

ACTIVE: A Tool for Integrating Analysis Contracts

Ivan Ruchkin¹ Dionisio De Niz² Sagar Chaki² David Garlan¹

¹Institute for Software Research, Carnegie Mellon University, USA, {iruchkin, garlan}@cs.cmu.edu

²Software Engineering Institute, Carnegie Mellon University, USA, {dionisio, chaki}@sei.cmu.edu

Abstract

Development of modern Cyber-Physical Systems (CPS) relies on a number of analysis tools to verify critical properties. The Architecture Analysis and Design Language (AADL) standard provides a common architectural model to which multiple CPS analyses can be applied. Unfortunately, interaction between these analyses can invalidate their results. In this paper we present ACTIVE, a tool developed within the OSATE/AADL infrastructure to ensure correct analysis interaction. We describe the problems that occur when multiple analyses are applied to an AADL model and how these problems invalidate analysis results. Interactions between analyses, implemented as OSATE plugins, are formally described in ACTIVE in order to enable automatic verification. In particular, these interactions are captured in an analysis contract consisting of inputs, outputs, assumptions, and guarantees. The inputs and outputs help determine the correct order of execution of the plugins. Assumptions capture the conditions that must be valid in order to execute an analysis plugin, while guarantees are conditions that are expected to be valid afterwards. ACTIVE allows the use of any generic verification tool (e.g., a model checker) to validate these conditions. To coordinate these activities our tool uses two components: ACTIVE EXECUTER and ACTIVE VERIFIER. ACTIVE EXECUTER invokes the analysis plugins in the required order and uses ACTIVE VERIFIER to check assumptions and guarantees. ACTIVE VERIFIER identifies and executes the verification tool that needs to be invoked based on the target formula. Together, they ensure that plugins are always executed in the correct order and under the correct conditions, guaranteeing correct results. To the best of our knowledge, ACTIVE is the first extensible framework that integrates independently-developed analysis plugins ensuring provably-correct interactions.

Keywords analysis contracts, cyber-physical systems, model checking, virtual integration

1. Introduction

Development of Cyber-Physical Systems (CPS) relies on analysis tools that use their own specialized abstractions. These abstractions, however, interact with each other. Neglecting these interactions leads to incorrect analysis results and design choices. For instance, selecting an otherwise valid scheduling policy may violate the assumptions of a functional correctness analysis (e.g., model checking) and lead to a deadlock-prone system being declared safe. Alternatively, such a scheduling policy could violate the assumptions of a processor frequency scaling analysis and lead to a non-schedulable allocation of tasks. Similarly, modifying a controller algorithm may alter its execution time and periodicity, thereby affecting schedulability [3].

Verification tools are particularly sensitive to their underlying domain abstractions, and incorrect application of these abstractions leads to invalid results. For instance, the original schedulability equations for Rate-Monotonic Scheduling (RMS) [6] use task abstractions that both restrict these tasks to be independent of one another and to forbid them to pause their computation. The application of this abstraction to tasks that do not honor these restrictions leads to incorrect schedulability results via RMS.

Virtual integration aims at addressing the issue of dependent abstractions by providing methods to resolve dependencies and conflicts among tools and models used in CPS engineering [7]. In our recent work we developed an approach for virtual integration based on the specification and verification of *contracts* between analysis tools [8]. Our approach allows us to describe and verify the interactions between analyses from different scientific domains.

Our contract verification approach relies on specification of *verification domains*, which are comprised of a set of symbols and their interpretation that enables us to describe analysis contracts from that domain in a precise manner. Examples of verification domains include the *scheduling* and *battery* domains, which are concerned with analyses of properties of real-time threads (e.g., schedulability) and batteries (e.g., thermal runaway), respectively. Formally, a verification domain is a signature to a mixed-logic language in which dependencies, assumptions, and guarantees of each analysis, that belongs to this domain, are specified as a contract. These contracts are algorithmically evaluated against a system architecture to verify if the system satisfies the assumptions of the analyses, and if the contracts of the analyses are compatible with one another.

Previously, we focused [8] on the theoretical foundations of the analysis contract approach. However, a number of significant practical issues were not addressed. For instance, the initialization and execution of CPS analysis and verification tools have many peculiarities. For example, many user-facing tools cannot be easily incorporated into a bigger toolchain. At the same time, specifications of verification domains and contracts must be reused to avoid unnecessary duplication, and, on the other hand, need to be adjusted correctly to each model’s context. The missing part is a virtual integration platform to manage tool interactions. One key requirement for such a platform is extensibility: adding new tools and verification domains should be simple so that analysis integration is practical.

In this paper, we present our tool ACTIVE¹ (Analysis Contract Integration Verifier), which implements the analysis integration approach. ACTIVE was developed on top of the OSATE2² toolkit, and uses the AADL [4] architectural description language. AADL offers a convenient way to represent the structural aspects of the system that determine critical quality attributes in real-time systems. Moreover, AADL provides an annex mechanism to embed custom sublanguages in its models. We use an AADL annex to define a mixed specification language that includes first-order logic and Linear Temporal Logic (LTL) [9] to describe both static and dynamic properties of the system. This allows us to use state-of-the-art verifiers that target different parts of the specification.

