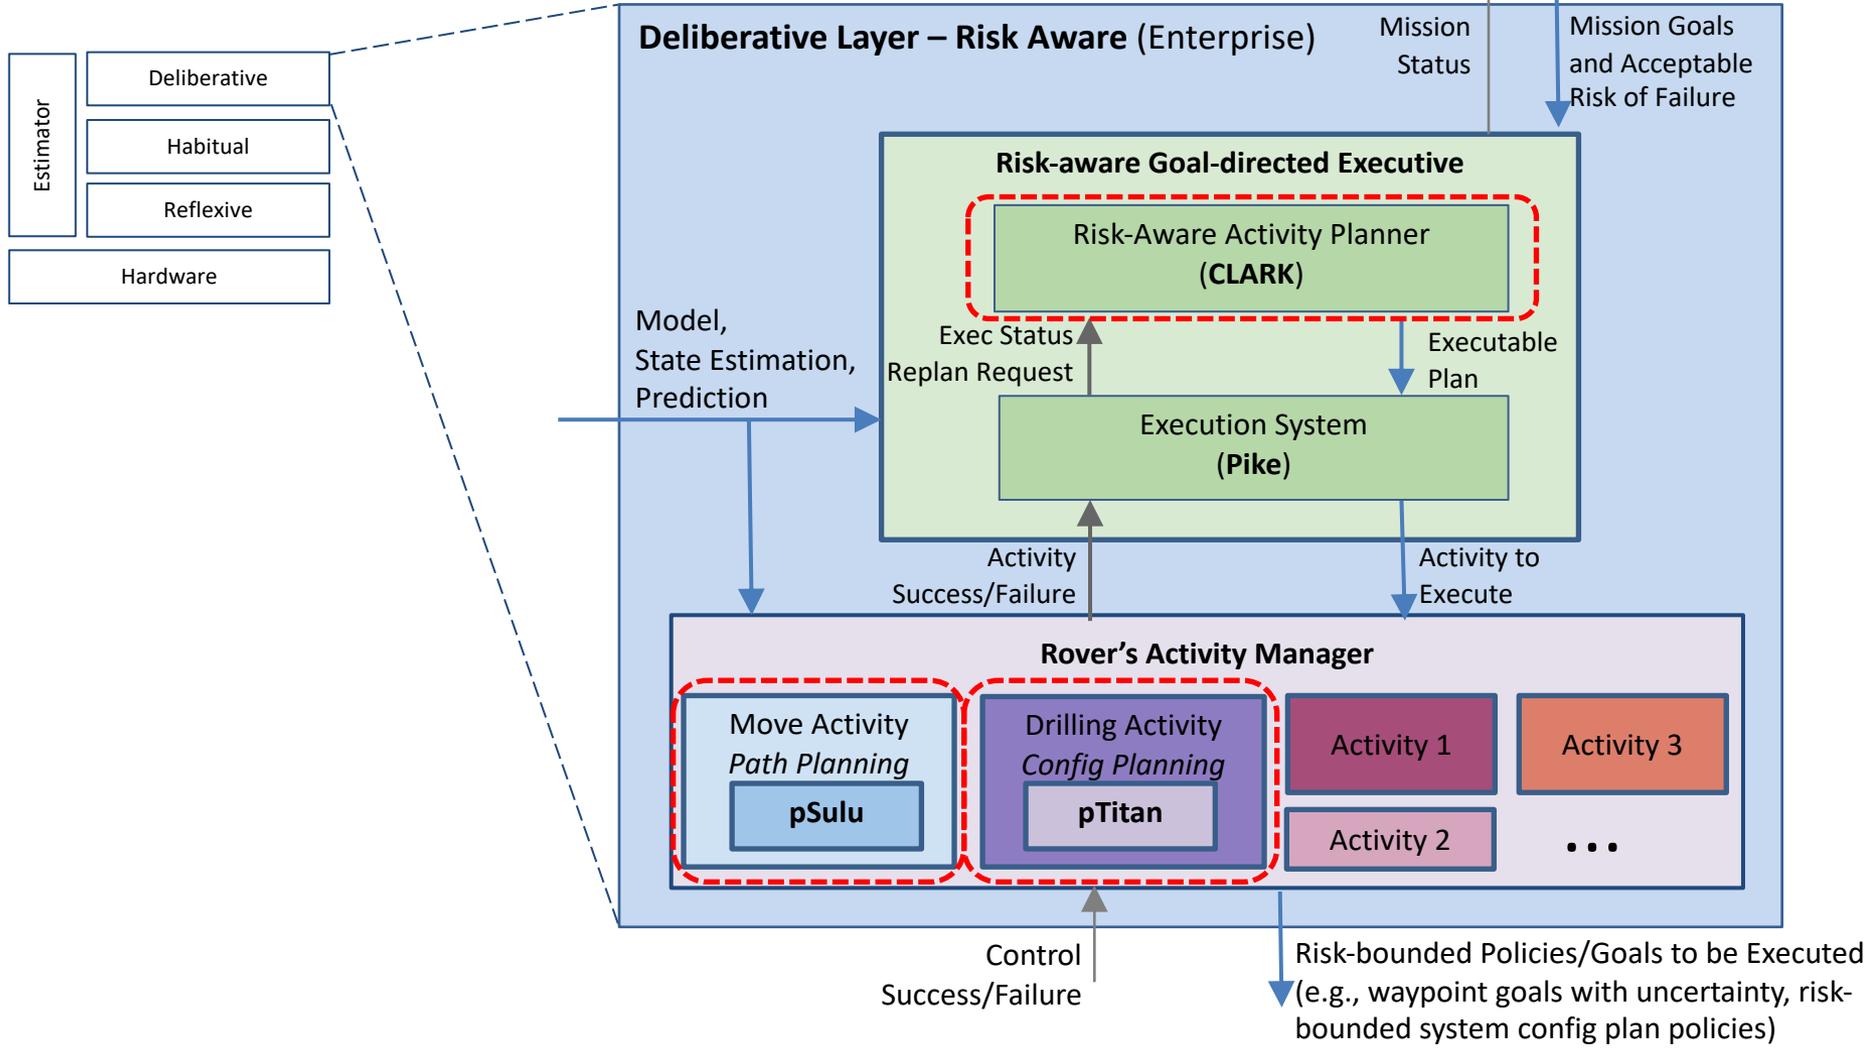
- Off-nominal Cases:**
- Unexpected terrain features/obstacles;
 - Component failure and degradation;
 - Unexpected delays;
 - Task failures;
 - ...

