

10-2014

Contract-Based Integration of Cyber-Physical Analyses

Ivan Ruchkin
Carnegie Mellon University

Dionisio de Niz
Carnegie Mellon University

Sagar Chaki
Carnegie Mellon University

David Garlan
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/isr>

 Part of the [Software Engineering Commons](#)

Published In

Proceedings of the International Conference on Embedded Software (EMSOFT), 2014, 1-10.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Institute for Software Research by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Contract-Based Integration of Cyber-Physical Analyses

ABSTRACT

Developing cyber-physical systems involves creating systems with properties from multiple domains, e.g., timing, logical correctness, thermal resilience, aerodynamics, and mechanical stress. In today’s industrial practice, multiple analyses are used to obtain and verify such properties. Unfortunately, given that these analyses originate from different scientific domains, they abstract away interactions among themselves, risking the invalidation of their results. Specifically, one challenge is to ensure that an analysis is never applied to a model that violates its assumptions. Since such violation can originate from the updating of the model by another analysis, analyses must be executed in the correct order. Another challenge is to do this soundly and scalably over models of realistic complexity and diverse set of analyses. To address these challenges, we develop an analysis integration approach that uses *contracts* to specify dependencies between analyses, determine their correct orders of application, and specify and verify applicability conditions across multiple domains. We present an implementation of our approach, and demonstrate its effectiveness, extensibility, and scalability.

General Terms

Verification, Design, Theory

Keywords

Cyber-physical systems, analysis, real-time scheduling, thermal runaway, model checking, battery scheduling, analysis contracts, virtual integration

1. INTRODUCTION

The development of today’s industrial-scale cyber-physical systems (CPS) is heavily driven by models [15] and analyses. This trend is expected to strengthen, since it enables CPSs to be developed, upgraded, integrated, and verified virtually through models before manufacturing starts. Modeling also supports collaborative development by different teams, and fosters early error detection, faster development, and lower costs. In particular, analyses enable model creation and verification at design time to guarantee important quality attributes, such as control stability, schedulability, power consumption, safety, and security.

These analyses emerge from different scientific domains, such as timing, logical correctness, and thermal resilience. Consequently, they focus on different CPS abstractions that interact in subtle ways. This leads to two problems that render analysis results untrustworthy: (i) one analysis modifies a system model in a way that violates the assumptions made by another, e.g., a real-time task-allocation algorithm [7] assigns a set of threads to a processor scheduled via a dynamic

priority algorithm, thus violating the fixed priority assumption made by a model checker [3]; (ii) the specification of such assumptions and the detection of their violation are left implicit in the hands of human designers that, more often than not, are unable to cope with their complexity and subtlety due to the inter-domain nature of the assumptions.

As a consequence, these problems are currently discovered very late during system integration, leading to costly fixes [25]. This issue is particularly prevalent in multi-tier industries, such as avionics and automotive, where systems are integrated from independently-developed parts, designs of which are analyzed with a mishmash of tools. It is currently mitigated in an ad-hoc and manual way, which is neither scalable, nor able to provide a high degree of assurance [14].

In this paper, we present and evaluate an alternative solution that is mathematically rigorous and automated. Our approach consists of two parts: (i) *contract* specifications for analyses in a language with well-defined syntax and semantics. A contract for an analysis expresses both the *assumption* under which it produces a sound result, and the *guarantee* about the resulting modified model. (ii) An algorithm that computes inter-dependencies between analysis contracts, ensures that they are consistent, and executes them in an order guaranteeing soundness of all analysis results. More specifically, we make the following contributions.

First, we present a new language for specifying analysis contracts, and define its syntax and semantics. Our language finds a balance between expressiveness and decidability to support contracts that capture both design-time and runtime system aspects. An example is the frequency scaling analysis [23], which assumes that the runtime scheduling of threads allocated at design-time to the same CPU is behaviorally equivalent to deadline monotonic. To this end, our contract language combines a many-sorted first-order logic with a variant of linear temporal logic, propositions of which are derived from predicates over the system state. The logical nature of the language also makes it applicable to multiple analysis domains. In addition, validity of contracts expressed in our language is efficiently decidable.

Second, we present a contract verification algorithm that takes a set of analysis contracts and an architectural CPS model as input, computes the dependencies between the contracts, and uses this dependency to execute the analyses over the model in an order guaranteed to produce sound results. During this process, our algorithm checks the validity of each contract as follows: (i) before an analysis is executed, the validity of its assumption is checked over the input model; and (ii) after the analysis completes, the validity of its guarantee is ensured over the output model. Moreover, the validity is checked in a sound and exhaustive manner via co-operative application of an SMT solver (to the design-time aspect) and a model checker (to the runtime aspect) of the contract.

Finally, we implement our approach and demonstrate its effectiveness, extensibility, and scalability with a case study that involves multiple analyses from two domains – real-

time thread scheduling and battery cell scheduling. Our implementation is based on OSATE [20], which enables us to handle CPS models described in the Architecture Analysis and Design Language (AADL) [8], an SAE standard. It also uses state-of-the-art tools Z3 [6] and Spin [13] for SMT solving and model checking, respectively. We show how our approach detects subtle bugs in interactions between analyses, and how the use of advanced tools enables it to scale to systems of realistic complexity.

The rest of the paper is organized as follows. Sec. 2 introduces a running example of a system and analyses used in its design. Sec. 3 presents the analysis contract language. Sec. 4 presents our contract verification algorithm. Sec. 5 and Sec. 6 present our implementation and evaluation, respectively. We wrap up by surveying related work in Sec. 7 and concluding in Sec. 8.

2. RUNNING EXAMPLE

Consider a reconnaissance aircraft that must satisfy certain security, power, thermal, timing, and logical requirements. It is controlled by a set of threads (a.k.a. tasks) with different security levels executing on several processors. Each thread executes an infinite sequence of jobs triggered periodically. A job is a finite computation, e.g., a control correction to aircraft stability. The system has dynamic multi-cell batteries with configurable connections between cells so that some cells recharge while others are discharging [12]. Thread scheduling policies, processor frequencies and voltages, battery cell scheduling policies, and allocation of threads to processors are adjustable parameters that different analyses set and use to verify the system’s requirements.

The aircraft system has to satisfy several requirements: (i) from the security viewpoint, parts of the system with sensitive information must be protected from less sensitive parts, therefore threads with different security levels should not execute on the same processor; (ii) from the timing perspective, all jobs must meet deadlines required by the control algorithms (e.g., for flight stability); (iii) from the electrical domain, processors must draw minimal power, switching to the minimal frequency possible, thereby maximizing battery lifetime; (iv) logically, threads must be free of deadlocks and race conditions; (v) and from the thermal domain, even if a battery cell overheats, it should not trigger a chain reaction called *thermal runaway* [4].

All of these requirements can be satisfied individually with different analysis tools. However, ensuring that they are all satisfied together is challenging, since one tool can select some design parameters that violate the assumptions of another, as discussed before. Similarly, creating a single overarching analysis that takes into account all requirements is intractable. Our approach uses each analysis as is, but complements them with analysis contracts that allow us to model their interdependencies and ensure their sound application. This enables sound and tractable verification of all the system requirements.

As an example for the rest of the paper, we consider the following analyses: (i) thread co-location security: computes permissible thread co-locations based on security classes; (ii) bin packing [7]: assigns threads on processors to ensure schedulability; (iii) frequency scaling: minimizes the processor frequency given the threads assigned to processors; (iv) REK model checking [3]: checks if threads satisfy

user-specified safety properties; (v) thermal runaway: determines patterns of battery cell connections that lead to thermal runaway; and (vi) battery scheduling: determines a battery scheduler given the required operation time and battery size.

Arbitrary independent use of these analysis can lead to unsound results. For example, running the bin packing analysis before the co-location security analysis may violate thread co-location constraints set by the latter but not present at the time we run the former. Similarly, using the frequency scaling algorithm that assumes a deadline-monotonic scheduler on a processor that uses the earliest deadline first scheduler may produce unsound results. In this paper we demonstrate how our approach enables systematic integration of these analyses ensuring satisfaction of interdependencies and assumptions of analyses.

3. CONTRACT SPECIFICATION

In this section we present our contract language, and use it to formalize contracts for the analyses from Sec. 2.

3.1 Analysis Domain

We formalize a analysis domain as a “signature” that “externalizes” the domain’s concepts.