Given that our goal is to develop an *extensible* platform for virtual integration of CPS analyses, we incorporate several analysis tools (e.g., a bin packing algorithm for thread-to-processor allocation) and contract verifiers (like SMT and Spin) into ACTIVE. To handle tool integration, ACTIVE includes (i) a language to describe analysis contracts, (ii) a mechanism to execute analyses in the correct order, and (iii) a contract verification engine to verify if the analysis contracts hold. All three parts are extensible with new verification domains (e.g., controller simulation) and types of verification (e.g., probabilistic model checking).

The rest of the paper is organized as follows. In Section 2 we break down and discuss the management of analysis tool interactions on the example of AADL analysis plugins. Section 3 focuses on the representation of analysis interactions in ACTIVE, while Section 4 and 5 demonstrate our solutions to behavioral challenges of using multiple analysis and verification tools. Section 6 concludes this paper.

2. Managing Interaction of CPS Analysis Plugins

AADL provides a description language to capture main software and hardware structures of a CPS (e.g., threads and processors) and subsystems they belong to, along with connections between structures and subsystems. Components are annotated with properties from various scientific domains. Examples of such properties include thread period, processor frequency, and scheduling policy (all from

the real-time scheduling domain), and battery dimensions, cell scheduling algorithm, and required battery voltage (from the battery domain). A system designer defines component types along with their subcomponents, interconnections, and properties to create the *AADL declarative model* of a system. This model is then transformed into an *AADL instance model* – an XML-based representation of the actual system, rather than component types.

In OSATE, analysis algorithms are implemented as plugins that have access to both the declarative and the instance AADL models. AADL allows designers to augment these models with tool-specific descriptions known as *annexes* that define sublanguages embedded in AADL descriptions. For example, the error model plugin [4] relies on an annex to specify error sources and propagations. It reads components and properties from the instance model and error annex specifications from the declarative model to produce, for example, a fault impact report.

While OSATE supports rapid development of analysis plugins, it does not support controlling undesirable interactions between different plugins in order to prevent incorrect results. In particular, OSATE provides extension points to add annexes (via sub-language grammars and compilers) and analysis invocation actions (via toolbar buttons or menu items). Once a plugin is invoked, it is given access to AADL models to carry out the desired operations and return the control back to the user. Each analysis plugin thus tends to focus on a specific technical concern, accessing components, properties, and annex clauses relevant to that concern, but ignoring the effects of other plugins.

The lack of support to capture and control the effects that plugins have on each other can lead to incorrect results. For instance, a plugin that modifies the thread-to-processor allocation might invalidate the result of any prior schedulability plugin applied to the affected threads and processors. Similarly, if a plugin depends on the scheduling policy of a processor (say Rate-Monotonic Scheduling) and another plugin modifies this policy (to, say, Earliest-Deadline First) the results could be invalidated. Worst of all, more often than not, these errors go unnoticed given the lack of an infrastructure to capture and verify these subtle plugin interactions. Unfortunately, OSATE currently lacks ways to manage the co-operative execution of analyses, let alone verify their assumptions formally.

To address the issue of CPS analysis interaction systematically, an *integration tool* needs to manage analysis plugins at the level of the abstractions that these analyses use. We identify three parts of this problem: (i) describing plugin interactions with enough detail to not only capture real conflicts but also avoid signaling false ones; (ii) executing plugins in a right sequence and at a right time; and (iii) verifying applicability assumptions of analysis plugins. All three parts need to be addressed in an *extensible way* so as to not lose the benefits of the flexible OSATE design. We go deeper into each part of the problem below.

2.1 Representing Plugin Interactions

In order to enable reasoning about plugin interactions, we must model these interactions by representing data flows

¹ Can be downloaded at <https://github.com/bisc/active>

² https://wiki.sei.cmu.edu/aadl/index.php/Osate_2

between the plugins and the AADL models. In other words, we must identify the parts of the model that different plugins change. This dataflow description would allow us to prevent the incorrect order of plugin execution where the output from one plugin is invalidated by the output from another plugin executed afterwards.

Also, we must represent the assumptions a plugin makes about the model under analysis. For example, several implementations of schedulability tests rely on rate-monotonic assumptions without stating them explicitly. More importantly, plugin assumptions may be related to the system behavior rather than its structure. Thus predicates over only the structural models in AADL are inadequate for expressing such assumptions. For example, one plugin may assume that a task may only be preempted by others with shorter deadline than its own. This assumption is trivially satisfied if we used Deadline Monotonic Scheduling. However, it could also be satisfied with other scheduling policies under certain specific system configurations. This can only be discovered with enough information about the behavior of the system with the different policies.

Finally, the way an analysis is identified in a specification must reflect the operations performed by the analysis on a model precisely. Many plugins come with several related operations; for example, a resource allocation plugin provides three operations: allocation of threads to processors (what we further call bin packing), a utilization-based schedulability test, and a priority inversion test. Just mentioning the name of the resource allocation plugin does not differentiate the operations properly.