Control Success/Failure

Risk-bounded Policies/Goals to be Executed (e.g., waypoint goals with uncertainty, risk-bounded system config plan policies)

Resilient Spacecraft Executive

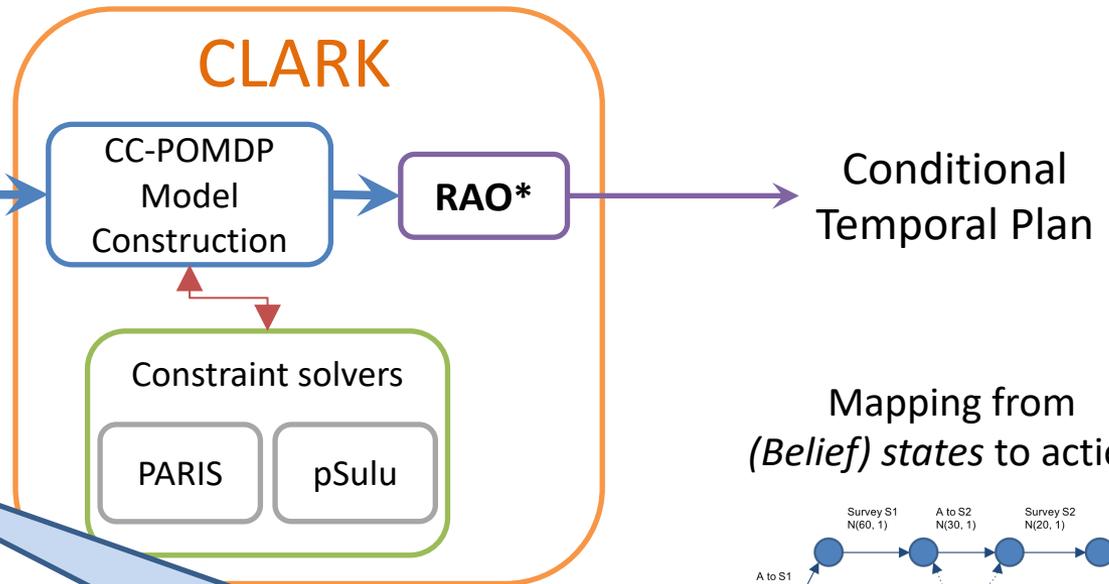
Risk-aware Goal-directed Executive



Risk-aware Activity Planner

- **Planner:** CLARK (Risk-aware, Optimal, Un/conditional and Temporal Planning)
- Core algorithm: Risk-bounded AO* (RAO*)

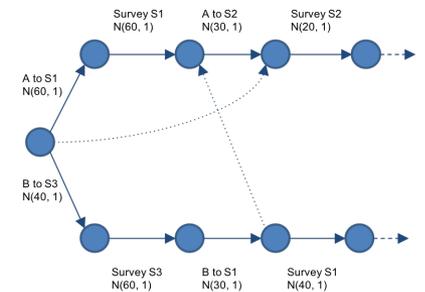
- Initial State
- Mission Goals
- States and Action Model (PDDL)
- Temporal constraints (time windows, deadlines)
- Acceptable risk levels (chance constraints)
- Additional objectives to be considered in optimization



Typical chance constraint:
 “Probability of violating constraints C during execution” $\leq \Delta$ b_{goal}

e.g.: minimize battery consumption,
 or minimize path length

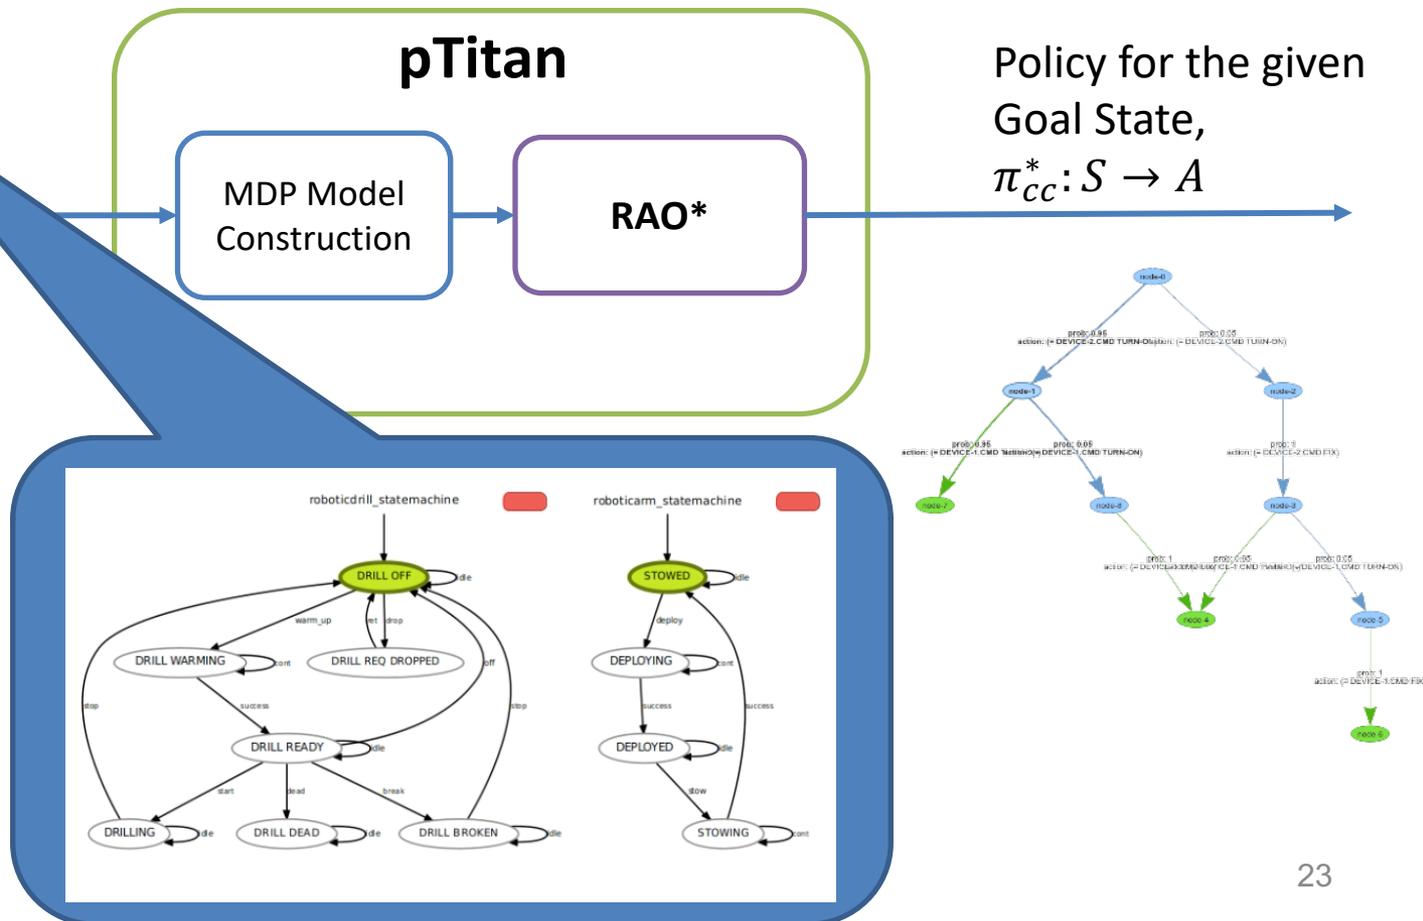
Mapping from
 (*Belief*) states to actions



