

Automated SysML v2 System Model to Memory-Safe Language Code Generation with Integrated AI Assistance

DARPA PROVERS



Collins Aerospace
An **RTX** Business

David S. Hardin, Ph.D.
Applied Research & Technology

Collaborators on this Presentation

- *Collins Aerospace*: Isaac Amundson, Junaid Babar, Darren Cofer, Saqib Hasan, Karl Hoech, Amer Tahat
- *Kansas State University*: Jason Belt, John Hatcliff, Robby
- *Aarhus University (Denmark)*: Stefan Hallerstedte



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

DARPA PROVERS

Pipelined Reasoning of Verifiers Enabling Robust Systems

Develop automated, scalable **formal methods** tools that are integrated into traditional development pipelines using “proof engineering” techniques

Enable traditional product engineers to incrementally produce and maintain **high-assurance** national security systems

So How Did We Get Here?

DARPA High Assurance Cyber Military Systems (HACMS)

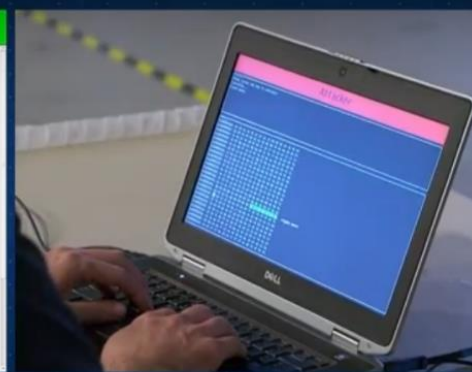
DEF CON 29



We brought a hackable quadcopter with defenses built on our HACMS program to [@defcon](#) [#AerospaceVillage](#). As program manager [@raymondrichards](#) reports, many attempts to breakthrough were made but none were successful. Formal methods FTW!



10:20 AM · Aug 9, 2021 · Hootsuite Inc.



9:59 / 25:07

Loonwerks.com/projects/hacms

2:19 / 3:43

Scroll for details



Collins Aerospace

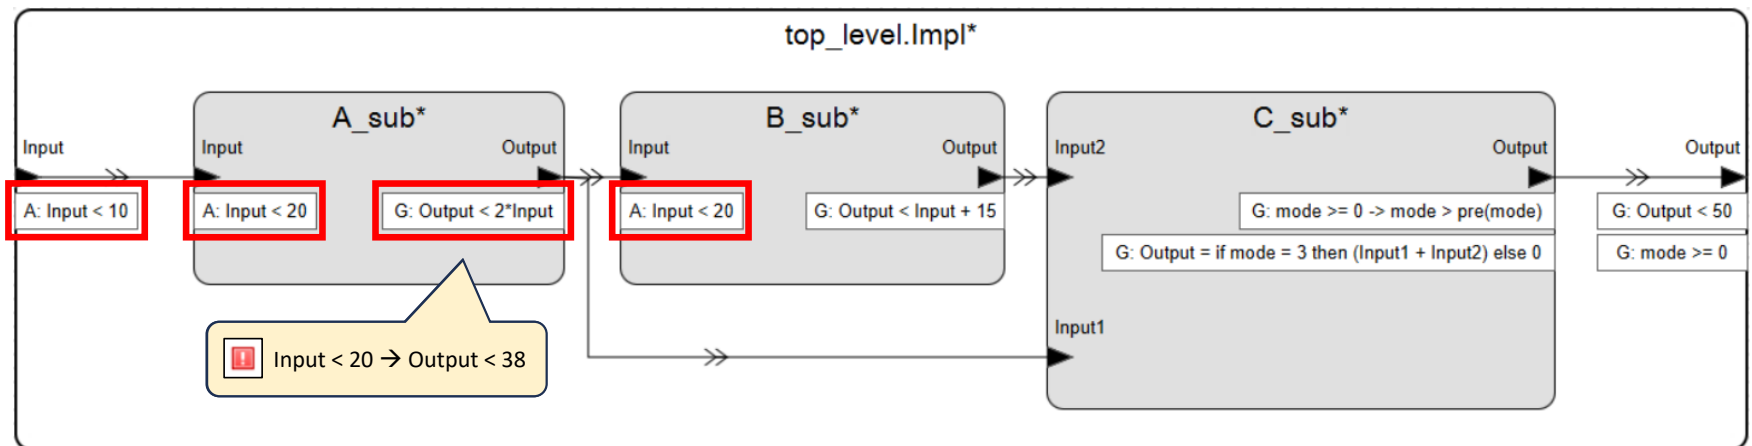
An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

Compositional Reasoning for Model-Based Systems Engineering

Assume Guarantee REasoning Environment (AGREE)

- Assume-Guarantee annex for AADL architecture models
 - Assumptions describe the expectations that a component has on the environment
 - Guarantees describe bounds on the behavior of the component when assumptions are valid
- Compositional analysis to prove correctness of:
 - Component interfaces (component assumptions are satisfied by upstream guarantees)
 - Component implementations (component assumptions and subcomponent guarantees satisfy guarantees)



Resolute

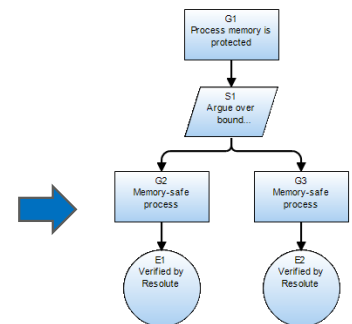
An Assurance Pattern Language and Evaluation Tool for Architecture Models