DEFINITION 1. An analysis domain σ is a many-sorted signature: $(\mathcal{A}, \mathcal{S}, \mathcal{R}, \mathcal{T}, [\cdot]_\sigma)$ where:

- \mathcal{A} is a set of sorts: $\mathcal{A} = \{A_i\}$.

Examples of sorts are Booleans (\mathcal{B}), integers (\mathcal{Z}), threads, processors, and scheduling protocols.

- \mathcal{S} is a set of static properties: $\mathcal{S} = \{S_i\}$. Each static property is a typed function $S_i : A_1 \times \dots \times A_j \mapsto A_k$.

Static properties capture system invariants set at design-time, e.g., standard operators like addition ($+ : \mathcal{Z} \times \mathcal{Z} \mapsto \mathcal{Z}$), conjunction ($\wedge : \mathcal{B} \times \mathcal{B} \mapsto \mathcal{B}$), and analysis-specific properties, such as thread periods and processor frequencies. Note that relations, like $<$, are functions with a Boolean range.

- \mathcal{R} is a finite set of runtime properties: $\mathcal{R} = \{R_i\}$, $R_i : A_1 \times \dots \times A_j \mapsto A_k$.

Runtime properties capture system aspects that depend on \mathcal{S} but change during execution, such as preemption between threads or connectivity between battery cells. Given a runtime state q , $q(R_i)$ means the evaluation of R_i in q .

- \mathcal{T} is the execution semantics of the domain.

Informally, \mathcal{T} is a set of infinite sequences of assignments to the runtime properties. The idea is that each sequence corresponds to an execution of the system, and specifically to the values of the runtime properties observed at the successive states of the execution. The executions are infinite since we are interested in analyzing reactive systems which, in general, do not terminate. \mathcal{T} is discussed in detail for specific domains in Sec. 4.2.3 and 4.2.4.

- $[\cdot]_\sigma$ is a partial interpretation of \mathcal{A} and \mathcal{S} . It assigns permissible values to some of the sorts in \mathcal{A} and value mappings to some elements of \mathcal{S} .

Intuitively, $\llbracket \cdot \rrbracket_\sigma$ interprets common sorts and operators in the standard way. For example $\llbracket \mathcal{B} \rrbracket = \mathbb{B}$, $\llbracket \mathcal{Z} \rrbracket = \mathbb{Z}$, $\llbracket + \rrbracket$ is integer addition, and $\llbracket \wedge \rrbracket$ is Boolean conjunction, and so on. In addition, $\llbracket \cdot \rrbracket_\sigma$ interprets domain-specific sorts, e.g., the set of permissible scheduling protocols.

Note that σ represents an analytic domain for the purpose of specifying its *external* contract, but does not limit *internal* reasoning of the analysis about the domain. Also, $\llbracket \cdot \rrbracket_\sigma$ leaves some elements of \mathcal{A} and \mathcal{S} (e.g., set of threads), and \mathcal{T} , uninterpreted. This is completed by the *architectural model* (or model, for short).

DEFINITION 2. An architectural model M is an interpretation $\llbracket \cdot \rrbracket_M$ of some elements of \mathcal{A} and \mathcal{S} , and \mathcal{T} .

For example, a set of three threads and their periods is specified by a model M as $\llbracket T \rrbracket_M = \{t_1, t_2, t_3\}$ and $\llbracket \text{Per} \rrbracket_M = \{t_1 \mapsto 40, t_2 \mapsto 50, t_3 \mapsto 60\}$. The value of runtime properties depends on the execution state, as follows. Let $\llbracket \cdot \rrbracket$ denote the combination of $\llbracket \cdot \rrbracket_\sigma$ and $\llbracket \cdot \rrbracket_M$. Formally, state q maps each runtime property $R_i : A_1 \times \dots \times A_j \mapsto A_k$ to a function $q(R_i) : \llbracket A_1 \rrbracket \times \dots \times \llbracket A_j \rrbracket \mapsto \llbracket A_k \rrbracket$. Let Q be the set of all states, and Q^ω be the set of all infinite sequences of states (i.e., executions). Then $\llbracket \mathcal{T} \rrbracket_M \subseteq Q^\omega$. In other words, $\llbracket \mathcal{T} \rrbracket_M$ is set of executions of the system defined by M . Note that \mathcal{R} is interpreted indirectly and modally via $\llbracket \mathcal{T} \rrbracket_M$. Specifically, each state q in each execution in $\llbracket \mathcal{T} \rrbracket_M$ gives an interpretation $q(R_i)$ to each $R_i \in \mathcal{R}$.

We require that $\llbracket \cdot \rrbracket$ interpret \mathcal{A} , \mathcal{S} , and \mathcal{T} completely and unambiguously. In general, analyses from multiple domains are applied to a single model M . Therefore, M must complete the interpretations for \mathcal{A} , \mathcal{S} , and \mathcal{T} from each such domain.

3.2 Contract Language

Our contract language consists of a combination of first-order and temporal logic formulas over the sorts and properties of a target domain. We present it in stages, beginning with the syntax of contract formulas.

3.2.1 Contract Formula Syntax

We first define the static fragment of contract formulas. Consider a domain $\sigma = (\mathcal{A}, \mathcal{S}, \mathcal{R}, \mathcal{T}, \llbracket \cdot \rrbracket_\sigma)$. Let V be a denumerable set of typed variables, $V = \{v_i\}, v_i : A_i$. Then the set of static formulas over σ , denoted ϕ , is defined by the following grammar:

$$\phi ::= v \mid f(e_1 \dots e_j),$$

where $v \in V, f \in \mathcal{S}$ with arity j , and $e_1, \dots, e_j \in \phi$.

Typing. Formulas are well-typed. Variables are typed with their sorts. If the type of f is $A_1 \times \dots \times A_j \mapsto A_k$, then the type of $f(e_1, \dots, e_j)$ is A_k and it is well-typed iff the type of each e_i is A_i . For simplicity, we use write $v_1 + v_2$ instead of $+(v_1, v_2)$, $e_1 \wedge e_2$ instead of $\wedge(e_1, e_2)$, etc.

Temporal Logic Formulas. We now turn our attention to the temporal fragment of contract formulas. Since we are interested in expressing properties of infinite executions of reactive systems, we use a variant of next-time-free linear temporal logic (LTL) [21]. The key difference with standard LTL is that propositions in our logic are not *atomic*, but are formulas constructed from runtime and static properties. The set of all such runtime formulas is defined by the following grammar:

$$RF ::= v \mid f(e_1, \dots, e_j),$$

where $v \in V, f \in \mathcal{S} \cup \mathcal{R}$ with arity j , and $e_1, \dots, e_j \in \phi$. Runtime formulas are also well-typed in the same way as static formulas. Then our LTL formulas, denoted ψ , are defined by the following grammar:

$$\psi ::= p \mid \neg\psi \mid \psi \wedge \psi \mid \psi \cup \psi,$$

where $p \in RF$ is a runtime formula with Boolean type. Note that other temporal operators, such as G and F , are defined in terms of \neg, \wedge , and \cup in the standard manner. Finally, we define contract formulas using ϕ and ψ .

DEFINITION 3. Given a domain σ , a set of contract formulas \mathcal{F}_σ is defined by the following grammar:

$$\begin{aligned} \mathcal{F}_\sigma ::= & \forall v_1 \dots v_j . \phi \mid \exists v_1 \dots v_j . \phi \\ & \mid \forall v_1 \dots v_j . \phi : \psi \mid \exists v_1 \dots v_j . \phi : \psi, \end{aligned}$$

where ϕ is a static formula of Boolean type, and ψ is a LTL formula – both over variables v_1, \dots, v_j . Thus, a contract formula can be purely first-order (i.e., the first two forms) or a combination of first-order and LTL (i.e., the last two forms). The meaning of the third form is that every assignment of V that satisfies ϕ must also satisfy ψ . The meaning of the last form is that at least one assignment of V satisfies both ϕ and ψ . Note that formulas without variables are also allowed. We are now ready to define analyses and their contracts.

3.2.2 Analysis Contracts

Functionally, an analysis An takes an architectural model M_I as input and produces a new architectural model $M_O = An(M_I)$ as output. The contract for An specifies restrictions on valid input models and valid output models, as well as the model properties its reads and modifies.