Risk-aware Configuration Planner

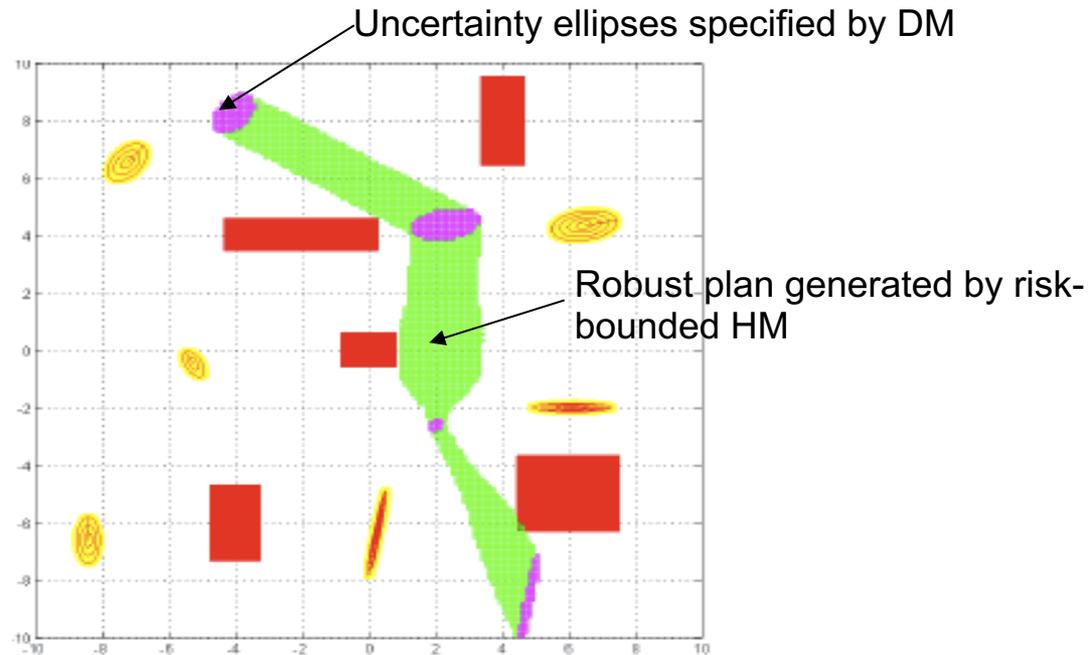
- **Planner:** pTitan (Conditional Configuration State Planning)
- **Core Algorithm:** RAO*

- Probabilistic Concurrent Automata
- Initial State and Goal State
- Chance Constraint (states to avoid)

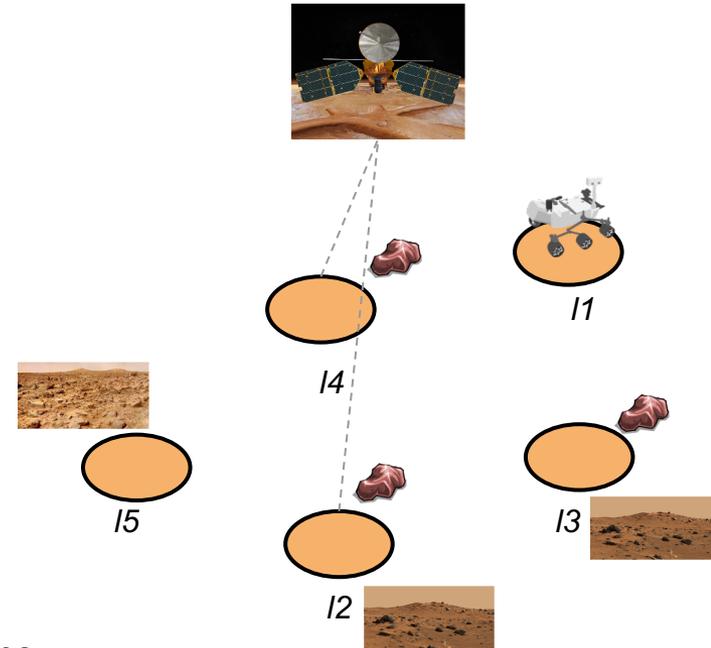


Risk-bounded Path Planning

- **Planner:** Risk-bounded RRT#
- Explicitly considers bounded uncertainty: detailed plan guaranteed to not violate specified constraints
- Based on a robust extension of the RRT# algorithm
- Applied to route planning problems

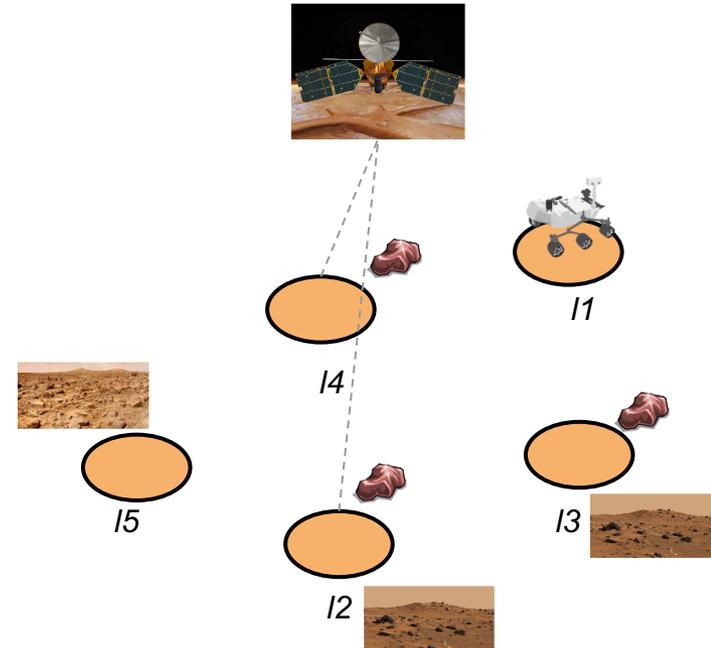


- **Environment:** Mars surface
 - Five target locations (*I1* ... *I5*) for science
- **Agent:** Mars rover
 - Activities:
 - *Navigate*
 - *Turn on cameras*
 - *Take pictures* (Mastcam or Hazcam)
 - *Survey* a location to detect interesting rocks
 - *Collect rock sample*
 - *Transmit data*
- **Temporal Constraints:**
 - Durations of the rover activities are uncontrollable: set-bounded and probabilistic (e.g., Gaussian or Uniform)
 - Data transmission has to happen within a predefined time window (controllable) either from *I2* or *I4*
- **Mission Goals:**
 - Take terrain pictures of two of the target locations.
 - Drill/Collect two rock samples from any two target locations
 - Send the data collected to the orbiter.
- **Risk:** bound on risk of failure (missing communication window, collision, uncompleted science goals)



Resilience Test Cases:

- **Environmental uncertainty / leverage plan flexibility:**
 - Fail to locate interesting rock sample to collect in a target location. Replan rock sampling activity to another location.
- **Onboard failures / use of functional redundancy:**
 - Mastcam fails to turn on. Replan imaging activity to use Hazcam instead.
- **Onboard failures / prudent management of limited resources:**
 - Rock sampling drill bit breaks. Depending on risk posture, retry with another drill bit (only 3 available) or replan rock sampling activity to another location.