- The **structure of the system architecture** dictates the structure of the **assurance case**
 - Design patterns → Assurance patterns
- Extension of AADL language
 - Assurance case **instantiated** with elements from AADL model
- Specify **logical rules** for **evaluating evidence**
 - Automated evaluation

```
goal memory_protection(p : process) <=

  ** "Process " p " memory is protected from alterations by other processes" **

  strategy "Argue over bound processes";
  property(p, OS) = "seL4" or
    forall (mem : memory) . bound(p, mem) =>
      forall (q : process) . bound(q, mem) => memory_safe_process(q)
```



AdvoCATE (NASA)

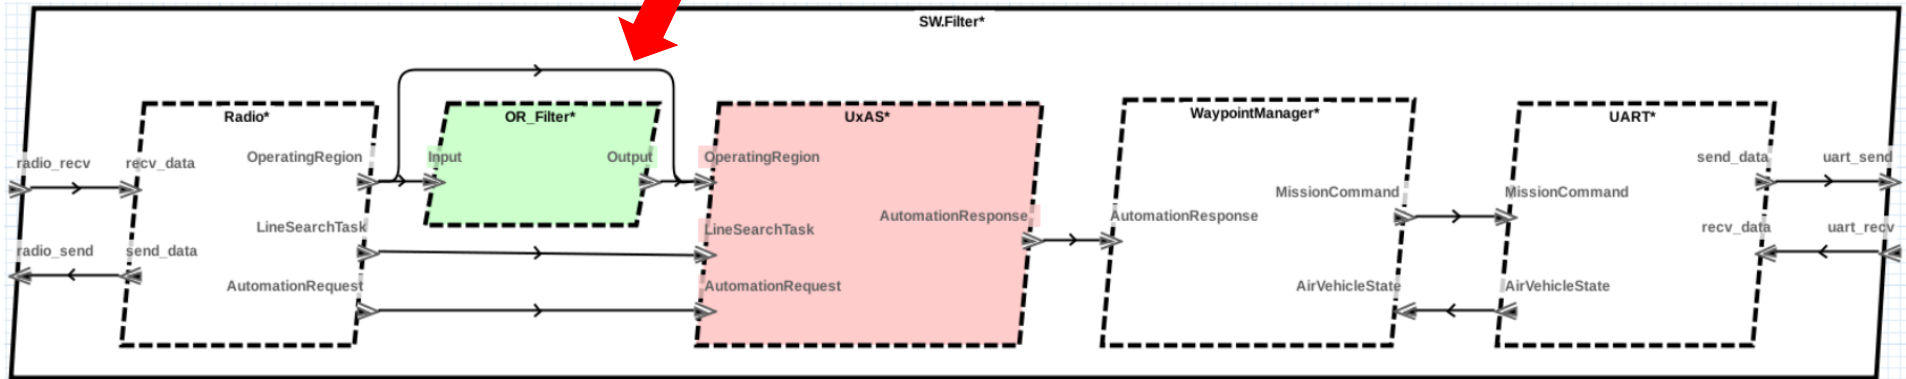


Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

Oops!



Problems Properties Console Assurance Case AGREE Results

Req_WellFormed_OperatingRegion(MissionComputer_Impl_Instance : MC::MissionComputer.Impl, "Req_Wellformed_OperatingRegion", MissionComputer_Impl_Instance : MC::MissionComputer.Impl)

- ❗ [permit_well_formed_OR_data] UxAS component shall only receive well-formed messages
 - ❗ AGREE properties passed
 - ❗ AGREE analysis was run
 - ❗ Filter OR_Filter : SW::OperatingRegion_Filter.Impl is properly added to component UxAS : Waterways_UxAS::Waterways.i
 - ✅ A CASE_Filter OR_Filter : SW::OperatingRegion_Filter.Impl is connected to component UxAS : Waterways_UxAS::Waterways.i by connection SW::SW.Filter.filter_out_connection
 - ❗ Component OR_Filter : SW::OperatingRegion_Filter.Impl cannot be bypassed
 - ✅ Component property implemented
 - ✅ OR_Filter : SW::OperatingRegion_Filter.Impl implementation is appropriate for OS
 - ✅ Component proof checked



Collins Aerospace

An RTX Business

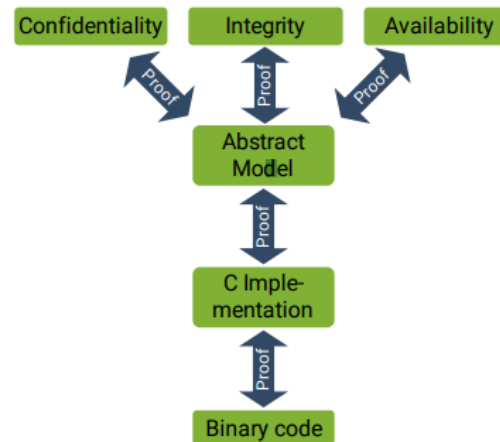
© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

seL4

Formally Verified Microkernel

- seL4 microkernel guarantees partitioning of components and communication, backed by computer-checked proofs
- seL4 guarantees no infiltration, exfiltration, eavesdropping, interference, and provides fault containment for untrusted code
- Ensures soundness of the MBSE design process – components can be analyzed separately and composed safely

