Integrating Formal Verification and Assurance: An Inspection Rover Case Study

Hamza Bourbouh[1], Marie Farrell[3], Anastasia Mavridou[1], Irfan Sljivo[1], Guillaume Brat[2], Louise A. Dennis[3], and Michael Fisher[3]

[1] KBR/NASA Ames Research Center, USA
[2] NASA Ames Research Center, USA
[3] Department of Computer Science, University of Manchester, UK

Abstract. The complexity and flexibility of autonomous robotic systems necessitates a range of distinct verification tools. This presents new challenges not only for design verification but also for assurance approaches. Combining the distinct formal verification tools, while maintaining sufficient formal coherence to provide compelling assurance evidence is difficult, often being abandoned for less formal approaches. In this paper we demonstrate, through a case study, how a variety of distinct formal techniques can be brought together in order to develop a justifiable assurance case. We use the AdvoCATE assurance case tool to guide our analyses and to integrate the artifacts from the formal methods that we use, namely: FRET, COCOSIM and Event-B. While we present our methodology as applied to a specific Inspection Rover case study, we believe that this combination provides benefits in maintaining coherent formal links across development and assurance processes for a wide range of autonomous robotic systems.

1 Introduction

The adoption of formal methods in industry has been slower than their development and adoption in research. One of the main pitfalls is the difficulty in integrating the results from formal methods with non-formal parts of the system development process. A central stumbling block is the formalisation of the (informal) natural language descriptions needed to perform the formal analysis, as well as the analysis and interpretation of the formal verification results.

The integrated formal methods approach relies on various tools cooperating to ease the burden of formal methods at various phases of system development. This often involves facilitating the use of one tool/formalism from within another (e.g. Event-B‖CSP [30]), the development of a tool/formalism that incorporates multiple others (e.g. Why3 [16], Circus [8]), or the construction of translation rules to systematically translate between tools/formalisms (e.g. EventB2JML [29]). Recent work has argued that, for autonomous robotic systems, the use of multiple formal and non-formal verification techniques is not only beneficial but actually necessary to ensure that such systems behave correctly [15, 23]. Central to this argument is the fact that the usually modular nature of robotic systems makes them more amenable to an integrated verification approach than

monolithic systems [7]. The inherent modularity in robotic systems usually stems from the use of a node-based middleware such as the Robot Operating System (ROS) [28]. However, similar middlewares such as NASA's core Flight System (cFS) [26] also support the development of similarly complex, modular systems.

In this paper, we study the support for integrating formal verification results at both system- and component-level in the design, implementation and assurance of a critical system, namely, an autonomous rover undertaking an inspection mission. In contrast to usual approaches to integrating formal methods, such as those described above, we use an assurance case as the point of integration rather than building bespoke tools or defining mathematical translations between specific formal methods. In this way, we harness the benefits of an integrated approach to verification without the usual overheads. Specifically, we use AdvoCATE [12] to perform safety engineering and assurance, FRET [18] to elicit and formalize requirements, and COCOSIM [5] with Kind2 to perform compositional verification of the system-level requirements. Further, we use Event-B [3] and Kind2 for the component-level formal verification. AdvoCATE facilitates the integration of the artifacts/evidence produced from these tools.

In summary, we contribute an inspection rover case study that demonstrates:
- how these tools can be linked via an argument in an assurance case.
- the benefit of using distinct tools due to limitations that each might have (e.g. Kind2 would timed on certain properties that were verified in Event-B).
- how developing with formal methods in mind from the outset can influence the design of the system, making it more amenable to formal verification.

2 Tool Support

In this section, we briefly discuss the tools that we used in the case study.

Assurance Case Automation Toolset (AdvoCATE) [11]: supports the development and management of assurance cases, which are composed of all of the assurance-related artifacts that are created during system development. To enable automation, AdvoCATE is built with a formal basis where all of the assurance artifacts can be defined and formally related. Some artifacts can be created directly in AdvoCATE (e.g., hazard log, bow tie diagrams), while others, such as formal verification results, can be imported. AdvoCATE uses Goal Structuring Notation (GSN) [2] to document assurance cases in the form of arguments.

Formal Requirements Elicitation Tool (FRET) [18]: is an open source framework [33] for the elicitation, formalization and understanding of requirements. Users enter system requirements in a structured natural language. FRET helps understanding and review of semantics by utilizing a variety of forms for each requirement: natural language description, formal logics, and informal diagrams. Requirements can be defined in a hierarchical fashion and can be exported in a variety of forms to be used by formal analysis tools such as COCOSIM [25].

Contract-based Compositional Verification of Simulink Models (CO-COSIM) [5]: is an open source framework [32] for Simulink/Stateflow formal verification. COCOSIM translates a Simulink model into Lustre code [20], which

can then be verified using the Kind2 model-checker [9]. COCOSIM annotates the model with assume-guarantee contracts. Verification can then be performed in a compositional way or by checking the contracts against component behavior.

Event-B [3]: is a formal method that is used predominantly in the verification of cyber-physical systems. Event-B uses a set-theoretic modelling notation and supports formal refinement. In Event-B, models are composed of machines, which model the dynamic components of a systems' specification, and contexts, which model the static components. Event-B has tool support via the Rodin Platform, an Eclipse-based IDE, which generates proof obligations corresponding to a given specification and provides support for automatic and interactive proof [4].

3 Assurance-based Formal Methods Integration

The objective of this work is *to study the integration of formal veri cation results via the development of an assurance case, as applied to a robotic system, using a tool palette that includes NASA Ames'* FRET, COCOSIM, *and AdvoCATE tools, as well as Event-B.* To this end, we provide a step-by-step methodology that builds on top of existing NASA guidelines [14, 21] that can be used in the design and development of mission-critical systems. In particular, existing guidelines [21] suggest the following phases: 1) characterization; 2) modeling; 3) specification; 4) analysis; and 5) documentation. Each phase consists of constituent processes and the overall process is iterative rather than sequential.