Initial Plan:

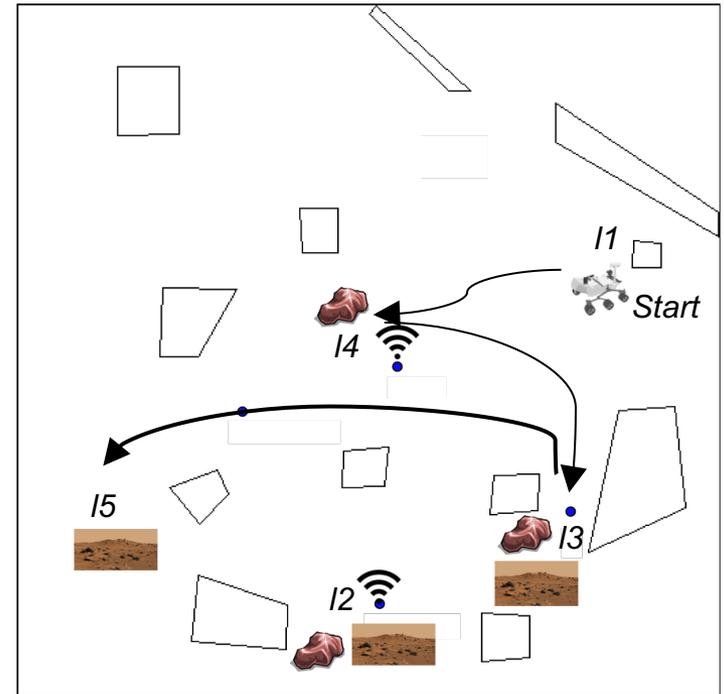
Starting at I1...

1. Visit location I4 to collect a rock sample
2. Visit location I3 for a picture of the terrain and collect a rock sample
3. Visit location I5 for a picture of the terrain
4. Communicate science data back to Earth

Resilience Demo:

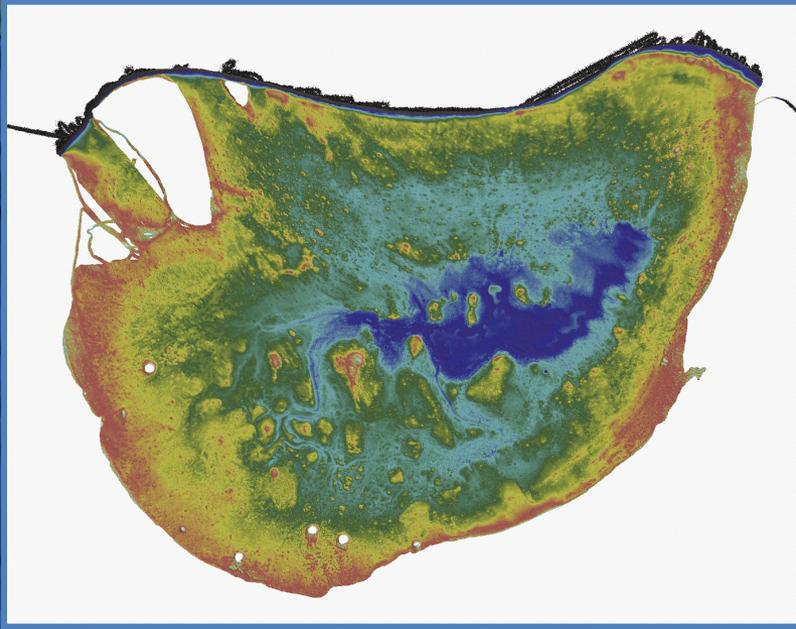
Rover is at location I4, performing the rock sample drilling activity and the drill bit breaks.

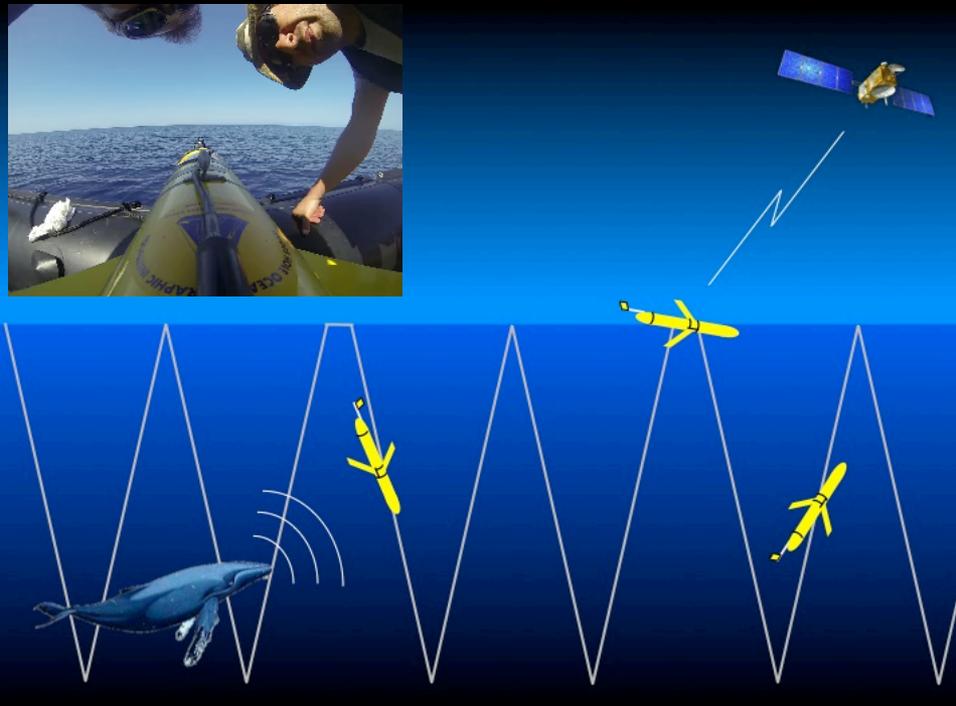
- **Contingency response:** Risk of losing more drill bits in this rock is unacceptably high, so DM requests a replan which will add a traverse to location I2 and a rock sample drilling activity there.