University of New South Wales
Proofcraft

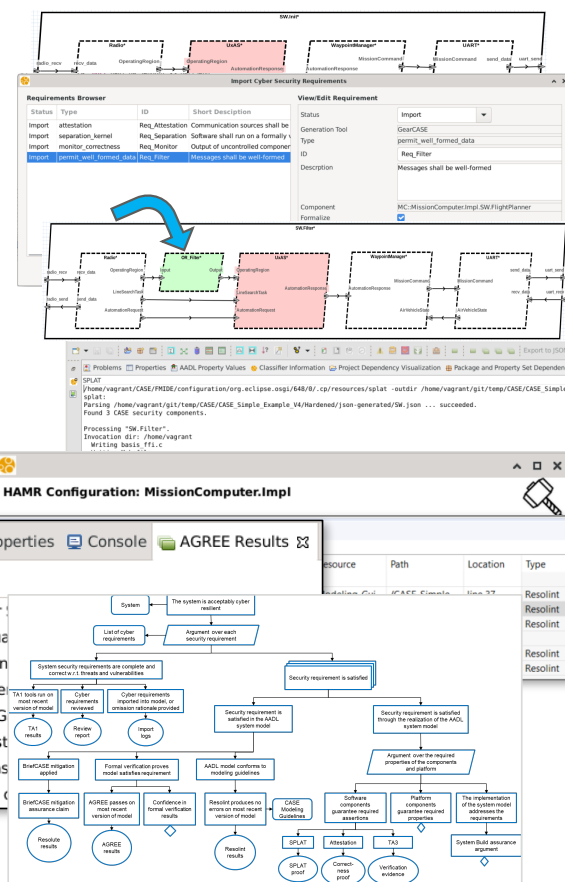


seL4 is...

- An operating system microkernel
- A hypervisor
- Proved correct
- Provably secure
- Fast

DARPA Cyber Assured Systems Engineering (CASE)

- Objective: develop the necessary design, analysis and verification tools that enable engineers to build cyber-resilient systems, including legacy elements
- BriefCASE**
 - Integrated **model-based systems engineering** tool suite based on AADL models
 - Analyze architecture models for cyber vulnerabilities and generate **cyber resiliency requirements**
 - Transform system architecture** models to satisfy cyber-resiliency requirements
 - Synthesize **high-assurance component implementations** from formal specifications
 - Generate **software integration code** directly from verified architecture models
 - Build to a **formally verified secure microkernel** target (seL4)
 - Assurance:**
 - Check **model conformance** to standards
 - Verify system design and implementation using **formal methods**
 - Document proof of correctness with an **assurance case**



Resolint

A linter tool for AADL models

- Define *rules* in Resolute that correspond to modeling guidelines
- Group rules into *rulesets* corresponding to organizational process, customer requirements, certification guidelines, and tool constraints
- Automatically check compliance with modeling guidelines in OSATE

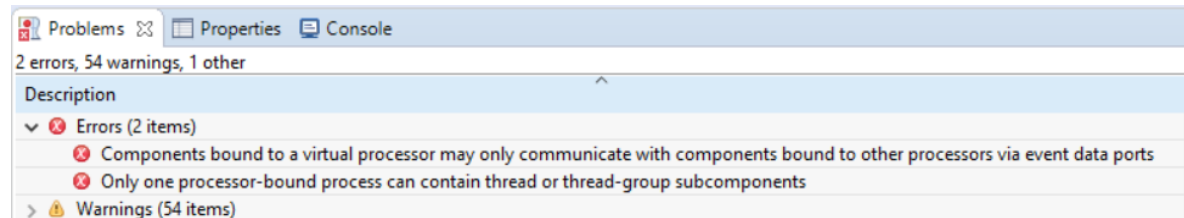
Modeling Guidelines

rule dispatch_protocol_specified	Threads should have the dispatch_protocol property specified	⚠
rule valid_dispatch_protocol	Threads can only specify a dispatch_protocol property of periodic or sporadic	❌

```
dispatch_protocol_specified() <=
  ** "Threads should have the Dispatch_Protocol property specified" **
  forall (t : thread) .
    lint_check(t, has_property(t, Dispatch_Protocol))

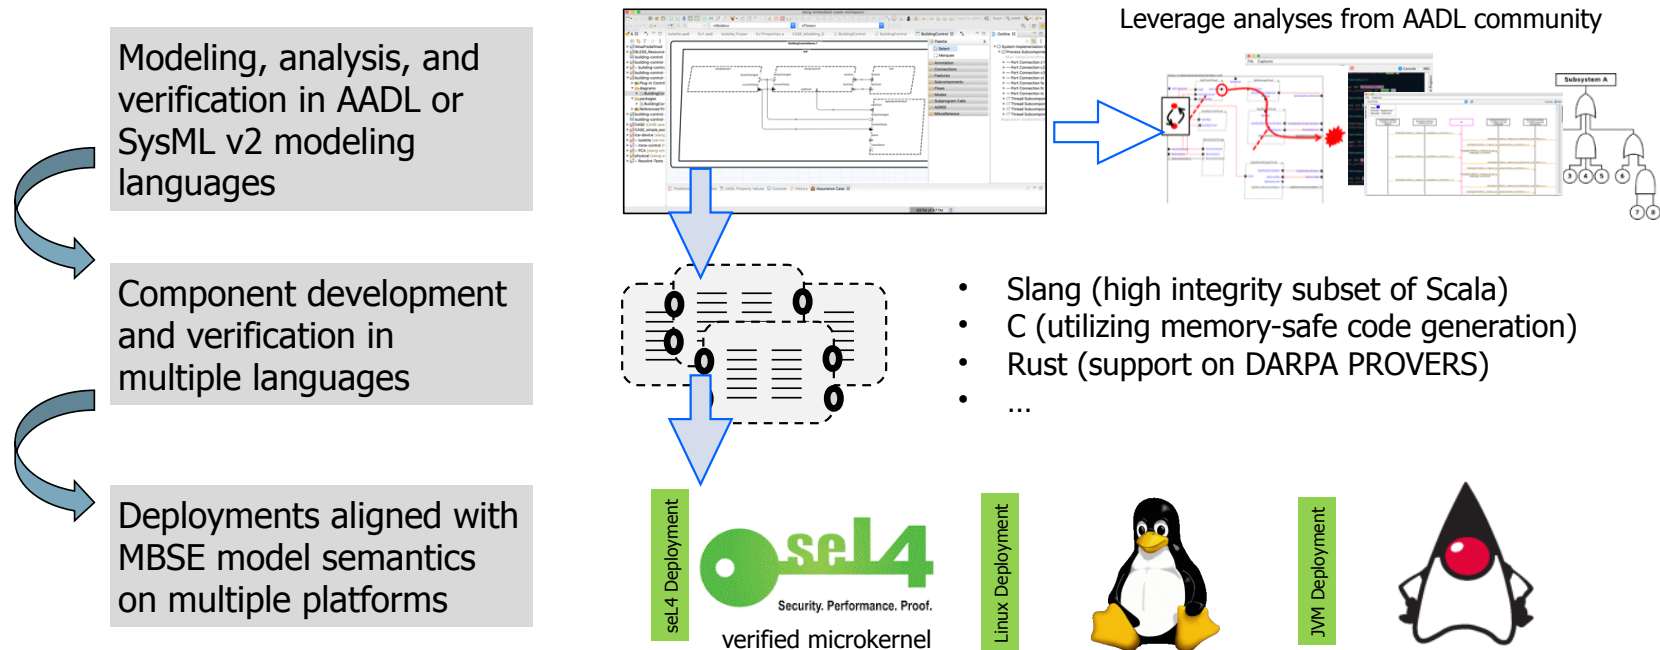
valid_dispatch_protocol() <=
  ** "Threads can only specify a dispatch_protocol property of periodic or sporadic" **
  forall (t : thread) . lint_check(t, has_property(t, Dispatch_Protocol) =>
    (property(t, Dispatch_Protocol) = "Sporadic" or property(t, Dispatch_Protocol) = "Periodic"))
```

```
1 package System_Build
2
3@public
4
5@
6  annex resolute {**
7
8    ruleset HAMR {
9      info (print("Linting HAMR ruleset"))
10
11      error (one_process())
12
13      warning (dispatch_protocol_specified())
14      error (valid_dispatch_protocol())
15
16      error (bounded_integers())
17      error (bounded_floats())
18
19      warning (data_type_specified())
20      warning (subcomponent_type_specified())
21
22      error (array_dimension())
23      error (one_dimensional_arrays())
24    }
25 }
```