One of the biggest challenges to represent plugin interactions come from the extensibility requirement. First of all, we must be able to change the plugin interaction specification independently from the plugin implementation. This is because the AADL semantics allows different plugins to have slightly different interpretations of the properties. For instance, we have seen AADL models where threads are assumed to be periodic by default while other models explicitly require the use of the `Periodic` value in the `Dispatch_Protocol` property, and yet others use the `Hybrid` value for the same property.

Another important aspect of the extensibility is the support of new domains that may introduce new types of components, properties, or even behavioral specification constructs. For example, to introduce analyses for multi-cell reconfigurable batteries, one would have to introduce a new device type, with new AADL properties (such as battery cell number and geometry), and runtime properties to capture dynamic cell connectivity. These changes must affect only AADL models, and not the plugin integration tool.

2.2 Correct Plugin Execution

To ensure the correct execution of a plugin it is necessary to respect its data dependencies and ensure that its assumptions are never violated as different plugins are executed. In practice, a *correct plugin execution* implies the following steps: (i) before the plugin is called we need to ensure that all the other plugins have finished their work and committed their changes to the AADL model; (ii) the plugin

assumptions must be validated on the model on which the plugin will be run; (iii) if an analysis plugin ends with an error, the sequence of execution has to be stopped.

Many OSATE plugin are made with human user interaction in mind: they expect to be run from toolbars and menus. Since OSATE discourages plugin interaction, it is challenging to invoke plugins programmatically. Once a plugin has been invoked, monitoring its progress is another challenge: many plugins use tools external to OSATE, without any feedback. The limited feedback mechanisms in OSATE were also designed with a human user in mind, making it difficult to monitor analysis execution and determine the time when it is safe to start the next analysis.

Like the specification, correct plugin execution needs to be achieved with minimal changes to existing plugins and the OSATE tool. Plugins still need to run individually on user's command, and major changes to their control flow are not acceptable. OSATE cannot be profoundly modified either. In particular, it is important to leave the option open to run plugins without the proper integration in cases when specifications have not been completed, or they do not hold.

2.3 Extensible Assumption Verification

An analysis plugin integration tool needs to use state-of-the-art verifiers to ensure that analyses are only used when they are known to produce correct results. Since the analysis plugin integration problem manifests itself in multiple domains that contribute to CPS, verifiers need to be tailored to domain-specific abstractions and applied in their corresponding contexts. For example, a dynamic model of a thread scheduler cannot be used to verify assumptions about a flight controller behavior. In addition, given the existence of specifications involving static and dynamic properties of the system it is necessary to enable the verification of both types of properties in a scalable way.

AADL models are hierarchical: a set of threads may be composed in a thread group, which in turn contributes to a software subsystem. A software subsystem may be part of a computational subsystem, which also includes processors and memory devices. Finally, the computational subsystem is part of the whole system, which also includes physical devices (such as rotors) and properties (such as mass). Choice of a hierarchy is left up to the designer, allowing multiple ways to describe the same system.

Unfortunately, many verifiers rely on their own system decomposition, which may not agree with AADL. Tools for timed automata verification, such as UPPAAL [5], use the refinement relation rather than an arbitrary composition. Hence, our challenge is to create mechanisms that use verifiers (some of which are unknown yet) for a custom-built AADL model at a proper level of hierarchy. Domain-specific verifiers, for instance, need to access their relevant abstractions without being aware of other components.

Management of analysis plugin interactions is a major obstacle to creating an extensible platform for virtual integration of CPS analyses. Roughly corresponding to the parts of this problem, the three following sections present the main components of our tool ACTIVE. We view our solutions from two positions: functionality – achieving their

goal correctly, and extensibility – allowing addition of new elements, such as analyses and verifiers, to the platform.

3. Contract Language as AADL Annex

The location of the information that defines the interaction between analysis plugins is critical to the extensibility of our platform. We evaluated three potential locations: (i) inside the analysis plugin itself; (ii) in the OSATE tool; or (iii) in the AADL model. Encapsulating the specification inside the plugin has the benefit that the description always travels with the plugin. On the other hand, storing the specification in a central database of specifications inside the OSATE tool facilitates collaboration and reuse of contracts. Unfortunately, both of these options diminish extensibility given that the exact specification of the interactions may change depending on the project and the plugins used in it, as discussed in Section 2.1. That is, semantic variations in the interpretation of AADL in a particular model need to be accounted for without changing the plugin or OSATE.