Our methodology is guided by the need to devise a detailed assurance case that integrates verification results from a number of distinct tools. The steps that we followed are the following:

Step 0: Characterize initial system.
Step 1: Create initial Simulink model.
Step 2: Perform preliminary hazard analysis.
Step 3: Define mitigations and safety requirements.
Step 4: Refine Simulink model according to mitigations.
Step 5: Formalize requirements and create formal specification(s).
Step 6: Perform verification and simulation at system- and component-levels.
Step 7: Document verification results and build safety assurance patterns.

Fig. 1 presents a detailed view of our methodology. The upper part of the figure shows the system-level concept, design, and assurance steps that are mainly performed by the AdvoCATE tool, while the lower part of the figure shows the formal methods application steps performed by the FRET, COCOSIM, and Event-B tools. In the analysis phase we perform two types of analysis. We use COCOSIM to perform *compositional system-level* analysis with the Kind2 analysis tool. We additionally perform verification at *component-level* against the system model using the Atelier-B and Kind-2 tools. Atelier-B is provided through Event-B. Finally, in the documentation phase, we use AdvoCATE to perform safety engineering and assurance by integrating the evidence that was produced by all tools.

Over the years, we have worked with a variety of formal approaches for the assurance of safety-critical systems. The goal of this study is to explore how

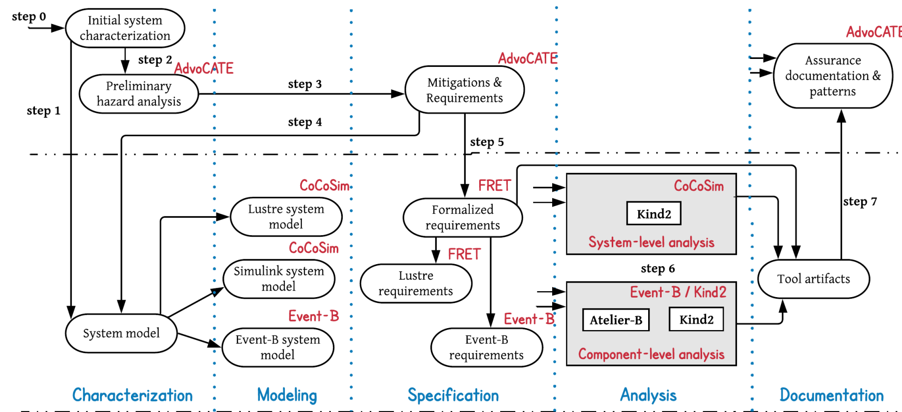

Fig. 1. Our methodology for integrating formal verification results via AdvoCATE.

such approaches can work together and be integrated within the development process of an autonomous system. To this aim, we developed a case study of a rover system. We target rovers for a variety of reasons. First, rovers are used in a large number of autonomous systems, and present challenges that are typical of autonomous applications. Second, some of the authors have prior experience with autonomous robotic systems that are deployed in hazardous environments, such as the nuclear, offshore, and space domains through their involvement in projects[4]. Third, our research group at NASA Ames is in the process of building some rover applications in order to experiment with AI technologies and their assurance techniques.

Our case study is not extracted from an actual mission. Rather, it is developed by iteratively using assurance approaches such as hazard analysis, requirements engineering, formal assurance cases, and model checking. The resulting Inspection case study has a reasonable complexity, and demonstrates a variety of challenges in formal methods techniques and their integration. Most importantly, we make the details of our case study publicly available [1], since we believe it can serve as a good basis for discussion and comparison of approaches and tools across the research community.

Four formal methods experts were involved: 1) a safety expert; 2) a requirements expert; 3) a Simulink and Lustre verification expert; and 4) a verification expert of robotic systems that also served as the domain expert. Step 0 was performed by the domain expert, step 1 was performed together by the Simulink and domain experts. Steps 2 and 3 were performed by the safety expert. Steps 4 and 6 were performed together by the safety and Simulink experts. Step 5 was performed together by the requirement and domain experts, and finally step 7 was performed mainly by the safety expert with contributions from all others.

[4] UKRI and EPSRC Hubs for “Robotics and AI in Hazardous Environments”.

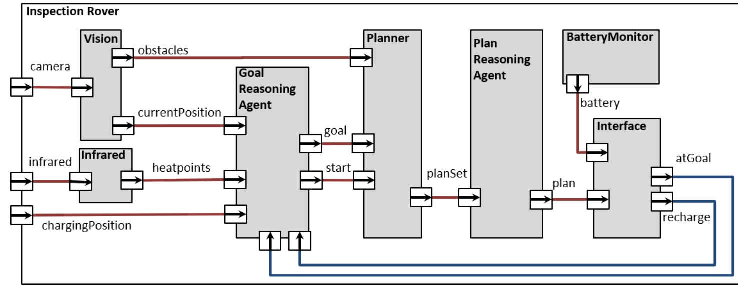

Fig. 2. Preliminary inspection rover system architecture.

4 The Case Study Step-by-Step

4.1 Step 0: Characterize Initial System

We performed our case study in the context of the navigation system for an autonomous rover undertaking an inspection mission. The objective of this rover is to explore a square grid of known size and to autonomously navigate to points of interest whilst avoiding obstacles and recharging as necessary. We assumed that this system would be operated indoors to minimize environmental uncertainty.

4.2 Step 1: Create Initial Simulink Model

Simulink is usually used by engineers for developing an executable specification of the whole system, which the user may use to simulate and validate before producing the source code. However, in this case study, we used Simulink in two different ways: 1) as an architecture description language, which allowed us to specify the architecture of the rover without providing implementation of low-level components (for compositionally verifying properties using assume-guarantee reasoning); 2) behavioral specification language for the implementation of some of the low-level components (for checking properties against component behavior).