HAMR

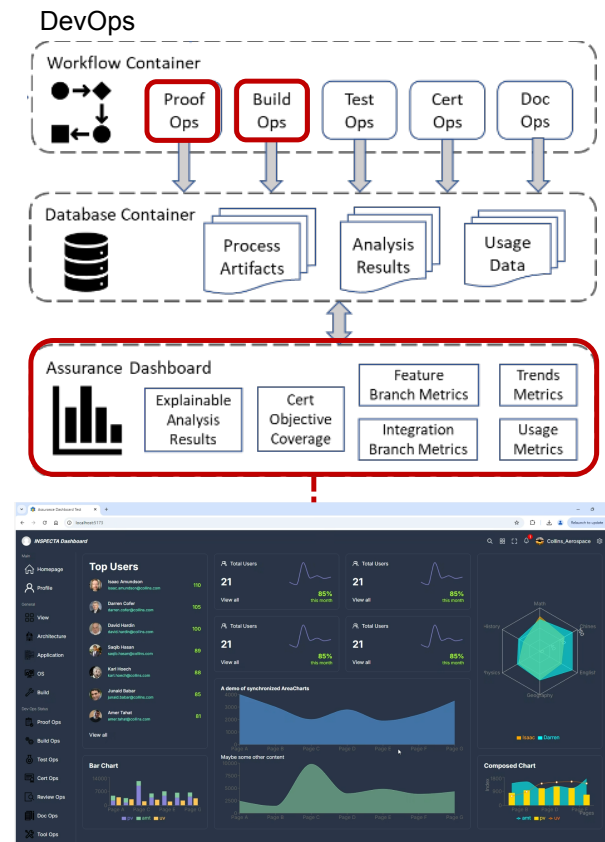
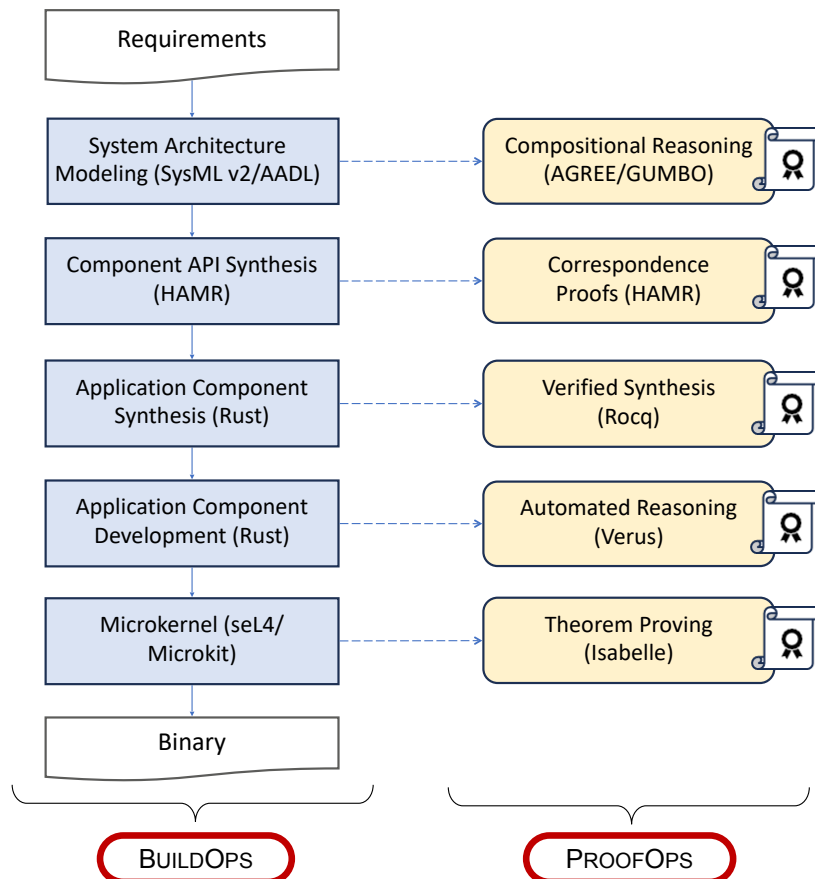
High-Assurance Modeling and Rapid Engineering for Embedded Systems