In ACTIVE we use an AADL annex (like many analysis plugins do) to represent analysis dependencies and assumptions. In this case, an annex instance is attached to a declarative AADL model and can be used whenever this model, or any derived instance model, is used. We refer to our analysis interaction specification as an *analysis contract*. An analysis contract has the following parts:

- **Name:** names the analysis contract and the plugin wrapper to be called when an analysis is invoked (in Section 4 we explain analysis plugin wrappers in detail).
- **Input:** a comma-separated list of elements: AADL component types (e.g., `thread`) and property names (e.g., `thread.Period`). The property names are prefixed with a component type to identify the dependency more precisely. By including a component type in its input, an analysis plugin declares that it accesses the set of components; and by including a property, it declares that it reads the property values from an instance model. If a property is included as part of the input or output, the related component is also included implicitly.
- **Output:** a list of same type of elements as that of the input. The only difference is semantic: the analysis declares that it changes the set of components or the values of a property. Although the specification of inputs and outputs is in terms of AADL types, the changes are meant to be done to the instance model.
- **Assumes:** a set of assumptions that must hold for an analysis to be applicable. Each assumption is a logical formula, as explained below. Given that we may not have a complete definition of when an analysis is applicable, assumptions describe at least the sufficient conditions of its applicability. If an analysis is always applicable, this part can be omitted.
- **Guarantees:** a set of formulas that must hold on the model after the analysis terminates. These formulas are syntactically equivalent to the assumption formulas. They can be used to satisfy the assumptions of other

contracts. In such a case, these assumptions do not need to be reverified. However, if the guarantees are not met, the assumptions of other contracts that depend on it must be reverified.

Assumes and guarantees contract formulas have the following syntax in ACTIVE:

```
ContractFormula ::=
(
  <Quan>
  (<Var>:<Type> ',')+
  ( '|' <PredicateExpression> )?
  ':'
)?
<LTLExpression>
```

Contract formulas consist of two subformulas: an optional first-order quantification and a mandatory main expression, often written in LTL. In the first subformula, `<Quan>` is a first-order quantifier, taking a value of either `forall` or `exists`. `<Var>` introduces a quantified variable name of an AADL component having type `Type`. Variables are quantified over `<PredicateExpression>`, which is a logical predicate over the AADL's model components and properties with the usual logical operators `and`, `or`, and `not`. The `<LTLExpression>` encodes a domain-specific behavioral property using a combination of logical operators above and the LTL modalities `Globally G`, `Eventually F`, and `Until U`. However, if necessary, `<LTLExpression>` can be limited to predicate logic.

The operator “:” is implicative when used in a `forall` formula (where it denotes that “all variable valuations that satisfy condition `<PredicateExpression>` should also satisfy `<LTLExpression>`”) and conjunctive when used in a `exists` formula (where it denotes that “all variable valuations that satisfy condition `<PredicateExpression>` should also satisfy `<LTLExpression>`”). We choose our mixed-logic language over Quantified LTL [9] because the latter prevents a cleaner split of formulas between general-purpose SMT solvers and domain-specific verifiers (see details in Section 5), and brings in unnecessary complexity.

Consider an example from the scheduling domain. Figure 1 shows a contract for a processor frequency scaling analysis. The goal of this analysis is to minimize the processor frequency to limit energy expenditure of the system. It reads threads, processors, thread deadlines, and thread bindings (allocations) to processors. It outputs the CPU frequency. An implicit assumption of this analysis is that threads run under the deadline monotonic scheduling policy. This is captured in a formula stating that “every pair of distinct threads allocated on the same processor should behave as if scheduled by a deadline-monotonic policy.” The first part of the formula (before the colon) indicates the condition for which all possible pairs of threads should be evaluated. The second part (after the colon) is an LTL expression that features a domain-specific predicate `CanPreempt`, which is true in any runtime state iff `t1` is executing but `t2` is ready to execute, but not executing.

```

-- Frequency scaling analysis
-- Adjusts processor frequency based on thread bindings
-- Assumes that threads on the same processor
--   behave equivalently to DMS
name FreqScalingAnalysis
input
  thread, thread.Actual_Processor_Binding,
  thread.Deadline, processor

output
  processor.Current_Frequency

assumes
  forall t1:thread, t2:thread | t1 != t2 &&
  t1.Actual_Processor_Binding[0] = t2.Actual_Processor_Binding[0]:
  G(CanPreempt(t1, t2) => t1.Deadline < t2.Deadline)

```

Figure 1: A contract for frequency scaling analysis.

```

-- LLrek thread model checking
-- Checks concurrency violations: deadlocks and race conditions
-- Assumes that preemption does not alternate
name LLrekAnalysis
input
  thread, thread.Period, thread.Deadline,
  thread.Source_Text, thread.Priority,
  thread.Compute_Execution_Time, thread.Priority,
  thread.Actual_Processor_Binding

output
  system.Is_Thread_Safe

assumes
  forall t: thread : t.Period = t.Deadline

assumes
  forall t1:thread, t2: thread | t1 != t2 :
  G(CanPreempt(t1, t2) => (G !CanPreempt(t2, t1)))

```

Figure 2: A contract for the LLREK analysis.

```

annex contract {**
  use contracts::SecureAllocationAnalysis,
  contracts::BinPackingAnalysis,
  contracts::FreqScalingAnalysis,
  contracts::LlrekAnalysis,
  contracts::ThermalRunawayAnalysis,
  contracts::BatterySchedulingAnalysis
**};

```

Figure 3: Annex subclause indicating the analyses to use.