In step 1, we created the initial Simulink model, which comprises a preliminary architecture of our rover as shown in Fig. 2. The rover's goal is to navigate to all heat positions on a 2D grid map of a given size. The *Vision* system is used by the rover to detect obstacles that it should avoid. The *Infrared* component identifies grid locations that are hotter than expected. Based on these heat locations, the autonomous *Goal Reasoning Agent* selects the hottest location as the goal, unless the *Battery Monitor* (via the *Interface*) indicates that it must recharge. The *Planner* returns a set of obstacle-free plans for navigating from the current position to the goal. The autonomous *Plan Reasoning Agent* selects the shortest plan. Finally, the *Interface* translates the navigation actions of the plan into the instructions required by the hardware components and alerts the *Goal Reasoning Agent* when it reaches the goal or that it does not have enough battery to execute the chosen plan so it must recharge.

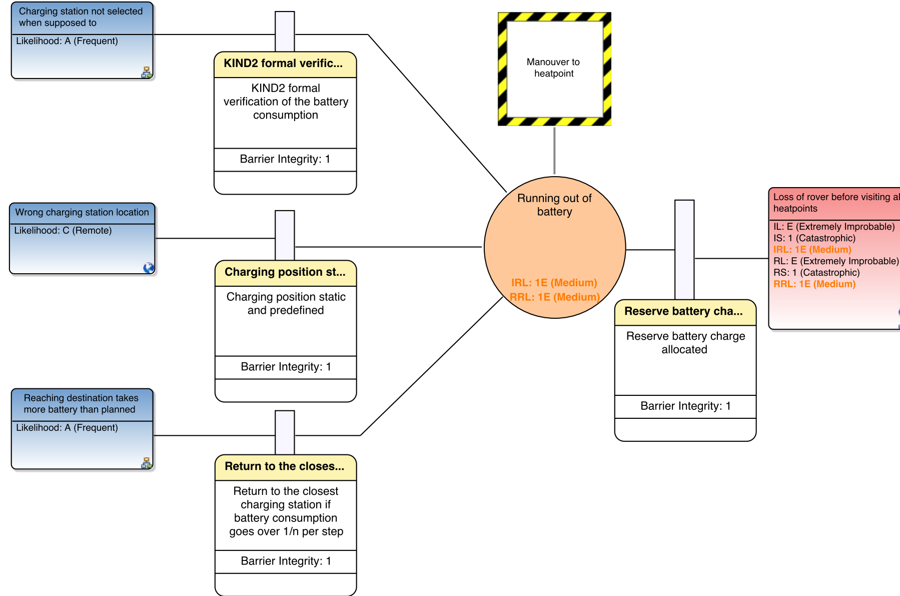

Fig. 3. Bow Tie Diagram corresponding to the *running out of battery* hazard.

4.3 Step 2: Perform Preliminary Hazard Analysis

To perform the preliminary hazard analysis in AdvoCATE, we defined a functional decomposition of the Inspection Rover based on Fig. 2 which was used as the basis for the traditional hazard analysis (FMEA [31]) that we documented in the AdvoCATE hazard log. In particular, we have identified the following two top-level hazards: 1) *loss of rover*, and 2) *inspection nished before visiting all of the heatpoints* In total, we identified 25 hazards including these two. For example, we identified the *running out of battery hazard* as a cause of *loss of rover*. Further causes of *running out of battery* are shown in Fig. 3.

4.4 Step 3: Define Mitigations and Safety Requirements

After identifying the hazards, including their causes and effects, we carried out a risk analysis where we qualitatively analysed the severity and likelihood of the identified hazards in order to estimate the risk level. Based on this, we defined mitigations to minimize the risk of those hazards and their consequences. For example, the *loss of rover* hazard is characterized with catastrophic severity, while its likelihood is calculated based on the events that cause it. The combination of the two defines the risk associated with the hazard.

Next, we identify and describe mitigations against these hazards. We performed mitigation planning using Bow tie Diagrams (BTD). Fig. 3 shows a BTD corresponding to the *running out of battery* hazard (orange circle). The BTD shows its causes (blue rectangles) and its consequence (red rectangle). For example, in order to minimize the risk of *running out of battery*: (1) we formally analysed the navigation system and battery controller, (2) we ensured that the

charging station position is predefined so that we can estimate at every point whether we have enough battery to go to recharge, and (3) if the basic assumptions about battery consumption are violated, then we abort and return to the charging station. Besides mitigating the causes to prevent the hazard from happening, we add the recovery barrier between the hazard and the consequence to reduce the severity of the consequence in case the hazard still occurs.

Based on the defined mitigations, we specified the corresponding mitigation requirements for each of the hazards. These are the parent (system-level) requirements that we have defined, and each was further decomposed into its child (component-level) requirements detailing the specific mitigation mechanisms.

[**R1:**] The rover shall not run out of battery.
[**R2:**] The rover shall not collide with an obstacle.
[**R3:**] The rover shall visit all reachable heat points.

4.5 Step 4: Refine Simulink Model According to Mitigations

Some of the identified mitigations required design modifications resulting in a refined system architecture (Fig. 4). The initial rover position and a static charging station position are given as user input. Note that the charging station position is static and the rover always starts its missions from a pre-defined initial position.

We modified the original architecture, adding the *MapValidator* component to check that the initial position, charging position, obstacles and heat points are valid, i.e., that the initial position and the charging station position are not obstacles or heat points, as well as that the obstacle and heat point sets are mutually exclusive. Furthermore, *MapValidator* checks that the initial position, as recognized by the *Vision* component, is in fact the same position as defined by the *initialPosition* pre-defined value.

Next, we defined the *NavigationSystem* which contains the *ReasoningAgent* and the *Battery_Interface* components. We emphasise these two components as we focus on formally verifying them. We further decompose these components.