HAMR: infrastructure code generation and target platform build tool

DARPA PROVERS: INSPECTA Team

Industrial-Scale Proof Engineering for Critical Trustworthy Applications



INSPECTA Team

- Collins Aerospace, Team Lead
 - *Darren Cofer, PI*
- Carnegie Mellon University
- Dornierworks
- Kansas State University
 - with Aarhus University
- Proofcraft
- University of Kansas
- University of New South Wales

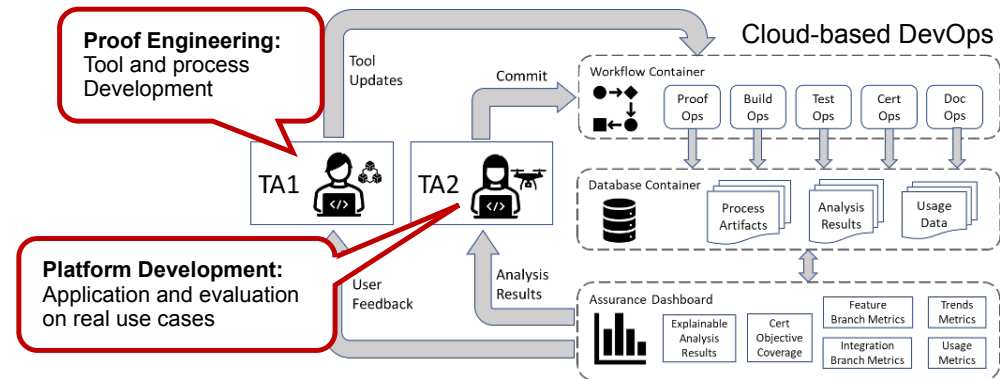
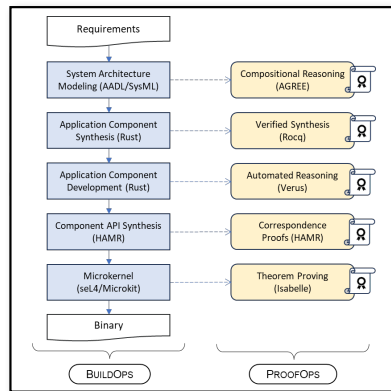


Collins Aerospace
An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

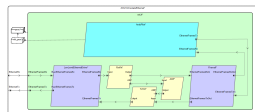
Technical Areas

TA1: Proof Engineering

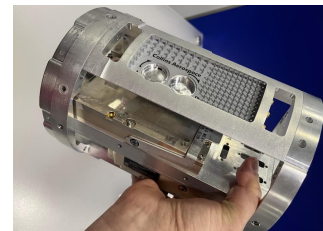


TA2: Platform Development

- Open Platform
 - Developed and supported by DornerWorks
 - Unrestricted UAV mission software, system model with formal properties, multiple VMs, Rust software components, seL4 kernel