DEFINITION 4. A contract C for an analysis An over a domain σ is a 4-tuple $C = (I, O, A, G)$, where:

- $I \subseteq \mathcal{A} \cup \mathcal{S}$ are sorts and properties read by An .
- $O \subseteq \mathcal{A} \cup \mathcal{S}$ are sorts and properties modified by An .
- $A \subseteq \mathcal{F}_\sigma$ are assumptions – contract formulas that must be satisfied by every valid input model to An .
- $G \subseteq \mathcal{F}_\sigma$ are guarantees – contract formulas over that must be satisfied by every valid output model from An .

In order to determine whether An satisfies its contract, we must first define what it means for An to satisfy a contract formula. This is the topic of the next subsection.

3.2.3 Contract Semantics

We start with the evaluation of static formulas and build up to defining satisfaction of a contract. Let the interpretation of sort A be $\llbracket A \rrbracket$ and the interpretation of $f : A_1 \times \dots \times A_j \mapsto A_k \in \mathcal{S}$ be $\llbracket f \rrbracket : \llbracket A_1 \rrbracket \times \dots \times \llbracket A_j \rrbracket \mapsto \llbracket A_k \rrbracket$. An *assignment* μ maps each variable $v : A$ to an element of $\llbracket A \rrbracket$. Given a static formula ϕ , and an assignment μ , $\llbracket \phi, \mu \rrbracket$ is the evaluation of ϕ under μ defined as:

$$\begin{aligned} \llbracket v, \mu \rrbracket &= \mu(v) \\ \llbracket f(e_1, \dots, e_j), \mu \rrbracket &= \llbracket f \rrbracket(\llbracket e_1, \mu \rrbracket, \dots, \llbracket e_j, \mu \rrbracket) \end{aligned}$$

Note that if the type of ϕ is A , then $\llbracket \phi, \mu \rrbracket \in \llbracket A \rrbracket$. The evaluation of a runtime formula RF under an assignment μ in a state q , denoted $\llbracket RF, \mu, q \rrbracket$, is defined as:

$$\begin{aligned} \llbracket v, \mu, q \rrbracket &= \mu(v) \\ \llbracket f(e_1, \dots, e_j), \mu, q \rrbracket &= \llbracket f \rrbracket(\llbracket e_1, \mu, q \rrbracket, \dots, \llbracket e_j, \mu, q \rrbracket), f \in \mathcal{S} \\ \llbracket f(e_1, \dots, e_j), \mu, q \rrbracket &= q(\llbracket f \rrbracket)(\llbracket e_1, \mu, q \rrbracket, \dots, \llbracket e_j, \mu, q \rrbracket), f \in \mathcal{R} \end{aligned}$$

An execution $\pi = q_0, q_1, \dots$ satisfies an LTL formula ψ under assignment μ , denoted $\pi, \mu \models \psi$, if:

- If $\psi \in RF$ then $\pi, \mu \models \psi$ iff $\llbracket \psi, \mu, q_0 \rrbracket = \top$.
- $\pi, \mu \models \neg \psi'$ iff $\pi, \mu \not\models \psi'$.
- $\pi, \mu \models \psi_1 \wedge \psi_2$ iff $\pi, \mu \models \psi_1$ and $\pi, \mu \models \psi_2$.
- $\pi, \mu \models \psi_1 \cup \psi_2$ iff there exists $i \geq 0$ such that for all $0 \leq j < i$, $\pi^j, \mu \models \psi_1$ and $\pi^i, \mu \models \psi_2$, where π^x is the (infinite) suffix of π starting with state q_x .

Consider a model M . Let $\llbracket \cdot \rrbracket$ be the combination of $\llbracket \cdot \rrbracket_\sigma$ and $\llbracket \cdot \rrbracket_M$. A contract formula $f \in \mathcal{F}_\sigma$ is satisfied by M , denoted $M \models f$, based on the form of f (see Def. 3) as follows:

- f is of the first form and $\forall \mu \in \mathcal{V}. \llbracket f, \mu \rrbracket = \top$.
- f is of the second form and $\exists \mu \in \mathcal{V}. \llbracket f, \mu \rrbracket = \top$.
- f is of the third form and $\forall \mu \in \mathcal{V}. \forall \pi \in \llbracket \mathcal{T} \rrbracket. \llbracket \phi, \mu \rrbracket = \top \Rightarrow \pi, \mu \models \psi$.
- f is of the fourth form and $\exists \mu \in \mathcal{V}. \forall \pi \in \llbracket \mathcal{T} \rrbracket. \llbracket \phi, \mu \rrbracket = \top \wedge \pi, \mu \models \psi$.

Thus, first-order quantification over static formulas is interpreted in a natural manner. Universal quantification over a static formula and an LTL formula holds if and only if all assignments that satisfy the static formula also satisfy the LTL formula. Existential quantification over a static formula and an LTL formula holds if and only if there exists an assignment that satisfies both static and LTL formulas.

Analysis Applicability. A model to which an analysis is applied should satisfy its assumption, and the resulting model should satisfy its guarantee. Formally:

DEFINITION 5. *A model M is applicable to an analysis An with contract $An.C = (I, O, A, G)$ – denoted $M \models C$ – iff $\forall a \in A. M \models a$ and $\forall g \in G. An(M) \models g$.*

In Sec. 4 we present an algorithm to decide $M \models C$. Now, to highlight our approach, we turn to formalizing analysis domains and contracts for the example from Sec. 2.

3.3 Scheduling Domain

The scheduling domain σ_{Sched} encodes the semantics of the real-time scheduling of threads along with their allocation to processors. Analyses in this domain decide valid thread allocations, priority assignments, check schedulability according to a selected scheduling policy, determine processor frequency, etc. The domain is defined as $\sigma_{Sched} = (\mathcal{A}_{Sched}, \mathcal{S}_{Sched}, \mathcal{R}_{Sched}, \mathcal{T}_{Sched}, \llbracket \cdot \rrbracket_{Sched})$. In addition to Booleans \mathcal{B} and integers \mathcal{Z} , \mathcal{A}_{Sched} has reals \mathcal{R} , threads T , CPUs C , thread security classes SecCl , and thread scheduling policies SchedPol . Sorts $\mathcal{B}, \mathcal{Z}, \mathcal{R}$ are interpreted by $\llbracket \cdot \rrbracket_{Sched}$ in a standard way. There are three security classes – normal, secret, and top-secret – $\llbracket \text{SecCl} \rrbracket_{Sched} =$

Name	Type	Description
Per	$T \mapsto \mathcal{Z}$	Thread's period.
Dline	$T \mapsto \mathcal{Z}$	Thread's deadline.
WCET	$T \mapsto \mathcal{Z}$	Thread's worst case execution time.
ThSecCl	$T \mapsto \text{SecCl}$	Thread's security class.
CPUSchedPol	$C \mapsto \text{SchedPol}$	CPU's scheduling policy.
CPUFreq	$C \mapsto \mathcal{R}$	CPU's normalized frequency ¹ .
NotColoc	$T \mapsto 2^T$	Thread t mapped to a set of threads that should not share the same CPU as t .
CPUBind	$T \mapsto C$	Thread-to-CPU binding.
ThSafe	$C \mapsto \mathcal{B}$	Flag whether CPU's threads are thread-safe.
Voltage	$() \mapsto \mathcal{R}$	Required system voltage ² .

Table 1: Static properties of σ_{Sched} .

$\{\text{normal, secret, topsecret}\}$ and three scheduling policies – rate monotonic scheduling (RMS), earliest deadline first (EDF) [17], and deadline monotonic scheduling (DMS) [1] – $\llbracket \text{SchedPol} \rrbracket_{Sched} = \{\text{rms, dms, edf}\}$. Sorts T and C are uninterpreted since they are model dependent.

In addition to Boolean, integer, and real arithmetic operators, the properties in \mathcal{S}_{Sched} interpreted by M are summarized in Tab. 1. Also, there is one runtime property $\mathcal{R}_{Sched} = \{\text{CanPrmpt} : T \times T \mapsto \mathcal{B}\}$ such that $q(\text{CanPrmpt}(t_1, t_2)) = \top$ iff t_1 can preempt t_2 in state q . Interpretation of \mathcal{T}_{Sched} and CanPrmpt is model dependent and presented in Sec. 4.