The *ReasoningAgent* takes as input the identified and validated obstacle locations, the current rover position, the heat points and the position of the charger. It outputs three things: (1) a plan from the current position to the goal (*plan2D*), (2) a plan from the goal to the charger location (*plan2C*), and (3) the list of *visited* locations which it is responsible for keeping track of. Within the *ReasoningAgent*, the goal reasoning agent (*GRA*) chooses the goal as either the next hottest heat point or as the charging location if the *recharge* flag has been set to true by the *Battery_Interface*. The *GRA* updates the *visited* locations.

Further, the *ReasoningAgent* contains *ComputePlan2Charging* and *ComputePlan2Destination* which both contain a *Planner* and plan reasoning agent (*PRA*). Specifically, *ComputePlan2Charging* returns the shortest plan from the goal to the charger, whilst *ComputePlan2Destination* returns the shortest plan from the current position to the goal.

The *Battery_Interface* is composed of a *BatteryMonitor* and a hardware *Interface*. The *Interface* takes the plans from *NavigationSystem* and the battery status from the *BatteryMonitor* as inputs, and returns two flags indicating whether the rover has reached the goal (*atGoal*) and the status of the battery

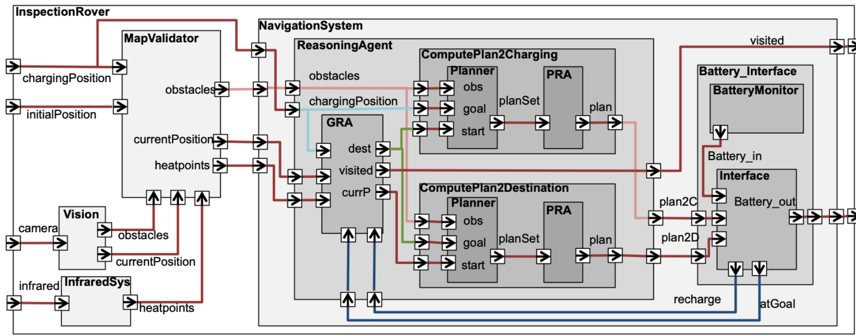

Fig. 4. Upgraded inspection rover architecture with additional components and data.

charge (*recharge*). The *recharge* output is set to true when the current battery charge is not sufficient to follow the plan to the goal (*plan2D*) and then return to the charging station (*plan2C*).

If the *recharge* flag is false, then the *Interface* executes the plan and returns *atGoal* as true once it reaches the goal. However, if the *recharge* flag is true and the rover first needs to recharge, then the *Interface* sets *atGoal* to false. We note that we actually do not need both of these outputs since we have always $recharge = not\ atGoal$. However, we include both for simplicity. These outputs are fed back to the *ReasoningAgent* that generates the next plan, and this loop executes until all of the heat points have been visited. Note that we assume that the frequency of execution for the *NavigationSystem* is similar to the *Interface*.

4.6 Step 5: Formalize Requirements and Create Formal Specifications

We formalized the natural language requirements in FRETISH, which is the restricted natural language provided by FRET which has a precise, unambiguous meaning. FRETISH requirements contain up to six fields: scope, condition, component*, shall*, timing, and response*, with mandatory fields indicated by '*'. Here, 'component' specifies the component that the requirement refers to. Next, 'shall' expresses that the component's behavior must conform to the requirement. The 'response' is a Boolean condition that the component's behavior must satisfy. 'scope' specifies intervals where the requirement is enforced. For example, 'scope' can specify system behavior *before* a mode occurs, or *after* a mode ends, or when the system is *in* a mode. The optional 'condition' defines a Boolean expression that triggers a 'response'. When triggered, the response must occur as specified by the timing, e.g., *immediately*, *always*, *after/for/within N time units*.

For each FRETISH requirement, FRET produces natural language and diagrammatic explanations of its exact meaning, and formalizes the requirement in temporal logics. The majority of the requirements that we formalized did not have scope or condition but they did have *always* timing, e.g.:

[R1]: Navigation shall always satisfy battery > 0.

Other requirements use the condition field and *immediately* timing, e.g.:

[R1.2]: if recharge GRA shall immediately satisfy goal $=$ chargePosition.

Notice that if recharge is a 'trigger': the requirement is only enforced when

the condition becomes true from false. The use of 'immediately' states that the response must hold simultaneously with each trigger point. The natural language version of **[R1]** was previously presented in Section 4.4, while the natural language version of **[R1.2]** is *"Charging station shall be selected as the next destination whenever the recharge flag is set to true"*.

Some requirements needed first-order temporal logic, which is not currently supported in FRET. For these, we used auxiliary variables that we instantiated with quantifiers at the Lustre level. For instance, the natural language version of requirement **[R3.3]** is *"The hottest heatpoint that was not visited before shall be the current goal when recharge flag is false."* was written in FRETISH as follows:

[R3.3]: GRA shall always satisfy if ! recharge then if forAll_i & i_inGrid then (if ! visited[i] then heatpoints[goal] >= heatpoints[i]))

In total, our case study contains 28 requirements, out of which 7 required first-order temporal logic formulas. We were able to formalize all of these in FRET. Although in most cases, the natural language requirements looked relatively straightforward, a closer study revealed many questions regarding their precise meaning. We performed several iterations in which we refined the FRETISH versions of the requirements. For this, the semantic explanations and simulation capabilities supported by FRET were very helpful.