Another example of an analysis is verification of safe concurrency based on the tool LLREK [1]: the tool takes source code of each thread annotated with safety assertions and determines whether the assertions are met. A contract for this analysis is shown in Figure 2. The analysis reads a number of thread properties and outputs whether the system was found to be safe with respect to its annotated assertions. LLREK has two assumptions: first, implicit deadline tasks, i.e., tasks whose relative deadlines are equal to their periods; second, it assumes fixed-priority scheduling, i.e., thread pre-emption is acyclic. In other words, if thread t_1 preempts thread t_2 , then t_2 never preempts t_1 . Other examples of contracts, for instance for electric and thermal analysis of multi-cell batteries, can be found in [8].

To improve convenience and reuse of contracts, we separate the definition of contracts (as in Figures 1 and 2), which we call a *library of contracts*, from the application of these contracts to a system, which we call a *usage subclause* (shown in Figure 3). Usage subclauses enable the association of analyses to models. In this way a user can control applicability of analyses at a macro-level, reusing the same contracts across different models.

Our approach explained in [8] relies not only on the specification of contracts (where the mapping between conceptual and practical aspects is more straightforward), but also on verification domains that define the formal underpinnings for both the specification and the analysis of the contracts within a verification tool (e.g., Spin). Formally, a verification domain σ is comprised of domain atoms \mathcal{A} , static functions \mathcal{S} , runtime functions \mathcal{R} , execution semantics \mathcal{T} , and domain interpretation for atoms and static functions $\llbracket \cdot \rrbracket_\sigma$. These elements are augmented by an architectural model that provides the interpretation $\llbracket \cdot \rrbracket_M$. Existence of a verification domain, with a semantics defined within a verification tool that automatically explores its behavior, guarantees correctness of our analysis contracts approach.

In ACTIVE, verification domains are not specified in one place, but are comprised of various elements of AADL and contract annexes. For some atoms $a \in \mathcal{A}$, $\llbracket a \rrbracket_\sigma$ is provided by OSATE. For example, integers, booleans, and reals are standard types in AADL. Other elements of \mathcal{A} are interpreted by the $\llbracket \cdot \rrbracket_M$, e.g., threads, processes, memory elements, processors, systems, and other sets of components.

Static functions \mathcal{S} map directly to AADL's properties, some of which are standard and some of which are defined by users in a declarative model. Only static functions can be used in `<PredicateExpression>` so that the semantics of `<PredicateExpression>` could be fully constructed based on the AADL instance model values. Standard types have a default value, which contributes to $\llbracket \mathcal{S} \rrbracket_\sigma$. For the most part, however, interpretation of static functions comes from $\llbracket \mathcal{S} \rrbracket_M$ in the form of values that properties have in a particular AADL instance model.

Unlike some static functions, runtime functions \mathcal{R} are strictly domain-specific, e.g., `CanPreempt`. Their interpretation comes from a domain-specific verifier and, as far as AADL models are concerned, these functions do not exist. Finally, the execution semantics \mathcal{T} is defined by a combination of static function specified by the model (e.g., thread periods and deadlines) and verifier-specific runtime behavior (e.g., how the state of system changes when a new thread arrives). Thus, all formal elements of σ are covered in ACTIVE, to which the formal conclusions of correctness can now be transferred.

4. ACTIVE EXECUTER

Ensuring that a set of analysis plugins is executed correctly requires more than the specification and verification of individual analysis contracts. Specifically, it requires coordination and proper sequencing of the execution of these analysis and their verifiers. This section describes ACTIVE EXECUTER— a plugin execution controller in our tool. The purpose of this controller is, on the one hand, to interact with the user and, on the other hand, to coordinate the execution of analyses. The scheme of ACTIVE EXECUTER is shown in Figure 4.

From the OSATE user's perspective, ACTIVE EXECUTER identifies dependencies between analyses, builds a dependency graph, and presents it to the user to allow him to select an analysis to run. When invoked on an instance

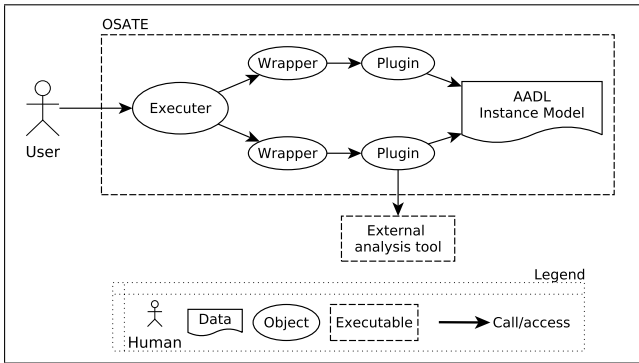


Figure 4: Operation of ACTIVE EXECUTER.

model, ACTIVE EXECUTER parses all usage subclauses that are in the scope of the system, retrieving the used contracts from annex libraries, and creating the dependency graph. In this graph, each vertex represents a contract with its corresponding analysis plugin. Edges in this graph represent input-output dependences between two contracts.