3.4 Battery Domain

The domain of battery design and usage $\sigma_{Batt} = (\mathcal{A}_{Batt}, \mathcal{S}_{Batt}, \mathcal{R}_{Batt}, \mathcal{T}_{Batt}, \llbracket \cdot \rrbracket_{Batt})$ is defined as follows. Sorts $\mathcal{B}, \mathcal{Z}, \mathcal{R} \in \mathcal{A}_{Batt}$ and their interpretations are identical to σ_{Sched} . The set of batteries is $B \in \mathcal{A}_{Batt}$, and $\text{ConnSchedPol} \in \mathcal{A}_{Batt}$ is the set of three battery scheduling policies [12][16]: unweighed round robin with fixed cell groups (**FGuRR**), weighed kRR with fixed parallel cell groups (**FGwRR**), and weighed kRR with cell group packing (**GPwRR**). $\text{Voltage} \in \mathcal{S}_{Batt}$ is the same as in the scheduling domain: our approach naturally captures the fact that both domains are concerned with the system voltage. \mathcal{S}_{Batt} contains the properties summarized in Tab. 2.

Informally, the battery execution consists of continuous charging, discharging, and resting of cells. The precise semantics $\llbracket \mathcal{T}_{Batt} \rrbracket$ is model-dependent and defined in Sec. 4. There is one runtime property: $\mathcal{R}_{Batt} = \{\text{TN} : \mathcal{Z} \rightarrow \mathcal{Z}\}$. Informally, when the system is in state q , $q(\text{TN}(i))$ denotes the number of cells with i thermal neighbors – cells that exchange heat conductively through a connector³. This is motivated by results [10]: there is a close connection between thermal neighbors and thermal runaway. Specifically, there exist constants $K(i) : i \in \mathbb{Z}$ such that a state q triggers a thermal runaway if it violates the condition:

$$\sum_i K(i) \times q(\text{TN}(i)) \geq 0 \quad (1)$$

However, the exact values of K are not known independently of a model, and experiments with a battery are

¹A real number between 0 and 1.

²**Voltage** is a nullary function, or a Boolean constant. We consider a simplified example where the system voltage is the maximum of required individual processor voltages.

³As opposed to electrical neighbors – cells that are connected to each other from the electrical viewpoint, no matter how far apart physically they are.

Name	Type	Description
Voltage	$() \mapsto \mathcal{R}$	Required system voltage.
BatRows	$B \mapsto \mathcal{Z}$	Battery's cell rows.
BatCols	$B \mapsto \mathcal{Z}$	Battery's cell columns.
BatConnSchedPol	$B \mapsto$ ConnSchedPol	Battery's cell scheduling policy.
SerialReq	$B \mapsto \mathcal{Z}$	Number of cells required to connect in series ⁴ .
ParalReq	$B \mapsto \mathcal{Z}$	Number of cells required to connect in parallel ⁵ .
K	$\mathcal{Z} \mapsto \mathcal{Z}$	Weight of cells with i thermal neighbors.
HasReqdLifetime	$B \mapsto \mathcal{B}$	Flag whether a battery has the lifetime required.

Table 2: Static properties of σ_{Batt} .

needed to estimate K .

3.5 Analysis Contracts for Running Example

Using the domains signatures σ_{Sched} and σ_{Batt} we define the contracts for analyses \mathcal{AN} from Sec. 2.

Secure thread allocation⁶ ($An_{SecAlloc}$) has contract $C_{SecAlloc} : I = \{T, ThSecCl\}$, $O = \{NotColoc\}$, $A = \emptyset$, and $G = \{g\}$ where g is:

$$\forall t_1, t_2. ThSecCl(t_1) \neq ThSecCl(t_2) \Rightarrow t_1 \in NotColoc(t_2)$$

Thus, $An_{SecAlloc}$ makes no assumptions, but guarantees that threads with different security classes are never co-located. We omit sorts, e.g., $t_1 : T$ and $t_2 : T$, since they are implied by typing rules.

Bin packing ($An_{BinPack}$) has contract $C_{BinPack} : I = \{T, C, NotColoc, Per, WCET, Dline\}$, $O = \{CPUBind\}$, $A = \emptyset$, and $G = \{g\}$ where g is:

$$\forall t_1, t_2. t_1 \in NotColoc(t_2) \Rightarrow CPUBind(t_1) \neq CPUBind(t_2)$$

Thus, $An_{BinPack}$ makes no assumptions but guarantees that threads that should not be co-located are never scheduled on the same CPU.

Frequency scaling (An_{FreqSc}) has contract $C_{FreqSc} : I = \{T, C, CPUBind, Dline\}$, $O = \{CPUFreq\}$, $G = \emptyset$, and

$$A \triangleq \{\forall t_1, t_2. CPUBind(t_1) = CPUBind(t_2) : \\ G (CanPrmpt(t_1, t_2) \Rightarrow Dline(t_1) < Dline(t_2))\}$$

Thus, An_{FreqSc} makes no guarantees but assumes that the scheduling used is semantically equivalent to a deadline-monotonic scheduling policy. Note that a scheduling policy can be DMS for a specific model, e.g., rate-monotonic scheduling (RMS) for a harmonic model, even though it is not deadline monotonic in general.

Model checking with REK (An_{REK}) has contract $C_{REK} : I = \{T, C, Per, Dline, WCET, CPUBind\}$, $O = \{ThSafe\}$, $G = \emptyset$, and $A = \{a_1, a_2\}$ where:

$$a_1 \triangleq \forall t. Per(t) = Dline(t),$$

$$a_2 \triangleq \forall t_1, t_2. G (CanPrmpt(t_1, t_2) \Rightarrow G \neg CanPrmpt(t_2, t_1)).$$

REK [3] takes threads and their marked source code files (which we didn't include into the formal example) as input

⁴SerialReq is a battery-specific form of the voltage output requirement.

⁵ParalReq is a battery-specific form of the electrical current output requirement.

⁶We omit the implied atoms in inputs of analyses, like SecCl in $C_{SecAlloc}.I$.

and verifies whether the system is safe, where safety is expressed as assertions embedded in the source code. An_{REK} assumes implicit deadlines and fixed-priority scheduling. Prior to this work, the only way to apply REK was to use RMS. However, our contract mechanism allows for a broader scope of applicability. Note that a_1 expresses implicit deadlines, while a_2 expresses fixed priority scheduling, i.e., if t_1 preempts t_2 , then t_2 should never be able to preempt t_1 .

Thermal runaway ($An_{ThermRun}$) has contract $C_{ThermRun} : I = \{B, BatRows, BatCols, Voltage\}$, $O = \{K\}$, $A = \emptyset$, and $G = \emptyset$. Note that $An_{ThermRun}$ has no assumptions or guarantees, but has dependency with defined below battery scheduling via I and O . It determines the patterns, which, given concrete battery characteristics, would result into a thermal runaway. In our example, we encode these patterns as $K(i)$ for $i : \mathcal{Z} \in [0, 3]$. $An_{ThermRun}$ determines K through experimentation, adjusting K so that acceptable heat propagation patterns satisfy (1), and unacceptable ones violate it.

Battery scheduling ($An_{BatSched}$) has contract $C_{BatSched} : I = \{B, BatRows, BatCols\}$, $O = \{BatConnSchedPol, HasReqdLifetime, SerialReq, ParalReq\}$, $A = \emptyset$, and $G = \{g\}$ where g is:

$$G (K(0) \times TN(0) + K(1) \times TN(1) + \\ K(2) \times TN(2) + K(3) \times TN(3) \geq 0).$$

$An_{BatSched}$ computes a battery cell connectivity scheduler that maximizes the battery lifetime given the battery characteristics and output requirements. It sets a flag indicating whether the battery with the selected scheduler meets the lifetime requirement. Since the scheduling is not aware of the thermal runaway, the determined scheduler needs to be verified against the thermal runaway pattern, hence the guarantee. $An_{BatSched}$ also sets cell group characteristics SerialReq and ParalReq that are used to verify its guarantee.