FRETISH to Verification Code: FRET formalized the requirements in future-time and past-time Linear Temporal Logic. We used the pure past-time logic variant since Lustre-based verification tools only accept pure past-time temporal logic specifications. For instance, **[R1]** and **[R1.2]** were formalized as follows:

[R1]: $H(battery > 0)$;

[R1.2]: $H((recharge \& ((Y(!recharge)) | FTP)) \Rightarrow (goal = chargePosition))$, where H, Y are the *Historically* and *Yesterday* past-time LTL operators and FTP means First Time Point of execution (equivalent to $\neg$ Y TRUE). We generated Lustre-based verification contracts that can be directly fed into COCOSIM for verification with the Kind2 model checker using the CoCoGen tool, which is integrated in FRET [24]. For example, **[R1]** was generated in Lustre:

```
1 guarantee "R1.2" (battery > 0);
```

If requirements were based only on model inputs, e.g., **[R1.2]**, then CoCoGen generated assumptions (instead of guarantees):

```
1 assume "R1.2" ((recharge and ((pre (not recharge)) or FTP)) => (goal =
      chargePosition));
```

where `FTP = true -> false`. As mentioned earlier, some requirements used first-order operators such as **[R3.3]** which was generated as follows:

```
1 guarantee "R3.3" not recharge => (forall (i:int) (0 <= i and i < width) => (
      not visited[i] => heatpoints[goal] >= heatpoints[i] )));
```

Notice that the `forAll_i` placeholder was replaced by `forall (i:int)`, and `i_inGrid` was replaced by `(0 <= i and i < width)` during generation.

Additionally, we specified the requirements in Event-B. Although Event-B does not support temporal logic, we used the FRETISH requirements to guide our modelling in Event-B. Notably, even though there was no LTL support, the

FRETISH requirements were simple enough and more useful as a starting point for formalization in Event-B than the original natural language requirements. For example, the natural language requirement [**R3.4**] is *"The shortest path to the current goal shall be selected"*. The FRETISH version is as follows: [**R3.4**]: Planner shall always satisfy if (planningCompleted & returnPlan) then (if (forAll_x & x_inPlanSet) then (card(chosenPlan) <= card(x))) and the corresponding Event-B invariant was based on the FRETISH version:
inv3: $(planningCompleted = TRUE) \wedge (returnplan = TRUE) \Rightarrow (\forall x \cdot x \in PlanSet \Rightarrow card(chosenplan) \leq card(x))$
where `card()` is native to Event-B and defines the cardinality (size) of a set.

Similarly, [**R2.5:**] *The calculated path to destination shall not include a location with an obstacle* was defined in Event-B as follows:
inv3: $\forall p, x \cdot p \in PlanSet \wedge x \in p \Rightarrow x \notin Obs$
where every element of `PlanSet` is a set of grid locations.

4.7 Step 6: Perform Verification and Simulation at System- and Component-Levels

Compositional Verification in CoCoSim: Our objective was to attach the component-level child requirements to the relevant component(s) and then to compositionally verify, using COCOSIM, that the system-level parent requirements hold. There were some requirements that we could not formally verify. For example, a requirement stating that the current position as recognized by the rover is its current physical position, cannot be *formally* modelled or verified but rather needs to be physically tested.

Compositional verification in COCOSIM involves defining a top system node with associated system-level assumptions and guarantees. During verification, the model checker attempts to show that these system-level properties can be successfully derived the component-level contracts. Using compositional assume guarantee reasoning, we were able to verify system-level requirements [**R1**] and [**R3**], defined in Section 4.4, using COCOSIM, but not requirement [**R2**].

Compositional verification of [**R1**] was achieved quite quickly (<20 seconds), this was because the model checker actually only had to analyse two components: the *Interface* and the *BatteryMonitor* to verify [**R1**]. [**R3**] was more complex since it involved a loop between the *Interface* and the *ReasoningAgent*. This required Kind2 to carry out a lot of unrolling to adequately assess this property as well as dealing with more complex contracts which included quantifiers and arrays. As a result, we were only able to prove [**R3**] for specific grid widths (3×3 took a few minutes, 4×4 took a few hours, and larger grids timed out).

Component-Level Verification Using Kind2 and Event-B: Previously, we used compositional verification to verify that the system-level parent requirements hold based on the component-level requirements. Here, our objective was to verify that the more detailed specification/implementation of individual components obey the associated component-level requirements. Recognising that, for autonomous robotic systems, it is often necessary to use a range of verification techniques for individual components, we used two distinct formal

methods in this section [15, 23]. Specifically, we used Kind2 to verify a simple implementation of the *GRA* and, Event-B to model and verify the *ComputePlan* component.

Specification and Verification of the GRA: We constructed a simple Lustre implementation of the *GRA* that we verified using Kind2. The full implementation can be found in [1]. The *GRA* computes the `start`, `goal` and the `visited` cells. The `start` is initialized as the `currentPosition`, if the goal was reached during the last execution (`atGoal` is true) then the start is the previous goal (`pre_goal`), if the `recharge` flag is active then the `start` is the previous start position since the rover did not move. The `goal` is set to `chargingPosition` if the `recharge` flag is active. Otherwise we choose the hottest heat point, the latter is computed with the help of the `hottestPoint` local array that keeps track of the hottest heat point traversing all `heatpoints`. We used Kind2 to verify all of the properties specified in the specification. We were able to verify most properties in less than 1 second. Due to state space explosion, there were some properties, e.g., requirement [**R3.3**] that were only provable for specific grid sizes. For example, we were able to verify [**R3.3**] for a grid size up to 4x4.

Specification and Verification of ComputePlan using Event-B: Our Event-B model contains three contexts (for modelling the static parts of the system) and two machines (for modelling the dynamic parts of the system). Event-B supports formal refinement and thus our contexts extend one another and our machines indicate refinement steps. Our most primitive context, `ctx0`, contains basic information about the size of the grid, valid grid locations, obstacles and heat points. We do not explicitly list the elements of these sets since this specification is for a generic planner. This is extended via `ctx1` which specifies functions that capture the behavior of the planning component.

The abstract machine, `mac0`, models a simple search-based planning algorithm which produces a set of plans containing the start and goal. Event-B uses sets as primitive so we ensure that these plans, represented as sets, can be linearized using the `adjacent` function specified in `ctx1`. The refinement, `mac1`, has the added functionalilty of a plan reasoning agent and chooses the shortest plan from `PlanSet`. We constructed another context, `ctx2`, to contain a constant which limits the number of plans that are generated.