Here we say that an analysis plugin *depends* on another if the former reads a property or a component set that the latter modifies. Since inputs and outputs are specified in terms of component types and properties, any plugin that writes a component type depends on any plugin that reads a property of this type because, formally, the former changes the function’s domain. The opposite is not true: changing a property does not constitute a change in components.

A user selects an analysis to run from the dependency graph. ACTIVE’s selection dialog is shown in Figure 5: the rectangles represent analysis plugins and the arrows show the dependency relationships. Once a selection is made, ACTIVE EXECUTER finds a correct sequence of analyses leading up to the selected one (the formal details can be found in [8]). For example, for the analyses in Figure 5, the frequency scaling plugin would first require the execution of the secure allocation plugin followed by the binpacking plugin. The analyses in the sequence are then executed one by one, exchanging their data through the instance model.

Currently, ACTIVE only supports acyclic dependencies. In practice a cyclic dependency can be a symptom of incorrect contract specifications. However, when this is not the case and the cycle is not an error it may be possible to execute each analysis in the cycle repeatedly till convergence (i.e., executing any analysis in the cycle does not change the system further, or in other words, a fixed point is reached). We leave this investigation as future research.

Invoking the right analysis plugin is not as trivial as the theoretical aspects of our work make it appear. If ACTIVE’s code were to call analysis plugins directly, ACTIVE would have a direct dependency on a concrete set of plugins and would not be deployable separately from these plugins. Even further, plugins developed for human user have external access points such as toolbar buttons and menus, that cannot be called programmatically from another OSATE plugin. Hence, a more complex approach to invoking plugins is needed to achieve the desired extensibility.

To overcome these limitations and use analysis plugins without substantial changes to their external interface, we

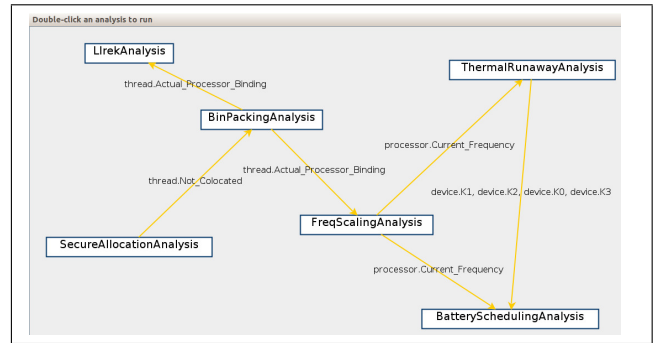


Figure 5: Analysis selection dialog in ACTIVE.

```
<!-- For integration with contracts: used to run a Binpack action -->
<extension
  id="binpack.command"
  point="org.eclipse.ui.commands">
  <category
    name="Resource management commands"
    id="org.osate.analysis.resource.management.commands">
  </category>
  <command
    categoryId="org.osate.analysis.resource.management.commands"
    id="org.osate.analysis.resource.management.commands.Binpack"
    description="Execute the binpacking action through the command framework"
    name="binpackcommand">
  </command>
</extension>
```

Figure 6: A plugin wrapper’s command interface for binpacking analysis.

developed *analysis plugin wrappers* – a method to execute analysis plugins using the Eclipse Command Framework³. A wrapper creates a command interface around the user action. A command, unlike an action, can be called programmatically, and results in a call of the associated action. Each plugin wrapper thus consists of a command interface – which is an addition to the plugin’s configuration exemplified in Figure 6 – and a direct command invocation that can be exercised by ACTIVE EXECUTER.

Another factor of correct plugin execution is understanding when it is safe to start the execution of the next plugin in a dependency chain. Starting a plugin too early would result in its reading of transient data that would be overwritten by the previous plugin. Therefore, the execution of every plugin needs to be monitored. Unfortunately, there is no direct way to do so in the original implementation of OSATE. Our ACTIVE EXECUTER relies on the plugin wrappers to perform this monitoring. Specifically, when a wrapper starts a command associated with an action, the progress of this command is tracked using the identity of the associated action. Once the wrapper reports that the plugin has finished, the next plugin is safely started.

To check whether the assumptions of a plugin hold before its execution, or if its guarantees hold afterwards, the ACTIVE EXECUTER calls ACTIVE VERIFIER – another key component of ACTIVE – passing over a contract and a formula to verify before and after the plugin’s execution. We discuss the details of this process in the next section.

5. ACTIVE VERIFIER

The verification of analysis assumptions is the most complex aspect of ACTIVE since it requires the combination

³wiki.eclipse.org/Platform_Command_Framework

of abstractions coming from diverse scientific fields into a common logic. Its goal is to take a contract formula and verify it against the current AADL instance model. This verification is, however, not limited to the AADL model, which abstracts away most of the behavioral system dynamics due to the model’s architectural nature. A further complication comes from assumption formulas with abstractions and properties from different verification domains, which manifest at different levels of the hierarchy in an AADL model. Last but not least, the addition of new verification tools and models, which is at the core of virtual integration, needs to be facilitated by minimizing changes from each such addition to the ACTIVE or verifier structure.