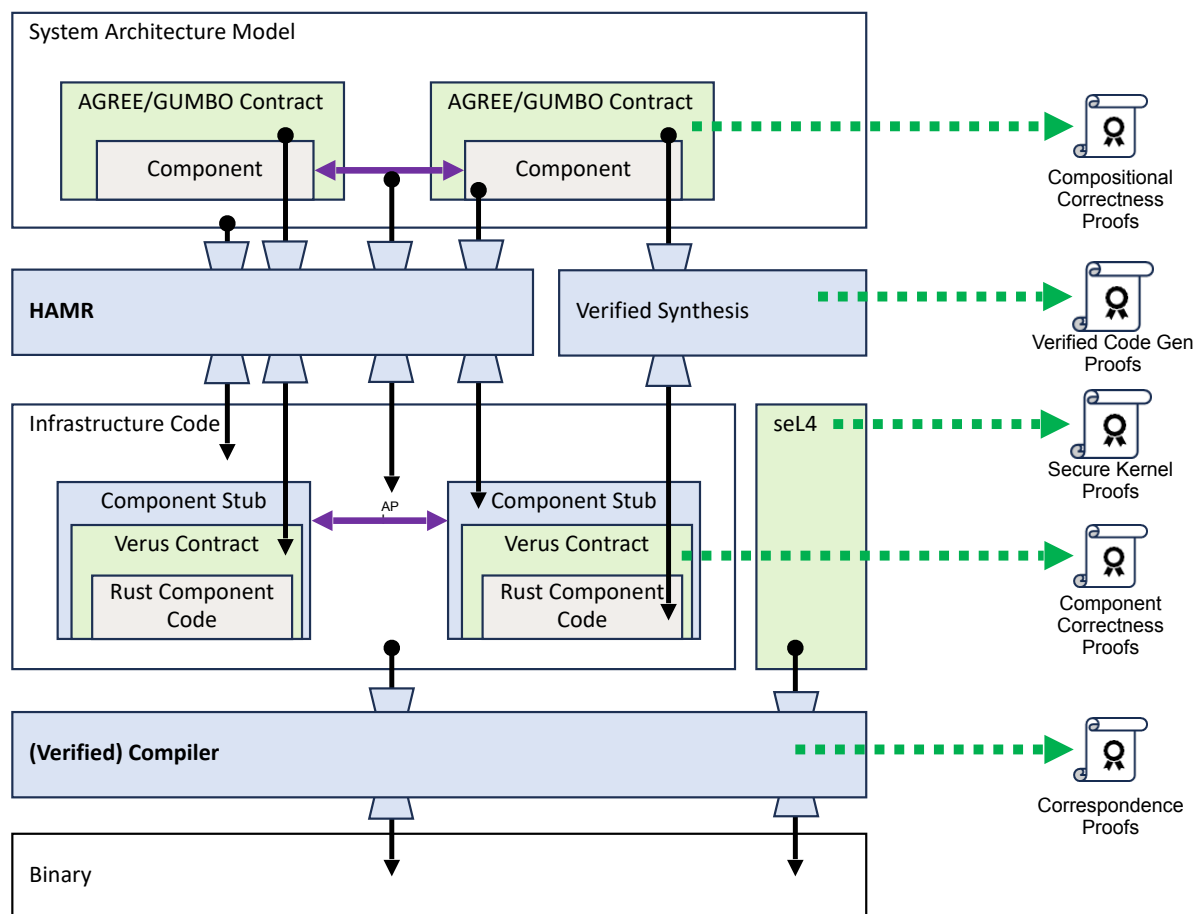


- Restricted Platform
 - Collins Launched Effects (LE) Mission Computer
 - Based on same computer hardware as the Open Platform



INSPECTA Proof Chain

With INSPECTA, engineers are able to generate comprehensive formal assurance across the entire development stack without requiring deep formal methods expertise



SysML v2

- SysML v2 is the second major version of the Systems Modeling Language
 - Standardized under the auspices of the Object Management Group
- Improved expressiveness relative to SysML v1
 - Now similar in expressiveness to AADL
- Standard textual form in addition to the graphical form
 - Promotes third party tool interaction
- Supported by major tool vendors: Siemens, The Mathworks, etc.
 - Necessary for mass adoption by the Defense Industrial Base

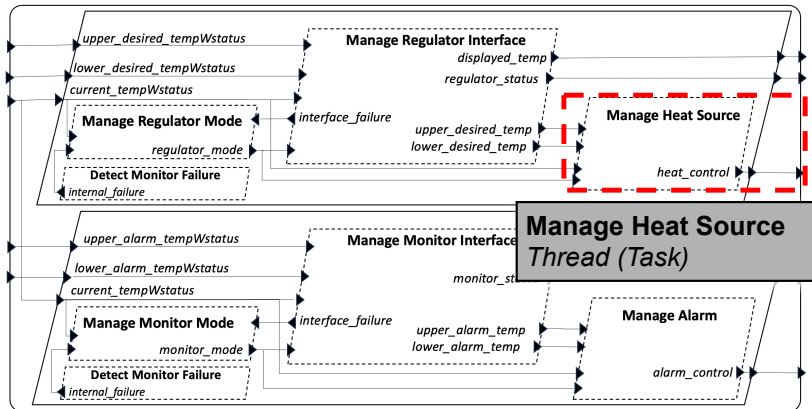


Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

AADL to SysML v2 Transition Example



AADL

```
thread Manage_Heat_Source
```

features

```
current_tempWstatus: in data port Isolette_Data_Model::TempWstatus.impl;
lower_desired_temp: in data port Isolette_Data_Model::Temp.impl;
upper_desired_temp: in data port Isolette_Data_Model::Temp.impl;
regulator_mode: in data port Isolette_Data_Model::Regulator_Mode;
heat_control: out data port Isolette_Data_Model::On_Off;
```

properties

```
Dispatch_Protocol => Periodic;
Period => Isolette_Properties::ThreadPeriod;
```

SysMLv2 + AADL Library

```
part def Manage_Heat_Source_i :> Thread {

    in port current_tempWstatus : DataPort { in :> type : Isolette_Data_Model::TempWstatus_i; }
    in port lower_desired_temp : DataPort { in :> type : Isolette_Data_Model::Temp_i; }
    in port upper_desired_temp : DataPort { in :> type : Isolette_Data_Model::Temp_i; }
    in port regulator_mode : DataPort { in :> type : Isolette_Data_Model::Regulator_Mode; }
    out port heat_control : DataPort { out :> type : Isolette_Data_Model::On_Off; }

    attribute :>> Dispatch_Protocol = Supported_Dispatch_Protocols::Periodic;
    attribute :>> Period = 1000 [millisecond];
    attribute Domain: CASE_Scheduling::Domain = 9;
```

Kansas State University



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

<https://github.com/loonwerks/INSPECTA-models/tree/main/isolette/sysml>

GUMBO Contract Language

- Inspired by AGREE and BLESS
- Aligns with MBSE run-time semantics
- Programming language independent
- Supports multiple quality assurance techniques
- Language Features:
 - Data type invariants
 - Port invariants (integration constraints)
 - Event-based / Shared-data based inter-thread communication
 - Local state declarations with invariants
 - Pre/Post conditions for thread code entry points
 - Support for fixed width scalars (e.g., Float32)

GUMBO contracts are specified in AADL/SysML v2 threads

```

thread Manage_Heat_Source
  features
    -- ===== INPUTS =====
    -- ("Current Temperature") - current temperature (from temp sensor)
    current_tempstatus: in data port Isolette_Data_Model::TempStatusImpl;
    -- ("Desired Range") - lowest and upper bound of desired temperature range
    lower_desired_temp: in data port Isolette_Data_Model::TempImpl;
    upper_desired_temp: in data port Isolette_Data_Model::TempImpl;
    -- ("Regulator Mode") - subsystem mode
    regulator_mode: in data port Isolette_Data_Model::Regulator_Mode;

    -- ===== OUTPUTS =====
    -- ("Heat Control") - command to turn heater on/off (actuation command)
    heat_control: out data port Isolette_Data_Model::On_Off;

  properties
    Dispatch_Protocol => Periodic;
    Period => Isolette_Properties::ThreadPeriod;

    Stack_Size => Isolette_Properties::StackSize;

  annex GUMBO {
    -- indicate that the component maintains an internal state (variables) that influence its behavior
    state
      lastCmd: Isolette_Data_Model::On_Off;

    -- ===== Initialize Entry Point Behavior Constraints =====
    initialize
      guarantee
        initlastCmd: lastCmd == Isolette_Data_Model::On_Off_Off;
      guarantee REQ_MHS_1 "If the Regulator Mode is INIT, the Heat Control shall be
        lset to Off";
        heat_control == Isolette_Data_Model::On_Off_Off;

    -- ===== Compute Entry Point Behavior Constraints =====
    compute
      -- assumption on set points enforced within the Operator Interface
      assume lower_is_lower_temp: lower_desired_temp.value <= upper_desired_temp.value;
  }