We encoded [**R2.1**], [**R2.4.1**], [**R2.4.3**], [**R2.4.4**], [**R2.5**] and [**R3.4**] in our Event-B model. Although we could not verify [**R2**] compositionally, its child requirements feature in our Event-B model (e.g. [**R.2.5**] above). This ensures that the planning components do not inadvertently cause the rover to collide with an obstacle. The vast majority of the Event-B proof obligations were discharged automatically by the Atelier B theorem prover in the Rodin Platform. Some required interactive proof but they were relatively straightforward.

We used Event-B because it is not limited by the state space explosion problem that was causing Kind2 to time out. We specified more complex component-level properties that would have been difficult to verify using a model-checker. The full Event-B specification is available in [1].

4.8 Step 7: Document Verification Results and Build Safety Assurance Patterns

All of the verification results produced by the tools are a part of the safety case that was constructed in AdvoCATE. Some artifacts were imported automatically into AdvoCATE, while others were added manually. In this way, we ensured that the necessary data was available in AdvoCATE. Since this case study did not include a full system implementation, the safety case that we report here is an interim version and contains the current safety assurance status.

Fig. 5 presents an argument fragment about mitigation of the *running out of battery* hazard that causes *loss of rover*. Similar arguments exist for other causes of *loss of rover* and the other hazards. For brevity, Fig. 5 only contains a fragment of the existing argument. This argument focuses on two aspects: the requirements directly related to this hazard (right branch), and the causes that lead to the hazard (left branch). We indicate goals that have been further developed elsewhere with a green rectangle, while those that are currently undeveloped are indicated with a green diamond. More details about the safety case can be found in [1].

The goal *G14* focuses on **[R1]** that was verified using COCOSIM. We have created a similar argument for each of the system-level requirements that were formalized and verified compositionally with COCOSIM. For each, we argued specific goals (shown at the base of Fig. 5):

- The formalisation of the natural language requirement is correct (*G3-A1*).
- The results from COCOSIM are trustworthy (the goal *G4-A1*).
- The different design representations are consistent (*G5-A1*).
- The COCOSIM verification result for **[R1]** is valid (*G6-A1*).

To ensure that the different design representations were consistent across the tools, we performed manual reviews where automated consistency validation was not available. For example, we carried out manual reviews to verify that the design as specified in AdvoCATE was consistent with the design modelled in Simulink as well as those used by Kind2 and Event-B.

The goal *G3-A1* relates to the correct specification of **[R1]** in FRETISH and the correct functioning of FRET. While we have to verify through a manual review that the natural language requirement is correctly represented in FRETISH, the correct FRET functioning and generation of the corresponding COCOSIM contracts is supported by the automated verification framework of FRET.

The goal *G6-A1* presents the rationale behind the COCOSIM results that **[R1]** is valid in the design. This part of the argument points out the dependencies to the properties of the other components, but also implicit assumptions on which these results rely. Finally, to have confidence in the results from COCOSIM, we argued the trustworthiness of COCOSIM in the goal *G4-A1*. Since COCOSIM relies on some model transformations and external tools for verification, the correctness of these has to be established. For example, we argued the correctness of the translation from Simulink to Lustre code that is used by Kind2.

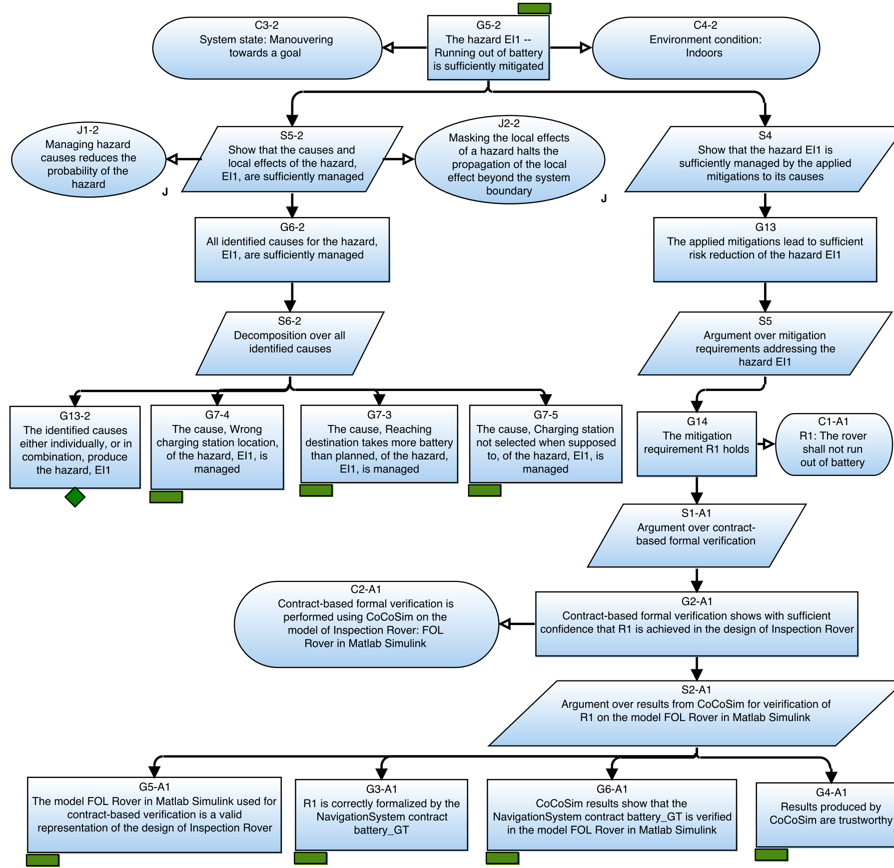

Fig. 5. The upper part of the argument assuring the *running out of battery* hazard (rectangles represent goals, parallelograms represent strategies, ovals with a 'J' represent justifications and rounded rectangles represent context statements).