4. CONTRACT VERIFICATION

This section presents our contract verification algorithm, which takes a model M and a set of analyses \mathcal{AN} and produces a correct execution of \mathcal{AN} on M , or aborts. The algorithm consists of the following steps: (i) determine an ordering \mathcal{O} of \mathcal{AN} that respects all inter-analysis dependencies; (ii) process each analysis $An \in \mathcal{AN}$ with contract $C = (I, O, A, G)$ in the order \mathcal{O} by: (iia) verifying $\forall a \in A. M \models a$, (iib) updating M by executing An , i.e., setting M to $An(M)$, (iic) and verifying $\forall g \in G. M \models g$; and (iii) output the final M as the result. The algorithm aborts if either an appropriate ordering \mathcal{O} cannot be computed, if any of the verifications in Steps (iia) and (iic) fail, or the execution of An or M fails. Note that the algorithm ensures that all analyses produce valid results since: (a) an analysis An executes successfully on M only if $M \models An.C$ (see Def. 5); and (b) the ordering \mathcal{O} ensures that once an analysis has been executed, future values of M do not violate its assumptions.

4.1 Analysis Ordering

We begin with Step (i), which uses contracts to determine a correct ordering of analyses execution. The contract of analysis An is denoted $C(An)$. For a contract $C = (I, O, A, G)$, $C.I$ means I , $C.O$ means O , etc.

DEFINITION 6. Analysis contract C_i is dependent on

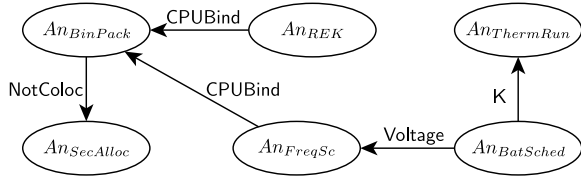


Figure 1: Analyses graph for the running example.

analysis contract C_j , denoted $d(C_i, C_j)$, if $C_i.I \cap C_j.O \neq \emptyset$. Analysis An_i is dependent on analysis An_j , denoted $d(An_i, An_j)$, iff $d(C(An_i), C(An_j))$.

Given a set of analyses \mathcal{AN} , an ordering $\mathcal{O} = \langle An_1 \dots An_n \rangle$ of \mathcal{AN} is *sound* if each analysis in the ordering is not dependent on any of its predecessors, i.e.,

$$\forall i \in [1, n]. \forall j \in [1, i]. \neg d(An_i, An_j)$$

Thus, the goal of step (i) is to produce a sound ordering of \mathcal{AN} . Consider the directed graph of analyses $\gamma = (\mathcal{AN}, d(\cdot, \cdot))$. It can be shown that: (i) if γ is cyclic then there is no sound ordering of \mathcal{AN} ; and (ii) otherwise γ is a DAG and any topologically sorted ordering of its nodes is a sound ordering of \mathcal{AN} . Therefore, \mathcal{O} is computed by: (i) constructing γ ; (ii) checking its cyclicity; (iii) aborting if γ is cyclic; and (iv) if it is acyclic, constructing any topological ordering of its nodes. The specific \mathcal{O} constructed does not affect the correctness of our algorithm since topological orderings of γ differ only in relative positions of mutually independent analyses. The γ for our example is shown in Fig. 1. Each edge is labeled by a property that causes the dependency between the corresponding analyses.

4.2 Analysis Applicability

We now focus on Steps (iia) and (iic) of our algorithm. Recall that the core problem here is to decide $M \models f$ where f is a contract formula. In Step (iia) f is one of the contract’s assumptions, while the Step (iic) it is one of the contract’s guarantees. The algorithm for verifying $M \models f$ depends on the form of f (see Def. 3), and we consider each separately.

4.2.1 Verifying Purely First Order Formulas

If f is a quantified first-order formula (i.e., the first two forms of Def. 3) we check $M \models f$ via Satisfiability Modulo Theories (SMT) solving. We first describe how to construct a SMT formula $\varphi(M, \phi)$ – using SMT v2 syntax – given a model M and a first-order formula ϕ . The first step in constructing $\varphi(M, \phi)$ is defining the basic types. These are obtained directly from the sorts (\mathcal{A}) of the analysis domain. Basic sorts like Boolean, integer, and float are already primitive types in SMT. Domain-specific sorts, like threads and processors, are declared as integers using the `define-sort` SMT command. Subsequently, the IDs of the actual threads and processors are used as concrete values for their corresponding types. For example, if M has five threads with IDs $[0, 4]$ and three processors with IDs $[0, 2]$, then $\varphi(M, \phi)$ has two types – `thread` and `processor` – defined as follows:

```
(define-sort thread() (Int))
(define-sort processor() (Int))
```

Subsequently, all variables of type T have five legal values – $[0, 4]$ – and all variables of type C have three legal values

```
1 (define-sort thread () (Int))
2 (define-sort processor () (Int))
3 (declare-fun Actual_Processor_Binding (thread) Int)
4 (define-fun Not_Colocated((x1 thread) (x2 thread)) Bool
5   (ite (and (= x1 0)(= x2 1)) true (ite (and (= x1 0)(= x2 2)) true
6     (ite (and (= x1 1)(= x2 0)) true (ite (and (= x1 2)(= x2 0)) true   false))))))
7 (assert (= (Actual_Processor_Binding 0) 0))
8 (assert (= (Actual_Processor_Binding 1) 1))
9 (assert (= (Actual_Processor_Binding 2) 1))
10 (assert (not (forall ((x1 thread) (x2 thread)) (=
11   (and (or (= x1 0) (= x1 1) (= x1 2))
12     (or (= x2 0) (= x2 1) (= x2 2))))
13   (= (and (not (= x1 x2)) (Not_Colocated x1 x2))
14     (not (= (Actual_Processor_Binding x1) (Actual_Processor_Binding x2)))))))
15 (check-sat)
```

Figure 2: SMT problem for verifying $C_{BinPack}.G$.

– $[0, 2]$. These legal values are enforced by adding appropriate SMT constraints to $\varphi(M, \phi)$. The functions in $\varphi(M, \phi)$ are obtained from the static properties of the domain that are used in the contract. For example, if the analysis references `Per`, then $\varphi(M, \phi)$ contains a function `Period` defined as follows:

```
(declare-fun Period (thread) Int)
```

Variables are also declared with appropriate types, e.g., variable $t_1 : T$ is declared as: `(declare-fun t1 () (thread))`. Finally, the constraints in $\varphi(M, \phi)$ are obtained from:

- The interpretation of static properties of the domain. For example, if thread 0 has a period 20, then we add the constraint `(assert (= (period 0) 20))`.
- The formula ϕ itself. For example if $\phi \triangleq \text{Per}(t_1) < \text{Per}(t_2)$ then we add the constraint `(assert (< (period t1) (period t2)))`.

Given an SMT formula φ as input, an SMT solver returns SAT if φ is satisfiable, and UNSAT otherwise. Now checking $M \models f$ reduces to two cases:

- $f = \forall v_1 \dots v_j . \phi$ (form 1 of Def. 3): In this case we return YES if the SMT solver returns UNSAT for input $\varphi(M, \neg\phi)$ and NO otherwise.
- $f = \exists v_1 \dots v_j . \phi$ (form 2 of Def. 3): In this case we return YES if the SMT solver returns SAT for input $\varphi(M, \phi)$ and NO otherwise.

The correctness of our algorithm follows from our semantics and the construction of $\varphi(M, \phi)$. An example of $\varphi(M, \phi)$ constructed for $C_{BinPack}.G$ is shown in Fig. 2. The guarantee states that non-colocated threads should be bound to different processors, and has the first form of Def. 3. The SMT solver returns UNSAT, hence $M \models C_{BinPack}.G$.

4.2.2 Verifying First Order+LTL Formulas

We now show how to verify f if it combines both first-order logic and LTL (i.e., the last two forms of Def. 3). This again has two cases:

- $f = \forall v_1 \dots v_j . \phi : \psi$ (form 3 of Def. 3): We first construct $\varphi(M, \phi)$. Next, we use the SMT solver iteratively to compute all satisfying solutions of $\varphi(M, \phi)$. To obtain all solutions, we use “blocking clauses”, i.e., once we obtain a solution, we add its negation to the formula before re-solving it. From each solution we construct the corresponding assignment μ to $\{v_1 \dots v_j\}$. For each such assignment μ , we check $\forall \pi \in \llbracket T \rrbracket . \pi, \mu \models \psi$ using a model checker. The model

checking step is domain-specific and described in the following subsections. We return YES if the model checker finds no ψ violations for every μ , and NO otherwise.