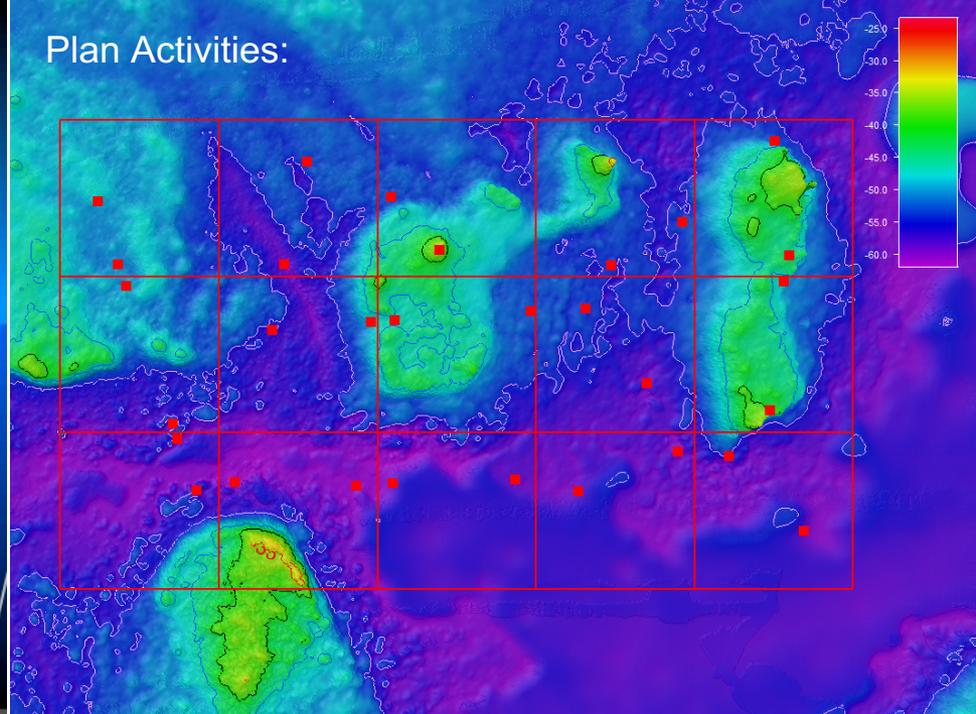
- Technology validation cruise with gliders in Scott Reef, Australia. (March 24 to April 6, 2015)
- Single vehicle mission in Santa Barbara, CA, USA (September 5 to 16, 2016)
- Multi-vehicle coordination in Cape Cod, MA, USA (November 17, 2016)

Goal-directed Planning and Execution in a Risky Environment: Falkor Cruise





Plan Activities:



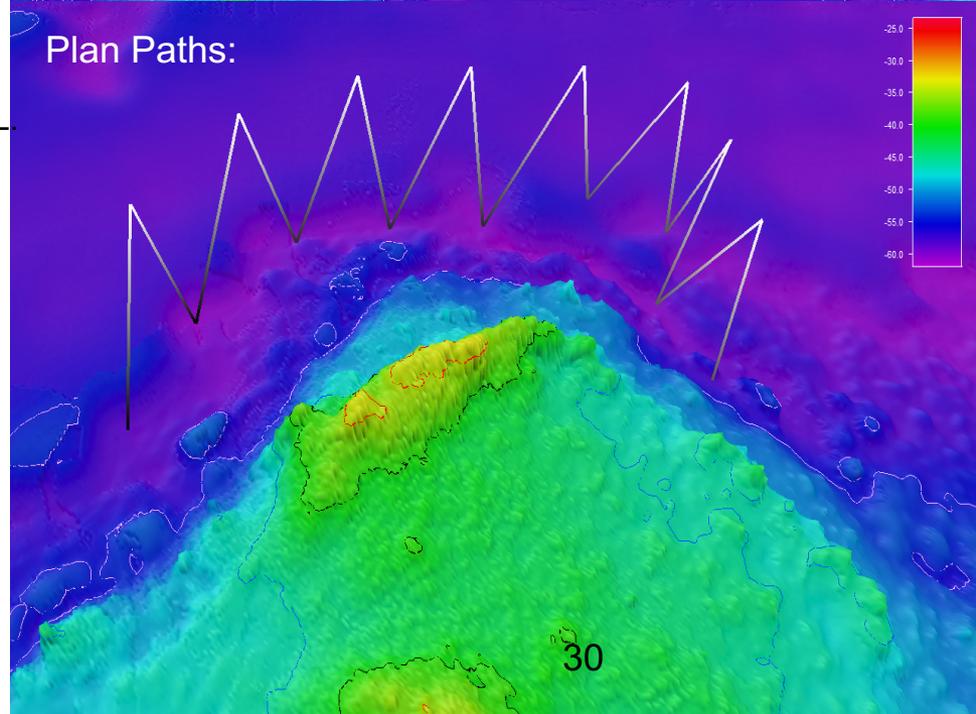
Combined Activity and Path Planning Rules:

- TRANSIT THROUGH BOTH GOAL POINTS IN EACH CELL.
- AVOID FIXED AND MOVING HAZARDS.
- MINIMIZE ENERGY EXPENDITURE.
(BY OPTIMIZE FOR: DEPTH BAND, LINEAR DISTANCE, AND TIDAL CURRENTS).
- STAY WITHIN 2KM COMMS RANGE OF MOTHER SHIP.

Missions Programs in RMPL:

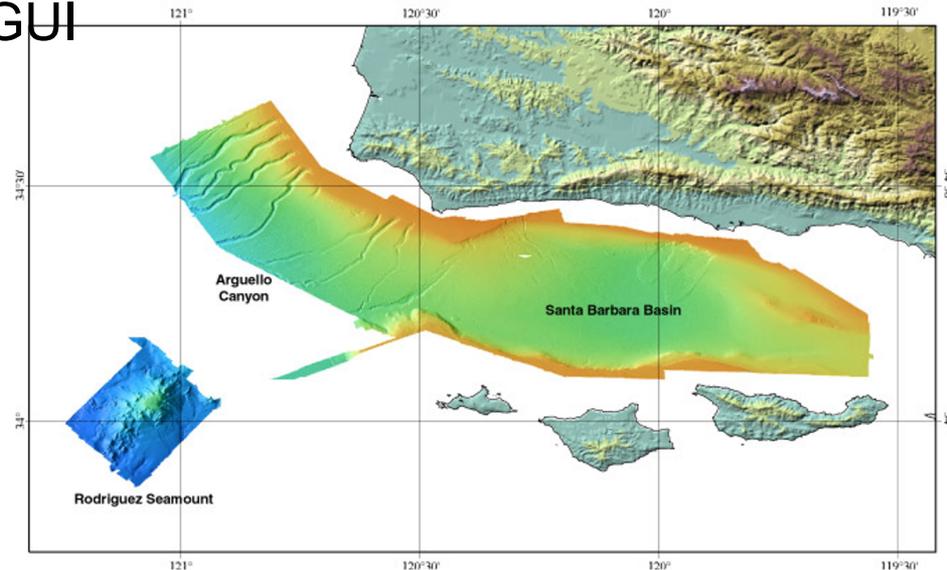
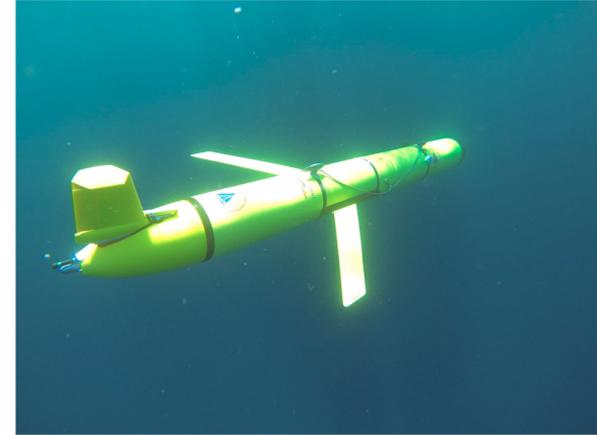
- Parallel threads;
- Goal locations;
- Flexible time bounds;
- Decision-theoretic choice.
- Safety margins, ⇒ more recently risk-bounds.