```

LLMs for MBSE Contract Verification: Counterexample Analysis/Resolution

- Counterexamples generated from MBSE contract verification can be difficult to analyze by non-experts
- We are utilizing LLMs to analyze these counterexamples, and suggest repairs
- Any LLM hallucinations are rejected, because assume/guarantee contract analysis is performed by a mathematically rigorous model checker
- We are also exploring use of LLMs for:
 - Proof repair
 - Documentation assistance
 - Model updates
 - Help writing formal properties

System Model

```
workspace - Car/packages/Car.aadl - OSATE2
Integer_Toy.aadl x Car.aadl x
property const_car_speed =
  true -> Target_Speed.val = pre(Target_Speed.val)
  guarantee G_car_1 "actual speed is less than constant target speed"
  guarantee G_car_2 "acceleration is limited"
  Agree_Nodes::abs(AcSpeed.val - prev)
end Car;
system implementation Car.Car_Impl
subcomponents
  THROT: system Transmission::Throttle.Throttle_Impl;
  CNTRL: system Transmission::Speed_Control.Speed_Control;
  AXL: system Steering::Axle.Axle_Impl;
  SM: system Transmission::State_Machine.State_Machine_Impl;
connections
  SpeedToThrot: port CNTRL.Actuator_Input -> THROT.Actuator_Input;
  AcSpeedToTop: port THROT.Actual_Speed -> Actual_Speed {Communication_Properties::Timing => immediate};
  AcSpeedToCntrl: port THROT.Actual_Speed -> CNTRL.Actual_Speed {Communication_Properties::Timing => immediate};
  TgSpeedToCntrl: port Target_Speed -> CNTRL.Target_Speed {Communication_Properties::Timing => immediate};
  AcPtchToAxl: port Target_Tire_Pitch -> AXL.Target_Tire_Pitch {Communication_Properties::Timing => immediate};
  AcPtchToCar: port AXL.Actual_Tire_Direction -> Actual_Tire_Direction {Communication_Properties::Timing => immediate};
  SSToSM: port SM.State_Out -> State_Signal {Communication_Properties::Timing => immediate};
end Car.Car_Impl;
end Car;
-- Stop Here
```

Contract counterexample

LLM prompts

```
Dash
THROT.Actual.val 11 10 9.0909090909090909
THROT.Actuator_Input 110 -10 -9.0909090909090909
###
AGREE-Dog: The counter-example provided by AGREE indicates that the guarantee 'G_car_1' ("actual speed is less than constant target speed") has been violated. Let's analyze the counter-example step by step:
### Counter-example Analysis:
1. **Variable Values:**
  - 'Actual_Speed.val': 11, 10, 9.0909090909090909
  - 'Target_Speed.val': 121, 0, 0
2. **Guarantees and Properties:**
  - 'const_tar_speed': 'true -> Target_Speed.val = pre(Target_Speed.val)'
  - 'G_car_1': 'const_tar_speed => (Actual_Speed.val <= Target_Speed.val)'
### Key Points:
1. **Initial State:**
  - At time 0, 'Actual_Speed.val' is 11 and 'Target_Speed.val' is 121. The guarantee 'G_car_1' holds true because '11 <= 121'.
2. **Time Step 1:**
  - 'Actual_Speed.val' is 10 and 'Target_Speed.val' is 0. Here, the guarantee 'G_car_1' is violated because '10 > 0'.
### Issue:
The violation occurs because the 'Target_Speed' is changing to 0 in the next time step, which makes the guarantee 'G_car_1' invalid as the 'Actual_Speed' is not allowed to catch up immediately to the new 'Target_Speed'.
### Solution:
To fix the problem, we need to ensure that the 'Target_Speed' remains constant for the duration over which the 'Actual_Speed' is being compared. This can be done by introducing a delay or ensuring that the 'Target_Speed' does not change abruptly. We can modify the property 'const_tar_speed' to accommodate a scenario where the
```

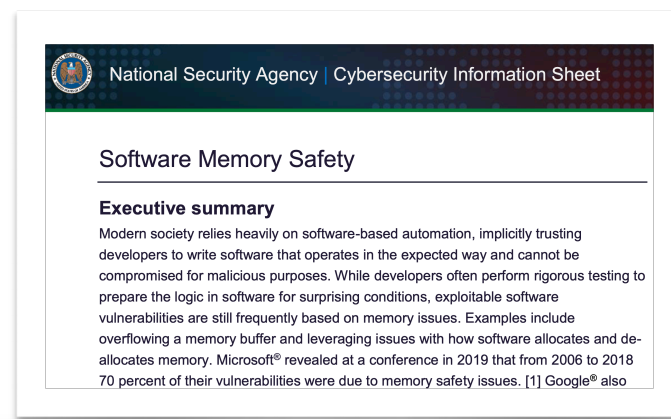
Analysis

Detection

Repair

Memory-Safe Programming Languages

- An emerging consensus amongst computer science thought leaders is that memory-safe programming language technology needs to be adopted more broadly:
- “NSA recommends using a memory safe language when possible.” (Nov. 2022)
- The White House has published a report championing the adoption of memory safe programming languages to enhance software security. (Feb. 2024)
- Microsoft, Google, and Amazon have all announced significant Rust initiatives.
- Memory-safe language requirements are beginning to appear in U.S. Government contracting.