To address the challenges of multi-domain verification, we created the component ACTIVE VERIFIER. The first step in its operation is to deconstruct the contract formula to subformulas that can be processed by an individual verifier. The first subformula, including `<Quan>`, `<Var>`, and `<PredicateExpression>`, can be processed with a general-purpose SMT solver: as we showed in Section 3, all atoms and operators of `<PredicateExpression>` are determined from the AADL models, thus rendering quantified `<PredicateExpression>` amenable to an efficient validity check. The practical purpose of `<PredicateExpression>` is to identify the part of a model that the formula should apply to, thus providing a convenient access to the hierarchy and bypassing irrelevant parts. This way, the frequency scaling plugin, for instance, can indicate that it targets only pairs of threads running on the same processor (Figure 1), even if the processor is located in a hardware subsystem, separated from threads by several AADL hierarchy levels.

ACTIVE VERIFIER reduces the search for variable valuations that satisfy `<PredicateExpression>`⁴ to a SMT formula χ generated from the AADL components and properties with an added assertion of negated `<PredicateExpression>`. To construct χ , ACTIVE VERIFIER only explores the subset of the AADL model that includes the components and properties mentioned in the contract formula. It then checks the satisfiability of χ using an off-the-shelf SMT solver (currently Z3 [2]). If χ is SAT, the solution is recoded and blocked for the next run. If χ is UNSAT, the search is stopped, and the verification moves to process the `<LTLExpression>`. Thus, by using “blocking clauses” incrementally, ACTIVE VERIFIER generates all solutions of χ .

The verification of `<LTLExpression>`, however, cannot rely on a general-purpose SMT solver since `<LTLExpression>` may contain domain-specific runtime functions like `CanPreempt` and modal LTL operators. Thus, an `<LTLExpression>` needs to be matched with a domain-specific verifier, based on the verifier’s fitness. Specifically, we say that a verifier *matches a formula* if and only if this verifier can give an interpretation to every atom (such as set or function) and every operator in the formula. Typically this matching is somewhat known to en-

gineers familiar with particular verifiers, but not explicitly documented in a machine-readable form. Without a proper representation of matching, we would risk an error of running a verifier on an inappropriate formula and producing invalid results. For example, a non-preemptive scheduler model would report non-schedulability on many systems where a preemptive scheduler would report schedulability.

The problem of matching verifiers to contract formulas led us to develop *verification engines* – an abstraction to simplify the access to verifiers and determine formula matching. Each verifier is augmented with a verification engine that governs the application and execution of the verifier through the following functions:

- Verifier’s initialization and parameter selection. For instance, to run the Spin verifier, used in the scheduling domain, it is necessary to translate AADL properties into Promela instructions to complete the template of a Promela program for preemptive schedulers.
- Declaration of atoms and operators that can be interpreted by the verifier. For example, the scheduling Spin verification engine reports atom `CanPreempt` and LTL operators `G`, `F`, and `U`. A formula *matches* a verifier iff the formula’s atoms and operators are a subset of the ones interpreted by the verifier.
- Declaration of AADL model parts that are required to achieve full a semantic interpretation of \mathcal{T} . E.g., to generate traces of a thread scheduler, a Spin program needs to read thread periods from an AADL model.
- Interpretation of the results from the verifier. While in theory a verifier could have only two answers \top and \perp , in practice other options are possible: a verifier may detect a syntax error or run out of memory. These results do not constitute a violation of contract, and verification engines report those as “verification not possible,” thus making user interaction more transparent.

Thus, verification engines allow ACTIVE to handle diverse verifiers using a common interface.

The pseudocode of the end-to-end algorithm for verifying a contract formula is shown in Figure 7. In this algorithm ACTIVE EXECUTOR first calls the function VERIFY, with a contract formula and a model. If the SMT solver can fully handle the formula (i.e., there are no domain-specific atoms or operators), ACTIVE VERIFIER takes a shortcut and delegates the verification exclusively to the SMT. If this shortcut is not possible, ACTIVE VERIFIER searches for a matching verifier (function MATCH) and runs the selected verifier on every valuation of the variables (function RUN).

Up to this point we have discussed coordinated application of a set of verifiers. However, another important factor in simplifying the addition of new verifiers is the access to the parts of the model referenced by contract formulas. In particular, without a generic model access approach it would be necessary to write custom code for each verifier to access parts of the model located at different levels of its hierarchy. To simplify data collection for verifiers, we introduced the *shared-data interface* that helps to decouple domain-specific verifiers from the hierar-