- $f = \exists v_1 \dots v_j . \phi : \psi$ (form 4 of Def. 3): We first construct $\varphi(\mathbf{M}, \phi)$. Next, we use the SMT solver iteratively to compute all assignments μ to $\{v_1 \dots v_j\}$ as in the previous case. For each assignment μ , we check $\forall \pi \in \llbracket \mathcal{T} \rrbracket . \pi, \mu \models \psi$ using a model checker. We return YES if the model checker find no ψ violations for at least one μ , and NO otherwise.

The correctness of our algorithm follows from our semantics and the construction of $\varphi(\mathbf{M}, \phi)$. Note that our algorithm always terminates if the sort of each quantified variable v_i is interpreted to a finite domain (e.g., threads and batteries, but not integers), since this means that there are only a finite set of assignments to $\{v_1 \dots v_j\}$. This is indeed the case for all analyses in our example. We now describe the model checking step for the scheduling and battery domains.

4.2.3 Model Checking for Scheduling

The execution semantics $\llbracket \mathcal{T}_{Sched} \rrbracket$ of the scheduling domain is defined as follows. Recall that each thread consists of an infinite and periodic sequence of jobs. A state q of the system corresponds to a point in time where a new job arrives or a currently executing job terminates. An execution consists of a infinite sequence of such states observed at runtime. Note that multiple executions are possible due to the non-determinism in the time required by each job to complete. Then $\llbracket \mathcal{T}_{Sched} \rrbracket$ consists of all such executions.

We model $\llbracket \mathcal{T}_{Sched} \rrbracket$ as a Kripke structure $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ composed of a “task” process for each thread. Task processes are periodic and their numeric characteristics – (Per, Dline, WCET) – are specified by the model \mathbf{M} . There are $\llbracket C \rrbracket_{\mathbf{M}}$ processors, and each running task is allocated to a processor dynamically. For each task process t , $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ has the following propositions:

- $\text{Prior}(t) : \mathcal{Z}$ – the priority of t .
- $\text{Run}(t) : \mathcal{B}$ – whether a job of t is dispatched on a processor.
- $\text{InQ}(t) : \mathcal{B}$ – whether a job of t has arrived but hasn’t been completed yet.

$\text{Prior}(t)$ is set by the scheduling policy and decides which tasks are executed. The last two propositions encode every possible state of t : idle if $\neg \text{InQ}(t) \wedge \neg \text{Run}(t)$, waiting for processor if $\text{InQ}(t) \wedge \neg \text{Run}(t)$, and executing if $\text{InQ}(t) \wedge \text{Run}(t)$. Also, for any state q of $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$, and threads t_1, t_2 , $q(\text{CanPrmpt})(t_1, t_2)$ is \top iff the following holds in q :

$$\text{Run}(t_1) \wedge \neg \text{Run}(t_2) \wedge \text{InQ}(t_2)$$

Recall that our model checking problem is $\forall \pi \in \llbracket \mathcal{T}_{Sched} \rrbracket . \pi, \mu \models \psi$, where μ is a variable assignment. We solve this by: (i) instantiating the LTL formula ψ to a propositional LTL formula ψ_{prop} using μ ; and (ii) using a model checker to verify $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket) \models \psi_{prop}$. For example, $C_{FreqSc.A}$ is expressed as the following propositional LTL formula:

$$\mathbf{G} (\text{Run}(\mu(t_1)) \wedge \neg \text{Run}(\mu(t_2)) \wedge \text{InQ}(\mu(t_2)) \Rightarrow \text{Dline}(\mu(t_0)) \leq \text{Dline}(\mu(t_1)))$$

The correctness of our algorithm follows from the semantics of our LTL formulas, the semantics of propositional LTL, and the correctness of model checking.

4.2.4 Model Checking For Battery

Let us now define the execution semantics of the battery domain $\llbracket \mathcal{T}_{Batt} \rrbracket$. A battery consists of a matrix of cells χ being continuously charged, discharged, connected and disconnected with each other. A state q of the system corresponds to a point in time when either the charge or the connectivity status of a cell changes. An execution consists of an infinite sequence of such states observed at runtime. Note that many such executions are possible due to the non-determinism in the time required by cells to charge and discharge. Then $\llbracket \mathcal{T}_{Therm} \rrbracket$ consists of all such executions.

We model $\llbracket \mathcal{T}_{Therm} \rrbracket$ as a Kripke structure $\mathcal{K}(\llbracket \mathcal{T}_{Therm} \rrbracket)$ with the following propositions for each cell $c = (x, y) \in \chi$, which is characterized by its physical coordinates $x \in [0.. \text{BatRows} - 1]$ and $y \in [0.. \text{BatCols} - 1]$:

- $\text{CellCharge}(c)$ is the charge of c . To simplify model checking we chose a Boolean abstraction for the cell charge, but other abstractions are possible too.
- $\text{CellSt}(c)$ is the status of c with possible values **discharging**, **charging**, and **idle**.
- $\text{Gr}(c)$ is the number of group of cells electrically connected in serial within which c is located. Groups are treated as electrically connected in parallel with each other. Every cell belongs to a group, but not every group or cell is discharging.

TN is encoded as follows. Cells c_1 and c_2 are thermal neighbors, denoted $\text{istnbr}(c_1, c_2)$, if: (i) $c_1 \neq c_2$; (ii) $\text{Gr}(c_1) = \text{Gr}(c_2)$; (iii) $|c_1.x - c_2.x| + |c_1.y - c_2.y| \leq \text{TNDIST}$ ⁷; (iv) $\text{CellCharge}(c_1) = \text{CellCharge}(c_2) = \top$; and (v) $\text{CellSt}(c_1) = \text{CellSt}(c_2) = \text{discharging}$. The number of thermal neighbors of cell c is $\text{ntnbr}(c) = |\{c' \in \chi . \text{istnbr}(c, c')\}|$. Finally, $\text{TN}(i) = |\{c' \in \chi . \text{ntnbr}(c) = i\}|$.

5. IMPLEMENTATION

In this section we present the implementation of our analysis contracts tool in OSATE – an open source environment for AADL modeling [8]. In the following we describe how our tool implements the concepts of this paper, and how Spin/Promela is used for model checking.

5.1 Contracts Framework

Let us start by describing how the different elements of our analysis contracts are represented. First, the architectural model \mathbf{M} is described in an AADL model. Secondly, the domain sorts, like SecCl and ConnSchedPol , are defined as AADL property types. Thirdly, thread T processor C sorts are derived from the corresponding threads and processors from the AADL model, and the device sort B is also derived from the corresponding AADL devices in the model. \mathcal{S}_{Sched} and \mathcal{S}_{Batt} are described using properties of the components. Finally, we specify the contracts as a sub-language *annex* in AADL to capture I , O , A , and G of analyses.

Fig. 3 depicts the architecture of our tool. Analysis contracts \mathcal{C} are associated with AADL component types, while

⁷For our calculations we use $\text{TNDIST} = 2$.

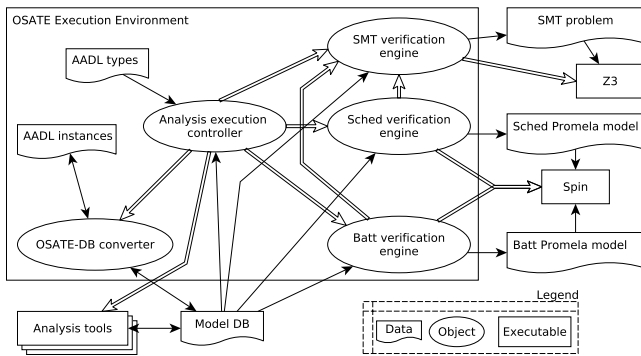


Figure 3: Contracts tool architecture. Solid arrow = data flow; hollow arrow = control flow.

M is derived from the AADL main system instance. Initially, our tool converts M from AADL into a database representation using the *OSATE-database converter*. All subsequent steps will be performed using this databased (model DB). The *analysis execution controller* constructs the analysis graph γ , as described in Sec. 4.1, and delegates the verification of A and G to an appropriate *verification engine*, which is determined by the form in Def. 3, along with \mathcal{A} , \mathcal{T} , and \mathcal{R} in the contract. A selected verification engine populates an SMT problem and a Promela model with values from the model DB, executes the verification via Z3 and Spin, and interprets the outputs. To verify forms 3 and 4 of Def. 3, σ_{Sched} and σ_{Batt} verification engines use Z3 to find assignments to $v_1 \dots v_j$ that satisfy ϕ , and use these assignments to verify ψ via Spin.