5 Discussion

Using the given formal verification tools we were able to fully verify that the Navigation System will not cause the rover to run out of battery. We could not verify that a collision will never occur at system-level with COCOSIM due to the specification complexity. However, we were able to verify with Event-B that the Navigation System will not generate plans that contain obstacles at component-level. Finally, we were able to verify that the rover will visit all of the heat points with COCOSIM, but only for a small grid size of 4x4. Verifying the property for greater grid sizes did not finish even after several days of analysis.

This case study showed us that by following our methodology we were able to leverage multiple formal tools and use them in a *complementary* fashion. In this way, we applied formal methods to small, manageable chunks of the system to ease the verification burden and to avoid becoming trapped by the limitations of any single tool. Using FRET to bridge the gap between the informal and formal

steps by formalizing our requirements was particularly useful because it helped us to clarify any details that were implicit in the natural language requirements.

The explanations produced by FRET were instrumental in ensuring that the FRETISH requirements captured our intended semantics. Even though FRETISH is intended to be intuitive, translating the natural language requirements into FRETISH was not always straightforward. However, most of the Inspection Rover requirements fall within a small number of patterns, an issue that we have observed in other studies within our organization. We were not able to encode everything in FRETISH. For example, some requirements had to be defined using first order temporal logic. We were able to get around this problem by using auxiliary variables as placeholders for the quantifiers at the requirements level, but a FRETISH-level solution would be desirable.

The choice of COCOSIM and in particular Kind2 greatly influenced our design decisions. For example, in our original design, we represented cells in the grid as (x, y)-coordinates. However, we subsequently simplified this by using indices so that they were easier to represent and reason about in formal tools. Our choice of a compositional verification approach caused us to output specific variables such as the remaining battery power to verify **[R1]** compositionally. Furthermore, we had to adapt the hierarchical structure of the system to accommodate compositional verification. If the choice of formal verification tools is made early on in the system development process, the design of the system can be created so that it is more suitable to formal verification using the chosen method(s).

Not all of the formalized requirements were formally verifiable, some described hardware constraints and some required physical testing. This supports the claim that the robotics domain requires both formal and informal verification processes [15]. For example, everything depends on the accuracy of the current position of the rover which is a property that we could not formally verify in this case. However, by formalizing the requirements that will be verified via testing, we can potentially incorporate run-time analysis. Specifically, the formalized properties can be used to generate formal run-time monitors, which will help with fault management during operation. These could potentially help to create recovery barriers in the bow-tie diagrams. In this way, we could include the development of fault management at design time.

Integrating the verification results from the different formal methods in an assurance case required intensive cooperation between the assurance and formal methods experts. The effort required identifying dependencies between different tools, understanding the techniques and the tool implementations, implicit assumptions on which analyses were ran and results interpretated. The activity was greatly performed ad hoc. A more systematic approach to gathering the assurance information from formal methods applications would be beneficial.

The case study helped us to identify limitations in the used tools (AdvoCATE, FRET, COCOSIM and Event-B) for robotics applications. In fact, this case study prompted an update to COCOSIM to incorporate abstract unimplemented components. In particular, COCOSIM now generates Lustre code for these components using the `imported` keyword when no implementation is available. Other limitations that were identified include the lack of support in FRET for abstract

data types which caused us to manually edit the FRET-generated COCOSIM contracts. We encountered some difficulties when attempting to automatically import verification artifacts directly from the tools into AdvoCATE which caused us to insert some information manually.

6 Related Work

Heterogeneous verification techniques were used to verify an autonomous Mars Curiosity rover simulation [7]. This work uses distinct verification methods for specific components but does explicitly link the verification artifacts produced. Recent work proposes first-order logic to unify heterogeneous formal methods via a compositional approach but this work currently lacks tool support [6].

Other approaches to compositional verification include AGREE [27] and OCRA [10]. We explored these as potential alternatives to COCOSIM in this work but neither offered the level of expressivity that we sought. They also did not accommodate for the use of distinct verification techniques at component-level.

Developers should choose the most appropriate formal method on a per-component basis based on the suitability of the formal method and the user's level of expertise. As such, there are many alternatives to Event-B and Kind2, including Gwendolen [13], TLA+ [34] and Dafny [22].

Isabelle/SACM [17, 19] extends the Isabelle proof assistant to support assurance case development. In Isabelle/SACM, a UTP semantics must be defined for each formal verification artifact that is to be included in the assurance case.

7 Conclusions and Future Work

This paper presented our methodology for integrating results from distinct formal methods via the development of an assurance case. We applied this methodology in the design of an Inspection Rover system and used the AdvoCATE, FRET COCOSIM and Event-B tools. This is the first effort to integrate the four aforementioned tools. We illustrated how the choice of verification method can impact the system design and discussed how a heterogeneous set of verification results can be linked during assurance with AdvoCATE. Further, we have made our case study artifacts publicly available to fuel discussion in the research community.

This work has opened up a number of avenues for future research. In particular, it is likely that requirements for complex robotic systems will increasingly contain probabilistic properties. Therefore, supporting the definition of these requirements in FRET is a future direction. Additionally, we intend to develop a DSL to facilitate the integration of AdvoCATE with different verification tools. In this case study, we mainly used tools that output artifacts in JSON or XML format, which are straightforward to parse in AdvoCATE, but this may not be the case for other analysis tools. Finally, we intend to explore the definition of a 'Taxonomy of Requirements' and classify those in this case study. This will help developers to design their system with verification in mind by demonstrating how to classify requirements based on the ways that they will be verified and argued in an assurance case early at design phase.