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

Why Memory-Safe Languages? Why now?

- Memory-safe languages are not new
 - For example, Collins successfully used Ada in major commercial and government avionics products in the 1980s and 1990s
 - Collins used SPARK effectively on high-assurance products for the intelligence community in the 2000s
- Recent improvements in compiler technology have made memory safety very low cost
- Additionally, novel memory ownership models (e.g, in Rust) have allowed references to be used safely
- Development organizations have tired of continual memory errors, causing a never-ending parade of security vulnerabilities, despite the use of increasingly sophisticated analysis tools



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

The Rust Programming Language

- The INSPECTA team is focusing our memory-safe language research on Rust
- Rust has several assurance advantages over C/C++, including:
 - Improved type safety
 - Vastly improved memory safety
 - No arbitrary pointer arithmetic
 - ...in short, **80% of C/C++ security flaws are eliminated outright!**
- Rust supports modern programming idioms such as a match primitive, traits, immutability by default, etc.
- Basic Rust syntax is familiar to C/C++ developers, easing the transition
- The Rust compiler produces code which is competitive in speed to C/C++



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

Verus: Rust Code Verification

- Verus is an open source Rust code verification environment under development by Carnegie Mellon University and numerous other researchers
- Verus has been utilized in a number of operating system, concurrent data structure, and distributed algorithm verification efforts
- Verus utilizes Rust syntax to express precondition and postcondition annotations, loop invariants, etc.
- Verus employs an SMT solver to attempt to prove postconditions, given the preconditions



Rust-Related Work on INSPECTA

- HAMR now supports the generation of Rust source code from SysML v2 models
 - For seL4, we use a new Rust userspace API developed by Nick Spinale
 - The KSU/Aarhus team is translating GUMBO system model contracts to the Verus Rust verification environment
- The University of Kansas is developing Rust code generation for their attestation protocol specifications written in the Rocq theorem prover
- Dornerworks is writing open model application code in Rust
- CMU is enhancing Verus to support INSPECTA, reducing fragility in their SMT backend, and creating a connection to the Lean theorem prover



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

SysML v2 model with GUMBO contracts translated to Rust/Verus by the KSU HAMR tool

```
inspecta > src > @ manage_heat_source.rs > {} impl ManageHeatSource > time_triggered
78 impl ManageHeatSource
130
131 // verus contract
132 requires
133   old(self).api.lower_desired_temp <= old(self).api.upper_desired_temp
134 ,
135 ensures
136   ((old(self).api.regulator_mode == RegulatorMode::INIT)
137    ==> (self.api.heat_control == HeatControl::OFF))
138   &&
139   ((old(self).api.regulator_mode == RegulatorMode::FAILED)
140    ==> (self.api.heat_control == HeatControl::OFF))
141   &&
142   ((old(self).api.regulator_mode == RegulatorMode::NORMAL &&
143     old(self).api.current_temp < old(self).api.lower_desired_temp)
144    ==> (self.api.heat_control == HeatControl::OFF))
145   &&
146   ((old(self).api.regulator_mode == RegulatorMode::NORMAL &&
147     old(self).api.current_temp > old(self).api.upper_desired_temp)
148    ==> (self.api.heat_control == HeatControl::OFF))
149   &&
150   ((old(self).api.regulator_mode == RegulatorMode::NORMAL &&
151     old(self).api.lower_desired_temp <= old(self).api.current_temp <= old(self).api.upper_d
152    ==> (self.api.heat_control == old(self).last_command))
153   &&
154   self.last_command == self.api.heat_control
155 ,
156 {
157   // rust implementation
158   self.api.heat_control = match self.api.regulator_mode {
159     RegulatorMode::INIT => HeatControl::OFF,
160     RegulatorMode::NORMAL =>
```

```
136 /
137   ((old(self).api.regulator_mode == RegulatorMode::INIT)
138    ==> (self.api.heat_control == HeatControl::OFF))
139   &&
140   ((old(self).api.regulator_mode == RegulatorMode::FAILED)
141    ==> (self.api.heat_control == HeatControl::OFF))
142   &&
143   self.last_command == self.api.heat_control
144   - failed this postcondition
145 /
146 {
147   // rust implementation
148   self.api.heat_control = match self.api.regulator_mode {
149     RegulatorMode::INIT => HeatControl::OFF,
150     self.last_command = self.api.heat_control;
151   }
152   ^ at the end of the function body
error: aborting due to 1 previous error
verification results: 8 verified, 1 errors
JB:inspecta junaid$
```



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

Carnegie Mellon University

25

Conclusion

- The INSPECTA team is **making formal verification** across the entire software development stack **accessible to non-formal methods experts** through automated analysis, DevOps integration, a ProofOps console, and improved user feedback
- Keys to achieving this goal include integration with the SysML v2 System Modeling Language, and support for modern memory-safe languages, specifically Rust
- Much important INSPECTA Research was not mentioned in this talk, including:
 - AGREE/GUMBO contract language harmonization (Collins / KSU)
 - Verified Component Synthesis (KU)
 - Lifecycle Attestation (KU / Collins)
 - seL4 proof engineering (Proofcraft), Microkit, and Lions OS (UNSW)
- Check it out – code, papers, links:
 - <https://loonwerks.com/projects/inspecta.html>



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.



Thank You!

Contact: david.hardin@collins.com

This work was funded by DARPA contract FA8750-24-9-1000.

The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.



Collins Aerospace

An RTX Business

© 2025 Collins Aerospace. | This document does not include any export controlled technical data.