Our tool can be extended with new domains. This is done by creating new domains σ , that define \mathcal{A} and \mathcal{S} in AADL, develop a verification engine that interprets \mathcal{T} and \mathcal{R} , and add a tool that carries out the verification for the new domain. In addition, it is possible to incorporate new verification tools (e.g., UPPAAL instead of Spin for σ_{Sched}) as well as new analysis tools (e.g., a higher-fidelity method of calibrating K to predict thermal runaways).

5.2 Scheduling Domain Implementation

We encode $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ as a Promela (the input language of the Spin [13]) program that computes $q(\text{CanPrmpt})(t_1, t_2)$ appropriately for each state q and pair of threads t_1, t_2 , as described in Sec. 4.2.3.

We implement each task t as a process and add a manager process that decides what priorities are assigned to threads and what threads are dispatched to processors. Thus, the manager process plays the role of a scheduler and a dispatcher. Our Promela program handles the events of job arrival and termination in an infinite cycle, interleaving each event with the manager execution.

The Promela program models non-deterministic job terminations without explicit time counters as follows. The manager process calculates possible upcoming events. Time is advanced in a greedy manner (i.e., whenever possible): if an arrival event happens, or the earliest of all the possible job termination events. To achieve a finite state space, all clock variables⁸ are reduced by the minimum value of all clock variables periodically. It can be shown that this approach

⁸Such as the next arrival of a job or the absolute system time.

maintains contract satisfaction as long as clock variables are not used in a contract. This condition holds for our contract language since the clock variables are not exposed in the language.

5.3 Battery Domain Implementation

We encode $\mathcal{K}(\llbracket \mathcal{T}_{Batt} \rrbracket)$ as a single-process Promela program. The program maintains the state CellCharge , CellSt , Gr discussed in Sec. 4.2.4. The program execution works in two steps: scheduling cells for discharge and charge (i.e., changing Gr and CellSt), and advancing the charge state (i.e., changing CellCharge).

The first step is deterministic: it imitates the logic of a selected scheduler. **FGURR** does not change Gr and rotates through groups, setting ParalReq groups to discharge at a time, and setting the rest to idle. **FGWRR** does not change Gr either, but instead of rotating the groups it sorts them in non-increasing order of charge (which, for us, is the number of cells with $\text{CellCharge}(c) = T$) and selects the top ParalReq groups. **GPWRR** tries to assemble groups by packing as many charged cells into each group as possible. Then it selects the top ParalReq most charged groups to discharge. Within each group, all schedulers select SerialReq number of charged cells.

The other step is non-deterministic: every discharging cell non-deterministically becomes discharged, every charging cell non-deterministically becomes charged; idle cells, however, do not change their charges. The program terminates when there is not enough charge to supply the output requirements.

This program is an overapproximation of high-fidelity battery models with precise measurements of the cell charge. This measurement is then used in these models to schedule the cells. Thanks to the non-determinism in the second step, our implementation accounts for possible cell failures (i.e., cell gets immediately discharged) and any high-fidelity model of charge. On the other hand, the program represents cell schedulers' logic precisely.

6. EVALUATION

In this section we evaluate two aspects of our approach: (i) the effective integration of analyses – detection of integration errors or demonstration of their absence – and (ii) the scalability of our tool for models of practical size.

Analysis Integration. Consider a concrete configuration for the sample aircraft: threads t_1, t_2, t_3 have $\llbracket \text{Per} \rrbracket_M = \{t_1 \mapsto 100, t_2 \mapsto 150, t_3 \mapsto 200\}$, $\llbracket \text{Dline} \rrbracket_M = \{t_1 \mapsto 100, t_2 \mapsto 90, t_3 \mapsto 200\}$, $\llbracket \text{WCET} \rrbracket_M = \{t_1 \mapsto 10, t_2 \mapsto 15, t_3 \mapsto 20\}$ are allocated to a single processor. Before analysis An_{FreqSc} is applied to determine processor frequencies, its assumption $C_{FreqSc.A}$ is verified. Recall that $C_{FreqSc.A}$ states that the scheduling policy must be semantically equivalent to DMS. Suppose first that the system uses RMS scheduling, i.e., $\text{Prior}(t_1) > \text{Prior}(t_2) > \text{Prior}(t_3)$. In this case, our tool detects a violation of $C_{FreqSc.A}$ via model checking because in this case DMS would assign $\text{Prior}(t_2) > \text{Prior}(t_1)$. Now suppose that the system uses EDF. As our Spin program indicates, it satisfies $C_{FreqSc.A}$. Thus, our approach not only prevents incorrect usage of An_{FreqSc} , but also extends its applicability to the case of EDF.

Next, suppose our system has a battery with $\text{BatRows} = \text{BatCols} = 4$, and a voltage requirement $\text{ParalReq} =$

Threads	DMS/RMS Time ¹⁰	EDF Time
3	0.01	0.01
4	0.01	0.52
5	0.07	33.4
6	0.37	2290.0
7	2.18	MEMLIM
8	12.4	MEMLIM
9	71.2	MEMLIM
10	421	MEMLIM
11	MEMLIM	MEMLIM

Table 3: Scalability of the $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ program.

$SerialReq = 3$. It has been observed [10] that heat-dissipating cells (i.e., those with many thermal neighbors) and heat-isolated cells (i.e., those with no thermal neighbors) tend to prevent thermal runaway, while cells with one thermal neighbor tend to accumulate heat and lead to runaway. An assignment of weights $K(0) = K(1) = K(2) = 2, K(1) = -1$ in (1) captures this intuition. After executing analysis $An_{BatSched}$, which picks a battery scheduler, our tool verifies its guarantee $C_{BatSched}.G$. Since $An_{BatSched}$ is not aware of thermal runaway, not every scheduler meets the guarantee. As our Spin verification indicates, **FGWRR** and **FGURR** satisfy it, but **GPWRR** fails because it causes the system to reach a configuration that violates (1) with $TN(0) = TN(3) = 0, TN(1) = 8, TN(2) = 1$. Thus, our approach detects possibility of thermal runaway even though the existing analysis $An_{BatSched}$ does not.

Scalability of Contract Verification. We evaluate the scalability of our approach by comparing it to an alternative based on a unified semantic model. We focus on model checking since it is by far the most expensive component of our algorithm. The execution semantics of a unified model would, at the very least, consist of the interleaving of the two Kipke structures – $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ and $\mathcal{K}(\llbracket \mathcal{T}_{Therm} \rrbracket)$. Model checking this interleaving would be intractable due to statespace explosion. On the other hand, our approach is compositional and always verifies $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ and $\mathcal{K}(\llbracket \mathcal{T}_{Therm} \rrbracket)$ in isolation.

We evaluated our Promela programs using a general-purpose Amazon EC2⁹ virtual machine with 8 cores and 30 Gb memory. The worst-case exploration times by scheduler for the full statespace $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ and $\mathcal{K}(\llbracket \mathcal{T}_{Therm} \rrbracket)$ are shown in Tab. 3 and Tab. 4, respectively. For the former we use threads with implicit harmonic periods, and for the latter we grow the battery size, fixing the output voltage requirement to $SerialReq = ParalReq = 3$. Although the complexity growth is exponential, $\mathcal{K}(\llbracket \mathcal{T}_{Sched} \rrbracket)$ is verifiable upto 6-10 threads (per CPU), and $\mathcal{K}(\llbracket \mathcal{T}_{Therm} \rrbracket)$ is verifiable up to batteries with 25 cells. We believe that this enables verification of realistic CPSs. We expect that other techniques, such as abstraction and symbolic model-checking, will help us to push these limits even further.

7. RELATED WORK

Contracts and assume-guarantee reasoning have been used extensively to enable modular verification. In particular, in the development and verification of CPS, contracts provide an important alternative to a unified semantic model. For instance, Torngren et al. [26] use architectural viewpoints

⁹<http://aws.amazon.com/ec2>

¹⁰ All times are in seconds. MEMLIM indicates that the verification exceeded the limit of 30Gb.

Cells	FGURR Time ¹⁰	FGWRR Time	GPWRR Time
9	0.13	0.15	0.15
12	0.61	2.34	3.94
16	44.0	31.4	127
20	1060	619	MEMLIM
25	MEMLIM	MEMLIM	MEMLIM

Table 4: Scalability of the $\mathcal{K}(\llbracket \mathcal{T}_{Batt} \rrbracket)$ program.

contracts as a coordination tool for designers from different domains when designing a mechatronic system where multiple tools are used for hardware, software, control, and mechanical design.