⁴The algorithm is the same for `<LTLExpression>` without domain-specific atoms or LTL operators.

```

1: function VERIFY(ContractFormula  $f$ , Model  $m$ )
2:   if  $\neg f$ .isDomainSpecific() then
3:     return SMT.isValid( $f$ )
4:    $v \leftarrow$  MATCH( $f$ ,  $m$ )
5:   if  $v \neq$  null then
6:     return RUN( $f$ ,  $v$ )
7:   else
8:     return error
9: function MATCH(ContractFormula  $f$ , Model  $m$ )
10:   $ltlExp \leftarrow$   $f$ .<LTLExpression>
11:  for all VerificationEngine  $v$  do
12:    if  $v$ .canInterpretAtoms( $ltlExp$ )  $\wedge$ 
13:       $v$ .canInterpretOperators( $ltlExp$ )  $\wedge$ 
14:       $v$ .hasFullInterpretation( $m$ ) then
15:        return  $v$ 
16:  return null
17: function RUN(ContractFormula  $f$ , VerificationEngine
18:   $v$ )
19:  if  $f$ .<Var> = null then
20:    return
21:       $v$ .VERIFY (null,  $f$ .<LTLExpression>,  $m$ )
22:   $varEvals \leftarrow$ 
23:    SMT.solve( $f$ .<PredicateExpression>,  $m$ )
24:  if  $f$ .<Quan> = "forall" then
25:     $res \leftarrow$  true
26:    for all  $ve \in$   $varEvals$  do
27:       $res \leftarrow$   $res \wedge$ 
28:       $v$ .VERIFY ( $ve$ ,  $f$ .<LTLExpression>,  $m$ )
29:  else if  $f$ .<Quan> = "exists" then
30:     $res \leftarrow$  false
31:    for all  $ve \in$   $varEvals$  do
32:       $res \leftarrow$   $res \vee$ 
33:       $v$ .VERIFY ( $ve$ ,  $f$ .<LTLExpression>,  $m$ )
34:  return  $res$ 

```

Figure 7: Algorithm of ACTIVE VERIFIER

chy of the AADL model. The data interface provides SMT and domain-specific verifiers with SQL-based access to the AADL instance model data. All components are given unique identifiers and stored in a single table. Each AADL property is represented with a table of the same name that lists the values in a format appropriate for the property type, as well as the owner component of the property. This way, all components and properties are easily accessible to verifiers without the need to traverse levels of hierarchy in AADL. The downside of the shared-data interface is that it does not support composition of data types of arbitrary depth (e.g., a sequence of records, each having a field that is a set of other records). However, we are yet to see a plugin that uses more than three levels of recursion. We use a MySQL database to implement the shared-data interface.

6. Conclusion

In this paper we presented ACTIVE, a tool addressing the problem of AADL analysis plugin interactions. In particular, we discussed how analysis plugin integration errors pose a risk of invalidating analysis results without user's

knowledge. Solving this problem entails, first of all, representing analysis interactions in a formal way to enable automatic reasoning. The specification of these interactions, in the form of inputs, outputs, assumptions, and guarantees, allows us to determine the correct order in which plugins must execute. Finally, the assumptions and guarantees need to be verified using a potentially wide variety of verification tools. Using an AADL annex-based contract language allows ACTIVE EXECUTER to manage the proper startup and monitoring of the analysis plugins, making the appropriate calls to ACTIVE VERIFIER, which in turn invokes general-purpose SMT solvers (e.g., Z3) and domain-specific model checkers (e.g., Spin) for in-depth behavioral verification. These major ACTIVE components were designed to be extensible to accommodate new verification domains, analysis plugins, and domain-specific verifiers. To the best of our knowledge ACTIVE is the first extensible framework able to integrate analysis plugins guaranteeing their correct interaction and execution.

Acknowledgments

We thank Julien Delange, Peter Feiler, Lutz Wrage, and Joseph Siebel for their help in implementing ACTIVE.

Copyright 2014 Carnegie Mellon University. This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN AS-IS BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution. Carnegie Mellon is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University. DM-0001717

References

- [1] Sagar Chaki, Arie Gurfinkel, and Ofer Strichman. Efficient Verification of Periodic Programs using Sequential Consistency and Snapshots. In *Proc. of FMCAD*, 2014.
- [2] Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: An Efficient SMT Solver. In *Proc. of TACAS*, 2008.
- [3] P. Derler, E. A. Lee, S. Tripakis, and M. Torngrén. Cyber-physical system design contracts. In *Proc. of ICCPS*, 2013.
- [4] P.H. Feiler and D.P. Gluch. *Model-Based Engineering with AADL: An Introduction to the SAE Architecture Analysis & Design Language*. Addison-Wesley Professional, 2012.
- [5] Kim G. Larsen, Paul Pettersson, and Wang Yi. Model-checking for real-time systems. In *Proc. of Fundamentals of Computation Theory*, 1995.
- [6] C. L. Liu and James W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM*, 20(1):46–61, January 1973.
- [7] Panagiotis Manolios and Vasilis Papavasileiou. Virtual integration of cyber-physical systems by verification. In *Proc. of AVICPS*, 2010.
- [8] Ivan Ruchkin, Dionisio De Niz, Sagar Chaki, and David Garlan. Contract-based integration of cyber-physical analyses. In *Proc. of EMSOFT*, 2014.
- [9] Aravinda Prasad Sistla. *Theoretical Issues in the Design and Verification of Distributed Systems*. PhD thesis, Harvard University, Cambridge, MA, USA, 1983. AAI8403047.