References

1. https://drive.google.com/drive/folders/1SMekVBudCcQE5vNQE58W_W3DqgDHKhqS?usp=sharing.
2. GSN Community Standard Version 2. Technical report, Assurance Case Working Group of The Safety-Critical Systems Club, Jan. 2018.
3. J.-R. Abrial. *Modeling in Event-B: system and software engineering*. Cambridge University Press, 2010.
4. J.-R. Abrial, M. Butler, S. Hallerstede, T. S. Hoang, F. Mehta, and L. Voisin. Rodin: an open toolset for modelling and reasoning in Event-B. *International Journal on Software Tools for Technology Transfer*, 12(6):447–466, 2010.
5. H. Bourbouh, P.-L. Garoche, T. Loquen, É. Noulard, and C. Pagetti. Cocosim, a code generation framework for control/command applications an overview of cocosim for multi-periodic discrete simulink models. In *European Congress on Embedded Real Time Software and Systems*, 2020.
6. R. C. Cardoso, L. A. Dennis, M. Farrell, M. Fisher, and M. Luckcuck. Towards compositional verification for modular robotic systems. In *Workshop on Formal Methods for Autonomous Systems*, pages 15–22. Electronic Proceedings in Theoretical Computer Science, 2020.
7. R. C. Cardoso, M. Farrell, M. Luckcuck, A. Ferrando, and M. Fisher. Heterogeneous verification of an autonomous curiosity rover. In *NASA Formal Methods Symposium*, pages 353–360. Springer, 2020.
8. A. Cavalcanti and P. Clayton. Verification of control systems using circus. In *International Conference on Engineering of Complex Computer Systems*, page 10. IEEE, 2006.
9. A. Champion, A. Mebsout, C. Sticksel, and C. Tinelli. The kind 2 model checker. In *International Conference on Computer Aided Veri cation*, pages 510–517. Springer, 2016.
10. A. Cimatti, M. Dorigatti, and S. Tonetta. Ocra: A tool for checking the refinement of temporal contracts. In *International Conference on Automated Software Engineering*, pages 702–705. IEEE, 2013.
11. E. Denney and G. Pai. Tool support for assurance case development. *Automated Software Engineering*, 25(3):435–499, 2018.
12. E. Denney, G. Pai, and J. Pohl. Advocate: An assurance case automation toolset. In *International Conference on Computer Safety, Reliability, and Security*, pages 8–21. Springer, 2012.
13. L. A. Dennis and B. Farwer. Gwendolen: A BDI Language for Verifiable Agents. In *Workshop on Logic and the Simulation of Interaction and Reasoning*, pages 16–23. AISB, 2008.
14. H. Dezfuli, A. Benjamin, C. Everett, M. Feather, P. Rutledge, D. Sen, and R. Youngblood. Nasa system safety handbook. volume 2: System safety concepts, guidelines, and implementation examples. 2015.
15. M. Farrell, M. Luckcuck, and M. Fisher. Robotics and Integrated Formal Methods: Necessity meets Opportunity. In *International Conference on Integrated Formal Methods*, pages 161–171. Springer, 2018.
16. J.-C. Filliâtre and A. Paskevich. Why3—where programs meet provers. In *European symposium on programming*, pages 125–128. Springer, 2013.
17. S. D. Foster, Y. Nemouchi, C. O'Halloran, N. Tudor, and K. Stephenson. Formal model-based assurance cases in isabelle/sacm: An autonomous underwater vehicle case study. In *Formal Methods in Software Engineering*. ACM, 2020.

18. D. Giannakopoulou, A. Mavridou, J. Rhein, T. Pressburger, J. Schumann, and N. Shi. Formal requirements elicitation with fret. In *Requirements Engineering: Foundation for Software Quality*, 2020.
19. M. Gleirscher, S. Foster, and Y. Nemouchi. Evolution of formal model-based assurance cases for autonomous robots. In *International Conference on Software Engineering and Formal Methods*, pages 87–104. Springer, 2019.
20. N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud. The synchronous data flow programming language lustre. *Proceedings of the IEEE*, 79(9):1305–1320, 1991.
21. J. C. Kelly. Formal methods specification and analysis guidebook for the verification of software and computer systems volume ii: A practitioner's companion. 1997.
22. K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *International Conference on Logic for Programming Arti cial Intelligence and Reasoning*, pages 348–370. Springer, 2010.
23. M. Luckcuck, M. Farrell, L. A. Dennis, C. Dixon, and M. Fisher. Formal Specification and Verification of Autonomous Robotic Systems: A Survey. *ACM Computing Surveys (CSUR)*, 52(5):1–41, 2019.
24. A. Mavridou, H. Bourbouh, P.-L. Garoche, D. Giannakopoulou, T. Pressburger, and J. Schumann. Bridging the gap between requirements and Simulink model analysis. In *26th Intl Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ-2020, Poster)*, 2020.
25. A. Mavridou, H. Bourbouh, P.-L. Garoche, and M. Hejase. Evaluation of the fret and cocosim tools on the ten lockheed martin cyber-physical challenge problems. Technical report, TM-2019-220374, NASA, 2019.
26. D. McComas. Nasa/gsfc's flight software core flight system. In *Flight Software Workshop*, volume 11, 2012.
27. A. Murugesan, M. W. Whalen, S. Rayadurgam, and M. P. Heimdahl. Compositional verification of a medical device system. In *SIGAda annual conference on High integrity language technology*, pages 51–64, 2013.
28. M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5, 2009.
29. V. Rivera and N. Cataño. Translating event-b to jml-specified java programs. In *ACM Symposium on Applied Computing*, pages 1264–1271, 2014.
30. S. Schneider, H. Treharne, and H. Wehrheim. A csp approach to control in event-b. In *International Conference on Integrated Formal Methods*, pages 260–274. Springer, 2010.
31. D. H. Stamatis. *Failure mode and e ect analysis: FMEA from theory to execution*. Quality Press, 2003.
32. C. D. Team. CoCoSim – automated analysis framework for simulink. https://github.com/NASA-SW-VnV/CoCoSim.
33. F. D. Team. FRET – formal requirements elicitation tool. https://github.com/NASA-SW-VnV/FRET.
34. Y. Yu, P. Manolios, and L. Lamport. Model checking tla+ specifications. In *Advanced Research Working Conference on Correct Hardware Design and Veri cation Methods*, pages 54–66. Springer, 1999.