Plan Paths:

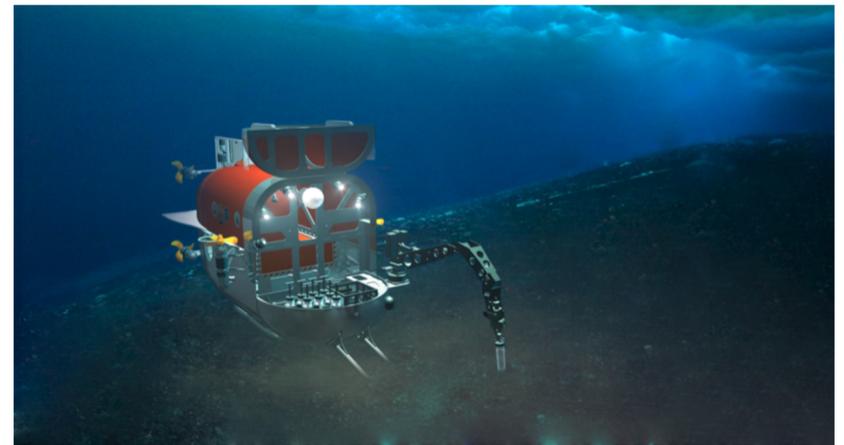
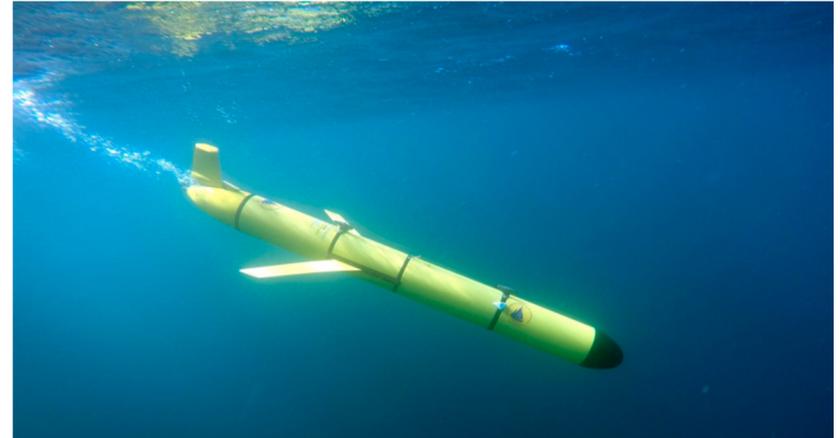
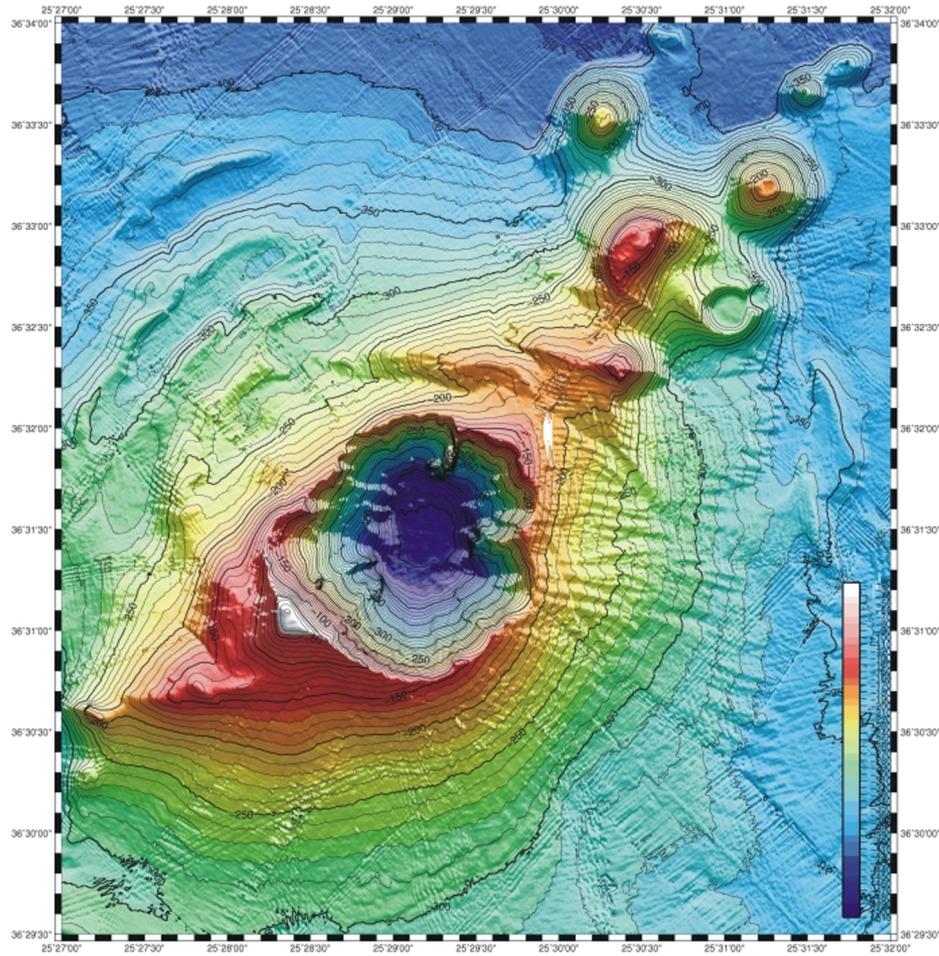


Slocum Glider off Santa Barbara

- Mission goal:
 - Use miniaturized **mass spectrometer** to find and characterize **oil seeps** off the coast of Santa Barbara.
- Major research goal:
 - Combined, activity planning and risk-bounded motion planning.
 - Intuitive user interaction - Web GUI



2019: Europa Analog Mission Demonstration Of Risk-bounded, Autonomous Exploration



Mission: Look for evidence of “extreme” life at
Kolumbo Deep-Sea Volcano near Santorini, Greece

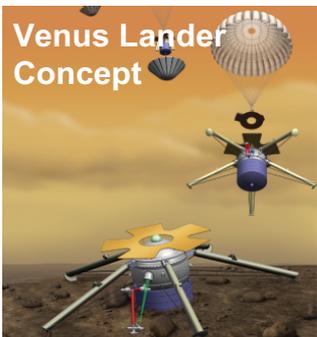
funded by NASA PSTAR Program
Team: WHOI, MIT, ACFR, U. Michigan

- Enable greater autonomy for robotic exploration of harsh, remote, and inaccessible destinations, e.g., Venus, Outer Planet Icy Moons, KBOs
- Reduce operational risk and associated cost for increasingly ambitious missions
- Specifically, resilient risk-aware autonomy can improve science gain by:
 - Adapting to component failures to allow ***graceful degradation***
 - Accommodating environments, science observations, and spacecraft capabilities that are ***not fully known in advance***
 - Making ***risk-aware decisions*** without waiting for ground-based reactions

Backup Slides

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

FY15-16 project jointly funded by JPL and the Keck Institute for Space Studies



Objective: Develop *Resilient Spacecraft Executive* that:

- adapts to component failures to allow graceful degradation
- accommodates environments, science observations, and spacecraft capabilities that are not fully known in advance
- makes risk-aware decisions without waiting for slow ground-based reactions

Why this is important to NASA and JPL:

- enables robotic explorations of harsh, remote, and inaccessible destinations
- reduces operational risk and associated cost

Expected Accomplishments:

FY15: Design and develop core algorithms of RSE; develop formal behavior models; validate algorithms through small-scale demo using simulation, rover testbed in Mars Yard, and AUV submarine.