Sangiovanni-Vincentelli et al. [24] use contracts between components and platform-based design to combine the semantics of multiple domains. A similar approach is used in the SPEEDS project [2] to enable speculative design to support teams of distributed designers. Specifically, the authors propose the use of “rich” components where functional and non-functional aspects of the system are combined. We do not require a model that semantically unifies multiple domains, and our focus is not on the interaction between *components*. Instead we use contracts to capture the semantics of, and interaction between, the *analyses* themselves.

Rajhans et al. [22] present an architectural multi-view approach where different modeling notations are captured in different views of an architectural model, using structural and semantic mappings to ensure consistency. In contrast, while we capture similar semantics on the analysis contracts, we only capture the interactions between analyses.

Assume-Guarantee reasoning for control theory has been widely explored. Frehse et al. [9] develop assume-guarantee reasoning for hybrid systems based on over-approximation by simulation to enable compositional reasoning. A similar approach is taken by Girard and Pappas [11] using bisimulation, simulation, and language inclusion to develop approximate system relationships. These approaches are focused on the intersection of control theory and computer science. In contrast, our work aims at capturing a larger set of domains with extensible representations of domains.

Cofer et al. [5] present an approach to add architectural contracts to AADL components to enable compositional verification. While we share the AADL platform, we use the AADL annexes to specify contracts on the analysis and within these contracts we specify the components accessed by these analyses. The FUSED [18] project defines a meta-language approach to merge notations from multiple analysis domains to enable a syntactic integration of multiple tools. It is limited to what the syntax can express and unable to reason semantically. For instance, it would not be able to reason about when different scheduling algorithms are equivalent depending on different parameters of the taskset. In our approach we model the semantics in the analyses algorithms and are able to reason about these behavioral differences that are hidden at the syntactic level.

The work closest to ours is by Nam et al. [19]. However, we go beyond them in a number of aspects. Their contracts are restricted to a single domain (resource allocation), the contract verification is unsound and incomplete (since it explores the system statespace only up to a finite depth), and their implementation is tied to a specific tool (Alloy) which cannot do model checking or SMT solving. In contrast, our contract language supports multiple domains, our algorithm is sound and exhaustive, and our implementation relies on

extensible co-operative use of SMT solvers and model checkers.

8. CONCLUSION AND FUTURE WORK

In this paper we presented an analysis integration approach for the development of cyber-physical systems. Our approach uses novel “analysis contracts” to formally specify and automatically verify interactions between analyses from different scientific domains. Our contract language can express both static and runtime aspects of models to which analyses are applicable. We presented the syntax and formal semantics of contracts, as well as an algorithm that orders a set of analyses in a correct way and checks the applicability of each analyses via co-operative use of SMT solving and model checking. We discussed how our approach captures the semantics of a new domain in order to support analysis specific to it. Contrasting with other approaches, we show how we avoided developing a unifying semantics (that leads to an intractable verification process), focusing instead on the interactions between domain-specific analyses. We then presented the integration of the analysis contracts into the AADL/OSATE toolkit with the use of an AADL annex. Finally, we used an example to demonstrate the application of our framework to realistic models. As future work we plan to explore more scalable contract verification tools and define contracts for the contract verification tools themselves.

9. REFERENCES

- [1] N. Audsley, A. Burns, M. F. Richardson, and A. J. Wellings. Hard real-time scheduling: The deadline-monotonic approach. In *in Proc. IEEE Workshop on Real-Time Operating Systems and Software*, pages 133–137, 1991.
- [2] L. Benvenuti, A. Ferrari, L. Mangeruca, E. Mazzi, R. Passerone, and C. Sofronis. A contract-based formalism for the specification of heterogeneous systems. In *Proc. of FDL*, 2008.
- [3] S. Chaki, A. Gurfinkel, S. Kong, and O. Strichman. Compositional Sequentialization of Periodic Programs. In *Proc. of VMCAI*, 2013.
- [4] Y. Chen and J. W. Evans. Thermal analysis of Lithium-Ion batteries. *J. of The Electrochemical Society*, 143(9):2708–2712, Sept. 1996.
- [5] D. D. Cofer, A. Gacek, S. P. Miller, M. W. Whalen, B. LaValley, and L. Sha. Compositional verification of architectural models. In *Proc. of NFM*, 2012.
- [6] L. M. de Moura and N. Bjørner. Z3: An efficient smt solver. In *Proc. of TACAS*, 2008.
- [7] D. De Niz and R. Rajkumar. Partitioning bin-packing algorithms for distributed real-time systems. *IJES*, 2(3), 2006.
- [8] P. H. Feiler and D. P. Gluch. *Model-Based Engineering with AADL - An Introduction to the SAE Architecture Analysis and Design Language*. SEI series in software engineering. Addison-Wesley, 2012.
- [9] G. Frehse, H. Zhi, and B. Krogh. Assume-guarantee reasoning for hybrid I/O-automata by over-approximation of continuous interaction. In *Proc. of CDC*, 2004.
- [10] Gi-Heon Kim and Ahmad Pesaran. Analysis of heat dissipation in li-ion cells & modules for modeling of thermal runaway. In *Proc. of 3rd International Symposium on Large Lithium ion Battery Technology and Application*, Long Beach, CA, 2007.
- [11] A. Girard and G. J. Pappas. Approximate bisimulation: A bridge between computer science and control theory. *Eur. J. Control*, 17(5-6), 2011.
- [12] K. Hahnsang and K. G. Shin. Scheduling of battery charge, discharge, and rest. In *Proc. of RTSS*, 2009.
- [13] G. J. Holzmann. The model checker spin. *IEEE Trans. Software Eng.*, 23(5), 1997.
- [14] X. Jin, J. V. Deshmukh, J. Kapinski, K. Ueda, and K. Butts. Powertrain control verification benchmark. In *Proc. of the 17th International Conference on Hybrid Systems: Computation and Control*, pages 253–262, New York, NY, USA, 2014. ACM.
- [15] G. Karsai and J. Sztipanovits. Model-integrated development of cyber-physical systems. In *Proc. of SEUS*, 2008.
- [16] H. Kim and K. Shin. On dynamic reconfiguration of a large-scale battery system. In *15th IEEE Real-Time and Embedded Technology and Applications Symposium, 2009. RTAS 2009*, pages 87–96, Apr. 2009.
- [17] C. L. Liu and J. W. Layland. Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment. *JACM*, 20(1), 1973.
- [18] M. B. M, M. Michalowski, A. Schwerdfeger, H. Shackleton, and S. Vestel. The FUSED Meta-Language and Tools for Complex System Engineering. In *Proc. of AVICPS*, 2011.
- [19] M.-Y. Nam, D. de Niz, L. Wrage, and L. Sha. Resource allocation contracts for Open Analytic Runtime models. In *Proc. of EMSOFT*, 2011.
- [20] Osate 2 -An Open-Source Tool Platform for AADL v2. <http://www.aadl.info>.
- [21] A. Pnueli. The Temporal Logic of Programs. In *Proc. of FOCS*, 1977.
- [22] A. Rajhans, A. Y. Bhave, I. Ruchkin, B. Krogh, D. Garlan, A. Platzer, and B. Schmerl. Supporting heterogeneity in cyber-physical systems architectures. *IEEE Transactions on Automatic Control's Special Issue on Control of CPS*, 2014. To appear.
- [23] S. Saewong and R. Rajkumar. Practical voltage-scaling for fixed-priority rt-systems. In *Proc. of RTAS*, 2003.
- [24] A. L. Sangiovanni-Vincentelli, W. Damm, and R. Passerone. Taming dr. frankenstein: Contract-based design for cyber-physical systems. *Eur. J. Control*, 18(3), 2012.
- [25] J. Sztipanovits, X. Koutsoukos, G. Karsai, N. Kottenstette, P. Antsaklis, V. Gupta, B. Goodwine, J. Baras, and S. Wang. Toward a science of cyber physical system integration. *Proc. of the IEEE*, 100(1):29–44, Jan. 2012.
- [26] M. Törngren, A. Qamar, M. Biehl, F. Loiret, and J. El-khoury. Integrating viewpoints in the development of mechatronic products. *Mechatronics*, 2013.