FY16: Integrate algorithms and behavior models; deploy RSE on simulator/hardware for Venus lander and/or Mars rover scenarios.

JPL Team Members

KISS-funded collaborators

Role	Name	Sec.
PI	Mitch Ingham	312
Co-I	Hiro Ono	347
Co-I	Tara Estlin	398
Co-I	Leslie Tamppari	322



Prof. Richard Murray
(Caltech)



Prof. Brian Williams
(MIT)



Dr. Richard Camilli
(Woods-Hole O.I.)

Overview of Approach and Year 1 Results:

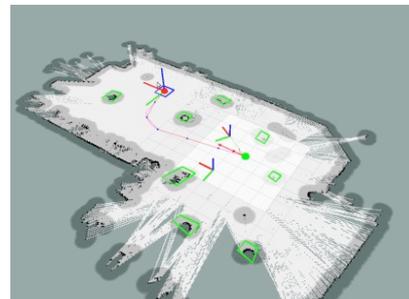
System adapts its behavior depending on acceptable level of risk



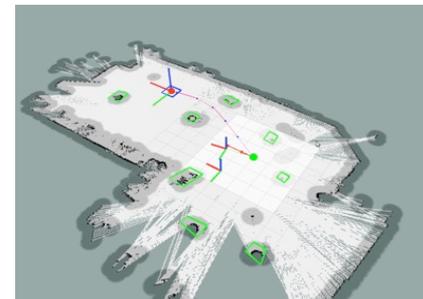
Keep the risk low.



You can take more risk.



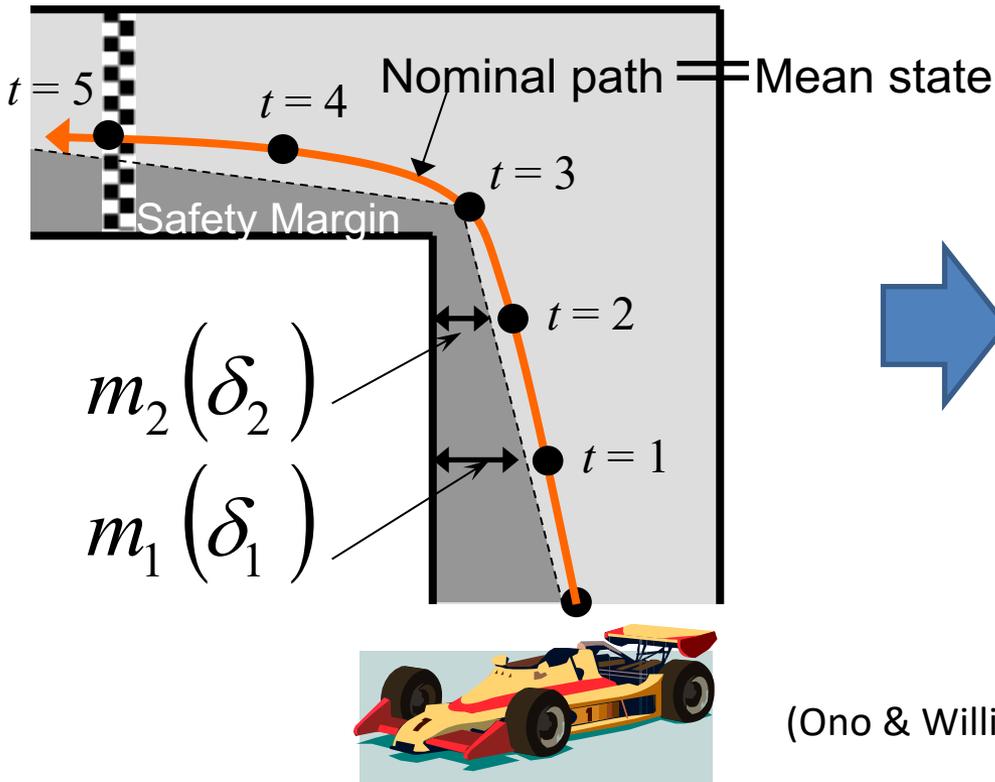
Low Risk



High Risk

Risk-bounded motion planner

Idea: pSulu reformulates risk-bounded path planning to a convex optimization using risk allocation.



$$\min_{\delta} \min_{u_{1:T} \in \mathbf{U}^T} J(u_{1:T})$$

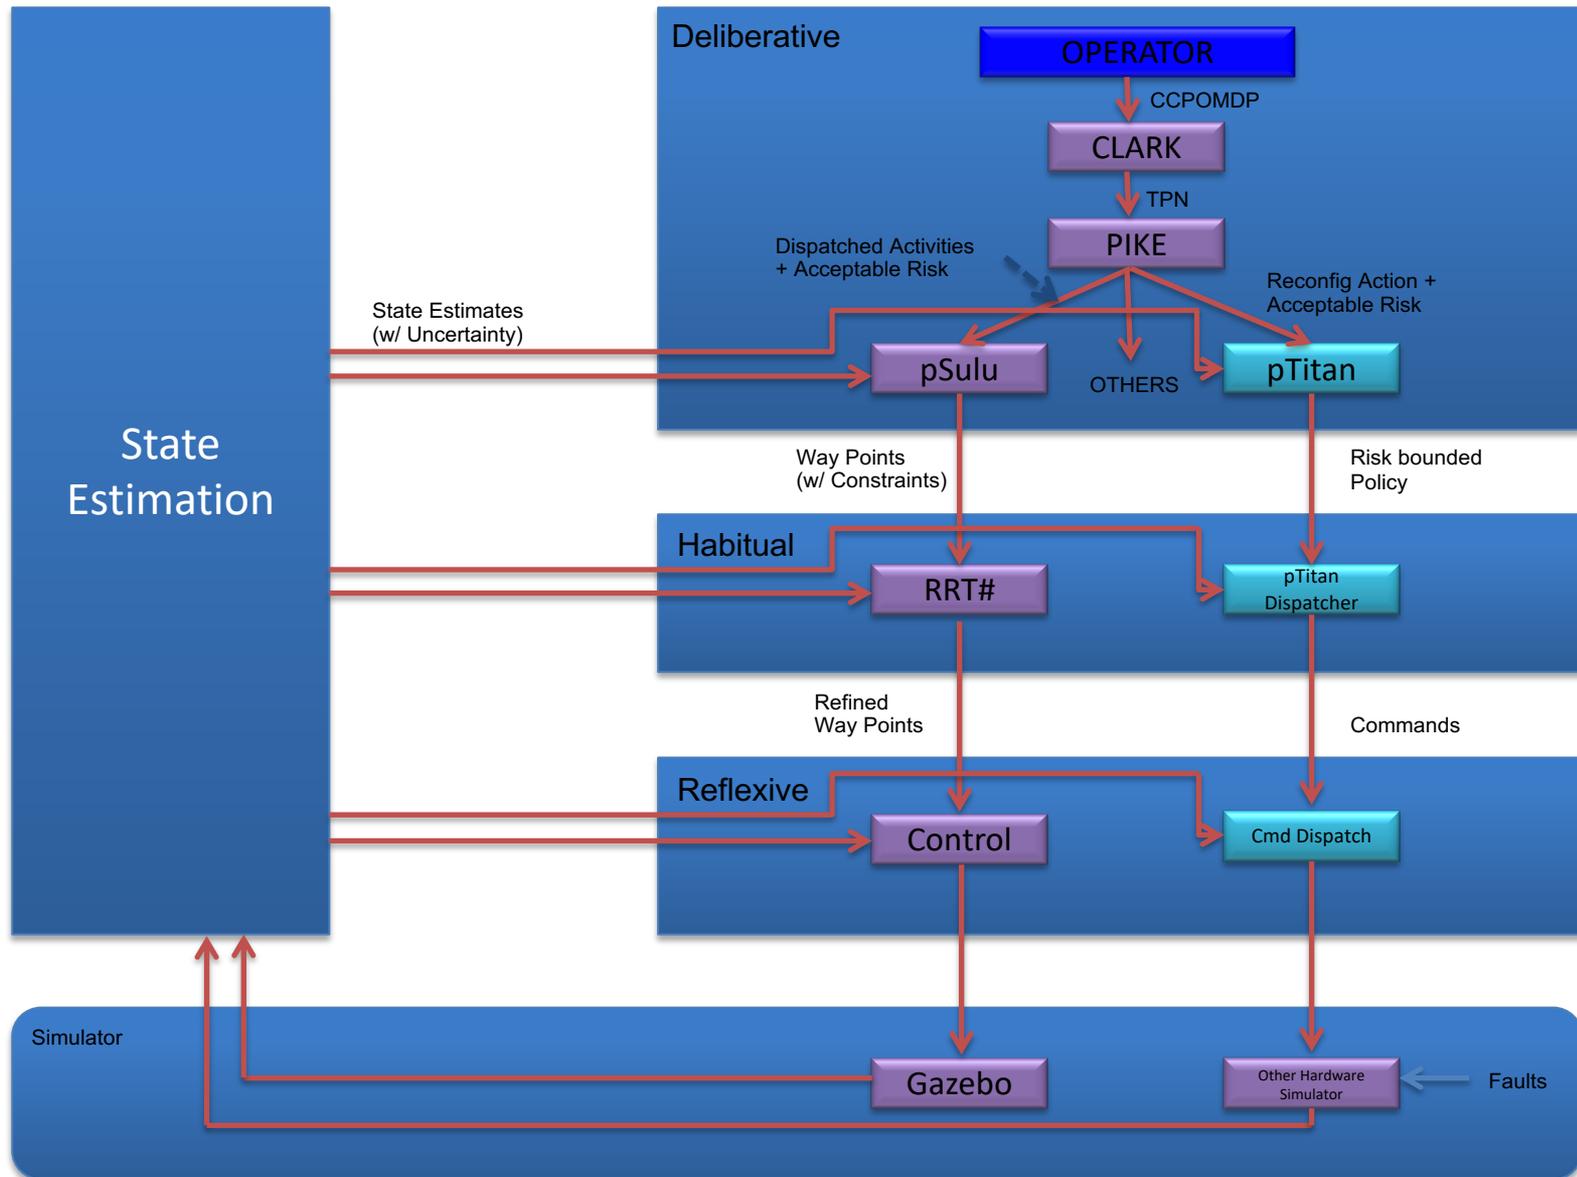
$$\bigwedge_{t=0}^{T-1} \mathbf{x}_{t+1} = \underbrace{A\bar{\mathbf{x}}_t + Bu_t}_{\text{Mean State}}$$

$$\bigwedge_{t=1}^T \bigwedge_{i=1}^I h_t^{iT} \bar{\mathbf{x}}_t \leq \underbrace{g_t^i - m_t^i(\delta_t^i)}_{\text{Safety Margin}}$$

$$\sum_{t,i} \delta_t^i \leq \Delta$$

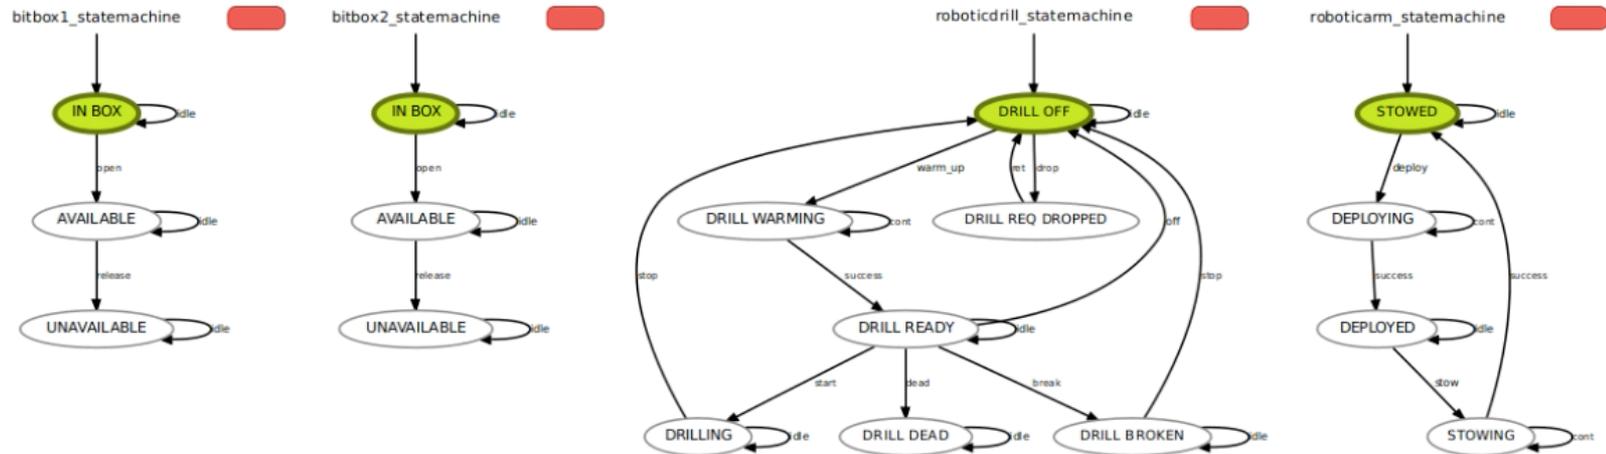
(Ono & Williams, CDC 08; Ono, Williams, & Blackmore JAIR 13)

Demo Architecture



Hardware Behavior State Machines

Drilling State Machine:



Camera State Machine:

