

A Compositional Approach for Schedulability Analysis of Distributed Avionics Systems

Pujie Han Zhengjun Zhai
School of Computer Science and Engineering
Northwestern Polytechnical University
Xi'an, China
{hanpujie,zhaizjun}@mail.nwpu.edu.cn

Brian Nielsen Ulrik Nyman
Department of Computer Science
Aalborg University
Aalborg, Denmark
{bnielsen,ulrik}@cs.aau.dk

This work presents a compositional approach for schedulability analysis of Distributed Integrated Modular Avionics (DIMA) systems that consist of spatially distributed ARINC-653 modules connected by a unified AFDX network. We model a DIMA system as a set of stopwatch automata in UPPAAL to verify its schedulability by model checking. However, direct model checking is infeasible due to the large state space. Therefore, we introduce the compositional analysis that checks each partition including its communication environment individually. Based on a notion of message interfaces, a number of message sender automata are built to model the environment for a partition. We define a timed selection simulation relation, which supports the construction of composite message interfaces. By using assume-guarantee reasoning, we ensure that each task meets the deadline and that communication constraints are also fulfilled globally. The approach is applied to the analysis of a concrete DIMA system.

1 Introduction

The architecture of Distributed Integrated Modular Avionics (DIMA) has been successfully applied to the aviation industry. A DIMA system installs standardized computer modules in spatially distributed locations[19] that are connected by a unified bus system[3] such as an AFDX network. Avionics applications residing on the modules run in ARINC-653[1] compliant operating systems. The generic distributed structure of DIMA significantly improves performance and availability as well as reduces development and maintenance costs, while it also dramatically increases the complexity of schedulability analysis. A schedulable DIMA system should fulfil not only the temporal requirements of real-time tasks in each ARINC-653 module but also communication constraints among the distributed nodes. As a result, the system integrators need to consider both computation and communication when analyzing the schedulability of DIMA architecture.

Currently, model checking approaches have been increasingly developed in the schedulability analysis of complex real-time systems. However, we found no studies that analyzed the schedulability of distributed avionics systems as a whole including the network by model checking. The related research isolates computation modules from their underlying network, thereby considering these nodes as independent hierarchical scheduling systems or investigating the network in isolation, which possibly leads to pessimistic results. There have been works using model-checking to analyze the temporal behavior of individual avionics modules in various formal models such as Coloured Petri Nets (CPN)[10], preemptive Time Petri Nets (pTPN)[5], Timed Automata (TA)[2], and StopWatch Automata (SWA)[16, 8], and verify schedulability properties via state space exploration. Unfortunately, when being applied to concrete avionics systems, all of them suffer from an inevitable problem of state space explosion. For hierarchical scheduling systems, some studies[6, 18, 4] exploit the inherent temporal isolation of ARINC-653

partitions[1] and analyze each partition separately, but they ignore the behavior of the underlying network or the interactions among partitions. Thus these methods are not applicable to DIMA environments in which multiple distributed ARINC-653 partitions communicate through a shared network to perform an avionics function together.

In this paper, we present a compositional approach for schedulability analysis of DIMA systems that are modeled as UPPAAL SWA, i.e. the TA extended with stopwatches. Compared with the clocks in TA, stopwatches can be blocked and resumed at any location and thus are effective in modeling task preemption. We decompose the system in such a way that we can check each ARINC-653 partition *including* a model of its communication environment individually and then assemble the local results together to derive conclusions about the schedulability of an entire system. Thereby, we verify a number of smaller, simpler, abstract systems rather than directly verifying a larger, more complex, concrete system including the details about all the partitions and the network. The main contributions of this paper are summarized as follows:

- A *compositional approach* performs assume-guarantee reasoning[12] to reduce the complexity of symbolic model-checking in the schedulability analysis of DIMA systems.
- An *abstraction relation*, timed selection simulation relation, allows users to create a set of abstract models that collectively describe the external behavior of a concrete model, thereby simplifying the abstraction in assume-guarantee reasoning.
- A notion of *message interfaces* decouples the communication dependencies between partitions. By composing any partition with its related message interfaces and verifying safety properties of the composition, we can conclude that these properties are still preserved at the global level.

The rest of the paper is organized as follows. Section 2 gives the necessary formal notions. The UPPAAL modeling of DIMA systems is presented in section 3. Section 4 gives the concept of timed selection simulation and its properties. In section 5, we detail the compositional analysis approach. Section 6 shows an experiment on a concrete DIMA system, and section 7 finally concludes.

2 Preliminaries

In this section, we present formal definitions including SWA with an input/output extension and its semantic object Timed I/O Transition Systems(TIOTSs)[9].

Suppose that C is a finite set of clocks and V is a finite set of integer variables. A *valuation* $u(x)$ with $x \in C \cup V$ denotes a mapping from C to $\mathbf{R}_{\geq 0}$ and from V to \mathbf{N} . Let $LC(C, V)$ be the set of linear constraints. A *guard* $g \in LC(C, V)$ is a linear constraint which is defined as a finite conjunction of atomic formulae in the form of $c \sim n$, $c - c' \sim n$ or $v \sim n$ with $c, c' \in C, v \in V, n \in \mathbf{N}$, and $\sim \in \{>, <, =\}$. Given any valuation u , we change the values of clocks and integer variables using an *update* operation $r(u) \in 2^R$ in the form of $c = 0$ or $v = n$ where $c \in C, v \in V$ and $n \in \mathbf{N}$, and R is the set of all possible update operations. In addition, we define an *action* set Σ . All the actions can be subsumed under two sets of unicast actions Σ^u and broadcast actions Σ^b . By contrast, $\tau \notin \Sigma$ denotes an internal action and $\Sigma^\tau = \Sigma \cup \{\tau\}$.

Definition 1 (Stopwatch Automaton[7]). *A stopwatch automaton is a tuple $\langle Loc, l_0, C, V, E, \Sigma, Inv, drv \rangle$ where Loc is a finite set of locations, $l_0 \in Loc$ is the initial location, C is a finite set of clocks, V is a finite set of integer variables, $E \subseteq Loc \times LC(C, V) \times \Sigma^\tau \times 2^R \times Loc$ is a set of edges, $\Sigma = I \oplus O$ is a finite set of actions divided into inputs(I) and outputs(O), Inv is a mapping $Loc \rightarrow LC(C, V)$, and drv is a mapping $Loc \times C \rightarrow \{0, 1\}$.*

From a syntactic viewpoint, SWA belongs to the class of TA extended with drv , which can prevent part of the clocks from changing in specified locations semantically. We now shift the focus to the semantic object TIOTS of SWA.

In a TIOTS, there are two types of transitions: delay and action transitions. We use the set $D = \{\varepsilon(d) \mid d \in \mathbf{R}_{\geq 0}\}$ to denote the delay, and refer to the 0-delay $\varepsilon(0)$ as $\mathbf{0}$.

Definition 2 (Timed I/O Transition System). *A timed I/O transition system is a tuple $\mathcal{T} = \langle S, s_0, \Sigma, \rightarrow \rangle$ where S is an infinite set of states, s_0 is the initial state, $\Sigma = I \oplus O$ is a finite set of actions divided into inputs (I) and outputs (O), $I \cap O \subseteq \Sigma^u$, and $\rightarrow \subseteq S \times \Sigma^\tau \cup D \times S$ is a transition relation. $s \xrightarrow{a} s'$ represents $(s, a, s') \in \rightarrow$, which has the properties of time determinism, time reflexivity, and time additivity[9].*

For any SWA, a state is defined as a pair $\langle l, u \rangle$ where l is a location and u is a valuation over clocks and integer variables. On the basis of TIOTSs, the operational semantics of SWA is defined as follows.

Definition 3. *The operational semantics of a stopwatch automaton $A = \langle Loc, l_0, C, V, E, \Sigma, Inv, drv \rangle$ is a timed I/O transition system $\mathcal{T}^A = \langle S, s_0, \Sigma, \rightarrow \rangle$ where S is the set of states of A , $s_0 = \langle l_0, u_0 \rangle$ is the initial state of A , Σ is the same set of actions as A , and \rightarrow is the transition relation defined by*

- $\langle l, u \rangle \xrightarrow{a} \langle l', u' \rangle$ iff $\exists \langle l, g, a, r, l' \rangle \in E$ ($u \models g \wedge u' = r(u) \wedge u' \models Inv(l')$)
- $\langle l, u \rangle \xrightarrow{\varepsilon(d)} \langle l', u' \rangle$ iff $l = l' \wedge (\forall v \in V \ u'(v) = u(v)) \wedge (\forall c \in C \ (drv(l, c) = 0 \Rightarrow u'(c) = u(c))) \wedge (\forall c \in C \ (drv(l, c) = 1 \Rightarrow u'(c) = u(c) + d)) \wedge u' \models Inv(l')$.

For any transition $s \xrightarrow{a} s'$, two symbols $a?$ and $a!$ denote the action a belonging to input I and output O respectively. Given $a \in \Sigma$, $s \xrightarrow{a}$ iff $\exists s' \in S$, s.t. $s \xrightarrow{a} s'$. $\xrightarrow{\tau^*}$ or $\xrightarrow{\mathbf{0}}$ denotes the reflexive and transitive closure of $\xrightarrow{\tau}$. $s \xrightarrow{\varepsilon(d)} s'$ iff $s \xrightarrow{\varepsilon(d)} s'$, or $\exists s_1, s_2, \dots, s_n \in S$, s.t. $s \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{n-1}} s_n \xrightarrow{\alpha_n} s'$ and $\forall i \in \{0, \dots, n\}$, s.t. $\alpha_i = \tau$ or $\alpha_i \in D$ and $d = \sum \{d_i \mid \alpha_i = \varepsilon(d_i)\}$.

The definition of parallel composition \parallel of TIOTSs is similar to that in [9]. Given two TIOTSs $\mathcal{T}_i = \langle S_i, s_{i,0}, \Sigma_i, \rightarrow_i \rangle, i \in \{1, 2\}$, they are *compatible* iff they satisfy the following conditions:

- (Unique output) $O_1 \cap O_2 = \emptyset$.
- (Deterministic-pair unicast) $I_1 \cap I_2 \cap \Sigma^u = \emptyset$.

Note that broadcast actions in the composition of TIOTSs are *input-enabled*: $\forall s \in S_i \ \forall a \in I_i \cap \Sigma^b \ s \xrightarrow{a}$.

Definition 4 (Parallel Composition). *Suppose two timed I/O transition systems $\mathcal{T}_1 = \langle S_1, s_{1,0}, \Sigma_1, \rightarrow_1 \rangle$ and $\mathcal{T}_2 = \langle S_2, s_{2,0}, \Sigma_2, \rightarrow_2 \rangle$ are compatible. The parallel composition $\mathcal{T}_1 \parallel \mathcal{T}_2$ is the timed I/O transition system $\langle S, s_0, \Sigma, \rightarrow \rangle$ where $S = S_1 \times S_2$, $s_0 = \langle s_{1,0}, s_{2,0} \rangle$, $\Sigma = I_{1 \parallel 2} \oplus O_{1 \parallel 2}$, $I_{1 \parallel 2} = (I_1 \setminus (O_2 \cap \Sigma^b)) \cup (I_2 \setminus (O_1 \cap \Sigma^b))$, $O_{1 \parallel 2} = O_1 \cup O_2$, and \rightarrow is the largest relation generated by the following rules:*

- *INDEP-L*: $\frac{s_1 \xrightarrow{a} s'_1 \quad a \in \{\tau\} \cup \Sigma_1 \setminus \Sigma_2}{\langle s_1, s_2 \rangle \xrightarrow{a} \langle s'_1, s_2 \rangle}$ *INDEP-R*: $\frac{s_2 \xrightarrow{a} s'_2 \quad a \in \{\tau\} \cup \Sigma_2 \setminus \Sigma_1}{\langle s_1, s_2 \rangle \xrightarrow{a} \langle s_1, s'_2 \rangle}$
- *DELAY*: $\frac{s_1 \xrightarrow{\varepsilon(d)} s'_1 \quad s_2 \xrightarrow{\varepsilon(d)} s'_2 \quad d \in \mathbf{R}_{\geq 0}}{\langle s_1, s_2 \rangle \xrightarrow{\varepsilon(d)} \langle s'_1, s'_2 \rangle}$
- *SYNC-IN*: $\frac{s_1 \xrightarrow{a} s'_1 \quad s_2 \xrightarrow{a} s'_2 \quad a \in I_{1 \parallel 2}}{\langle s_1, s_2 \rangle \xrightarrow{a} \langle s'_1, s'_2 \rangle}$
- *SYNC-BIO*: $\frac{s_1 \xrightarrow{a} s'_1 \quad s_2 \xrightarrow{a} s'_2 \quad a \in (I_1 \cap O_2) \cup (O_1 \cap I_2) \cap \Sigma^b}{\langle s_1, s_2 \rangle \xrightarrow{a} \langle s'_1, s'_2 \rangle}$

- *SYNC-UIO*:
$$\frac{s_1 \xrightarrow{a} s'_1 \quad s_2 \xrightarrow{a} s'_2 \quad a \in I_{1||2} \cap O_{1||2}}{\langle s_1, s_2 \rangle \xrightarrow{\tau} \langle s'_1, s'_2 \rangle}$$

We use Ω to denote the set of TA and SWA in our modeling framework. For any $A, B \in \Omega$, we define the composite model $C = A||B$ iff their TIOTSs satisfy $\mathcal{T}^C = \mathcal{T}^A||\mathcal{T}^B$.

3 Avionics System Modeling

We focus on a generic DIMA architecture including a set of ARINC-653 modules connected by an AFDX network, as shown in Fig.1. There is a three-layer structure in the DIMA system that consists of scheduling, task, and communication layers.

The *scheduling layer* is defined as the scheduling facilities for generic computation resources of a DIMA system, where standardized computer modules execute concurrent application tasks in partitioned operating systems. In this operating system, partitions are scheduled by a Time Division Multiplexing (TDM) scheduler and each partition also has its local scheduling policy, preemptive Fixed Priority (FP), to manage the internal tasks[1]. The scheduling layer is modeled as two TA templates PartitionSupply and TaskScheduler in UPPAAL¹. The PartitionSupply depicted in Fig.2 provides the service of TDM partitioning for a particular partition pid. The TaskScheduler implementing FP scheduling allocates processor time to the task layer only when the partition is active.

The *task layer* contains all the application tasks executing avionics functions. A task is regarded as the smallest scheduling unit, each of which runs concurrently with other tasks in the same partition. The execution of a task is modelled as a sequence of commands that are either computing for a duration, locking/unlocking a resource, or sending/receiving a message. We consider two task types: *periodic tasks* and *sporadic tasks*. A periodic task has a fixed release period, while a sporadic task is characterized by a minimum separation between consecutive jobs. The task layer is instantiated from two SWA templates PeriodicTask and SporadicTask in UPPAAL. Since the tasks in a partition are scheduled by a task scheduler, we use a set of binary channels as scheduling actions to communicate between task models and TaskScheduler.

The *communication layer* carries out inter-partition communication over a common AFDX network. The AFDX protocol stack realized by an End System(ES) interfaces with the task layer through ARINC-653 ports. Based on the AFDX protocol structure, the communication layer is further divided into UDP/IP layer and Virtual Link layer, where a Virtual Link (VL) ensures an upper bound on end-to-end delay. In UPPAAL, the UDP/IP layer is divided into two TA templates IPTx and IPRx, which calculate the latency of the UDP/IP layer in a transmitting ES and a receiving ES respectively. Similarly, two TA templates VLinkTx and VLinkRx model the delay of a VL in opposite directions.

From a global view of the system, its schedulability is also affected by the communication layer. According to the ARINC-653 standard[1], there are two types of ARINC-653 ports, sampling ports

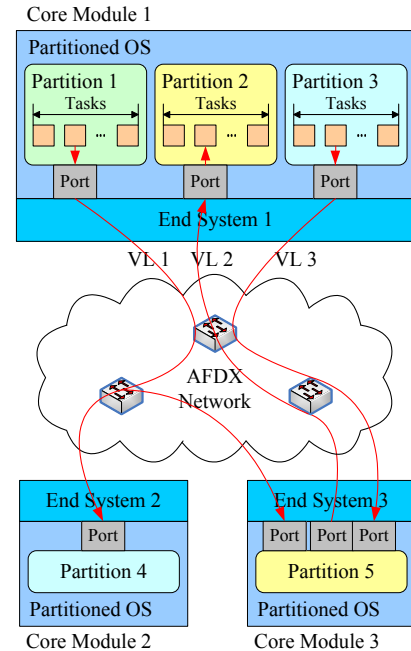


Figure 1: An Example of DIMA systems

¹Models available at <http://eptcs.web.cse.unsw.edu.au/paper.cgi?MARSVPT2018:2>




Figure 2: The UPPAAL Template of an ARINC-653 Partition Scheduler

and queuing ports. A sampling port can accommodate at most a single message that remains until it is overwritten by a new message. A refresh period is defined for each sampling port. This attribute provides a specified arrival rate of messages, regardless of the rate of receiving requests from tasks. In contrast, a queuing port is allowed to buffer multiple messages in a message queue with a fixed capacity. However, the operating system is not responsible for handling overflow from the message queue.

In this paper, we verify the following three typical schedulability properties:

- All the tasks meet their deadlines in each partition.
- The refresh period of any sampling port is guaranteed.
- The overflow from any queuing ports must be avoided.

The schedulability of an avionics system is described and verified as a safety property of the above TA/SWA models. We add a set *Err* of *error locations* to the templates. Once schedulability is violated, the related model will lead itself to one of the error locations immediately. Thus, the schedulability is replaced with this safety property φ :

$$A[] \neg(\bigvee_{loc \in Err} loc), \quad (1)$$

which belongs to a simplified subset of TCTL used in UPPAAL.

However, since the verification algorithm inside UPPAAL for SWA introduces a slight over-approximation[7]², UPPAAL may sometimes give the verification result “Maybe satisfied” or “May not be satisfied”. To further refine the result in this case we manually analyse the possible counter example using UPPAAL’s concrete simulator to determine if the system is unschedulable. Alternatively, the statistical model-checking (SMC) engine could be invoked to attempt an automatic falsification. In our experiences, the result only appears when the system is on the very borderline of being schedulable.

4 Timed Selection Simulation

We propose a notion of timed selection simulation relation to support assume-guarantee reasoning. Compared with some other abstraction relations like timed simulation[15] and timed ready simulation[14], timed selection simulation only abstracts a selected subset of actions from the concrete model. Applying timed selection simulation to the abstraction of a concrete system, one can pay attention to part of the system, individually model the behavior of each component, and thereby obtain a composite abstract model rather than a monolithic one.

Considering the semantic object \mathcal{T}^A of an automaton $A \in \Omega$, we denote the *error states* of \mathcal{T}^A by the set $\mathcal{E} = \{ \langle l, u \rangle \mid l \in Err \}$ where *Err* is the error-location set of A . Thus, for any TIOTS $\mathcal{T} = \langle S, s_0, \Sigma, \rightarrow \rangle$,

²Exact reachability for SWA with more than 3 stopwatches is known to be undecidable[7].

its error states are defined as a set $\mathcal{E} \subseteq S$, and the following function $g : S \rightarrow \{true, false\}$ indicates whether a state $s \in S$ has violated schedulability properties:

$$g(s) = \begin{cases} true & \text{if } s \in \mathcal{E} \\ false & \text{if } s \notin \mathcal{E}. \end{cases} \quad (2)$$

Given two compatible TIOTSs $\mathcal{T}_i, i \in \{1, 2\}$ with the error-state set \mathcal{E}_i , their composition $\mathcal{T}_1 \parallel \mathcal{T}_2$ has the error-state set $\mathcal{E}_{\mathcal{T}_1 \parallel \mathcal{T}_2} = \{(s_1, s_2) \mid s_1 \in \mathcal{E}_1 \vee s_2 \in \mathcal{E}_2\}$ and the function $g(\langle s_1, s_2 \rangle) = g(s_1) \vee g(s_2)$.

Based on the function $g(s)$, the formal definition of timed selection simulation is given as follows.

Definition 5 (Timed Selection Simulation). *Let $\mathcal{T}_1 = \langle S_1, s_{1,0}, \Sigma_1, \rightarrow_1 \rangle$ and $\mathcal{T}_2 = \langle S_2, s_{2,0}, \Sigma_2, \rightarrow_2 \rangle$ be two timed I/O transition systems with $\Sigma_2 \subseteq \Sigma_1$. Let R be a relation from S_1 to S_2 . We call R a timed selection simulation from \mathcal{T}_1 to \mathcal{T}_2 , written $\mathcal{T}_1 \preceq \mathcal{T}_2$ via R , provided $(s_{1,0}, s_{2,0}) \in R$ and for all $(s_1, s_2) \in R$, $g(s_1) = g(s_2)$ and*

1. if $s_1 \xrightarrow{a^?} s'_1$ for some $s'_1 \in S_1$, $a \in \Sigma_2$, then $\exists s'_2 \in S_2$ such that $s_2 \xrightarrow{a^?} s'_2$ and $(s'_1, s'_2) \in R$
2. if $s_1 \xrightarrow{a^!} s'_1$ for some $s'_1 \in S_1$, $a \in \Sigma_2$, then $\exists s'_2 \in S_2$ such that $s_2 \xrightarrow{a^!} s'_2$ and $(s'_1, s'_2) \in R$
3. if $s_1 \xrightarrow{a} s'_1$ for some $s'_1 \in S_1$, $a \in (\Sigma_1 \setminus \Sigma_2) \cup \{\tau\}$, then $\exists s'_2 \in S_2$ such that $s_2 \xrightarrow{\mathbf{0}} s'_2$ and $(s'_1, s'_2) \in R$
4. if $s_1 \xrightarrow{\varepsilon^{(d)}} s'_1$ for some $s'_1 \in S_1$, $d > 0$, then $\exists s'_2 \in S_2$ such that $s_2 \xrightarrow{\varepsilon^{(d)}} s'_2$ and $(s'_1, s'_2) \in R$.

Definition 6. *Let $A_i, i \in \{1, 2\}$ be stopwatch automata. We say that $A_1 \preceq A_2$, if and only if their corresponding timed I/O transition systems \mathcal{T}_i satisfy $\mathcal{T}_1 \preceq \mathcal{T}_2$.*

We now give some necessary properties of timed selection simulation.

Theorem 1. *Timed selection simulation \preceq is a preorder.*

For any automaton $A \in \Omega$, by construction, the reachability of its error locations is equivalent to that of the error states in the corresponding TIOTS \mathcal{T}^A . Hence the following theorem shows that timed selection simulation can preserve the satisfaction of the safety properties in the form of Eq.(1).

Theorem 2 (Property preservation). *Let $\mathcal{T}_i, i \in \{1, 2\}$ be timed I/O transition systems and \mathcal{E}_i be the set of error states of \mathcal{T}_i . Given a safety property $\varphi : \neg reach(\mathcal{E}_i)$ that any error states are not reachable, if $\mathcal{T}_1 \preceq \mathcal{T}_2$ and $\mathcal{T}_2 \models \varphi$, then $\mathcal{T}_1 \models \varphi$.*

Theorem 3 (Abstraction compositionality). *Let $\mathcal{T}_i, i \in \{1, 2, 3\}$ be timed I/O transition systems. If $\mathcal{T}_1 \preceq \mathcal{T}_2$, $\mathcal{T}_1 \preceq \mathcal{T}_3$, and \mathcal{T}_2 and \mathcal{T}_3 are compatible, then $\mathcal{T}_1 \preceq \mathcal{T}_2 \parallel \mathcal{T}_3$.*

Theorem 4 (Compositionality). *Let $\mathcal{T}_i = \langle S_i, s_{i,0}, \Sigma_i, \rightarrow_i \rangle, i \in \{1, 2, 3, 4\}$ be timed I/O transition systems. Suppose $\mathcal{T}_1 \parallel \mathcal{T}_3$ and $\mathcal{T}_2 \parallel \mathcal{T}_4$ are the parallel compositions of compatible timed I/O transition systems. If (1) $\mathcal{T}_1 \preceq \mathcal{T}_2, \mathcal{T}_3 \preceq \mathcal{T}_4$, and (2) $O_1 \cap I_4 \subseteq \Sigma_2 \subseteq \Sigma^b, I_2 \cap O_3 \subseteq \Sigma_4 \subseteq \Sigma^b$, then $\mathcal{T}_1 \parallel \mathcal{T}_3 \preceq \mathcal{T}_2 \parallel \mathcal{T}_4$.*

5 Compositional Analysis

We apply assume-guarantee reasoning to the schedulability analysis, and describe the schedulability goal as a safety property φ (Eq.(1)). As shown in Fig.3, our compositional analysis is comprised of the following four steps:

1. *Decomposition:* The system is first decomposed into a set of communicating partitions modeled by TA and SWA. The global property φ is also divided into several local properties, each of which belongs to one partition.

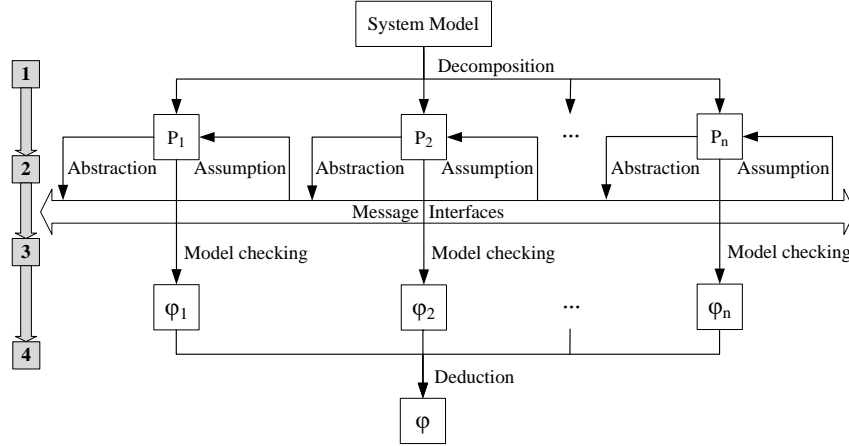


Figure 3: Compositional Analysis Procedure

2. *Construction of message interfaces:* We define message interfaces as the assumption and abstraction of the communication environment for each partition. In general, the templates of message interfaces should be built manually by the engineers.
3. *Model checking:* The local properties under the assumptions and the abstraction relations are verified by model checking.
4. *Deduction:* From the assume-guarantee rules, we finally derive the global property by combining all the local results.

The procedure can be performed automatically except for the first construction of message interfaces. We assume that a task never blocks while communicating with other partitions, which is commonly used in avionics systems[11, 6]. Otherwise a loop of communication dependency will cause circular reasoning, because the assumptions of a partition might be based on its own state recursively.

5.1 Decomposition

Assume that there are n constituent partitions in a system. Let $P_i, i \in \{1, 2, \dots, n\}$ be the SWA composite model of a partition. Let Err_i be the error-location set of P_i . The safety property $\varphi_i: A[\] \neg(\bigvee_{loc \in Err_i} loc)$ denotes the schedulability of P_i . The global property φ is therefore written as $\varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_n$, and the goal of our schedulability analysis is expressed as the verification problem:

$$P_1 \parallel P_2 \parallel \dots \parallel P_n \models \varphi \quad (3)$$

that can be further divided into n satisfaction relations:

$$P_1 \parallel P_2 \parallel \dots \parallel P_n \models \varphi_i, i \in \{1, 2, \dots, n\}. \quad (4)$$

Since the error-location set Err_i is only allowed to be manipulated by P_i , we check each partition model P_i independently for the corresponding *local property* φ_i instead of the original verification problem with a large and complex system. However, the communication environment of P_i , which denotes the behavior that P_i receives messages from other partitions, may affect the satisfaction of the schedulability property φ_i . Hence when performing the verification for partition P_i , one needs to give the *assumptions* of its communication environment and verifies the local property φ_i under these assumptions.




Figure 4: An Example of a Message Interface

5.2 Construction of message interfaces

A set of TA models is created to describe the message-sending behavior of a partition. Each of the TA is called a *message interface* of this partition and associated with a particular message type. Suppose there are a number of messages sent from partition P_j to another partition P_i and their corresponding message interfaces make up a composite TA model $A_{i,j}$. When we analyze P_i in the compositional way, it should be safe for $A_{i,j}$ to replace P_j . Hence, we say that a message interface of P_j is an *abstraction* of P_j .

Our abstraction of the message delivery between a partition and its underlying network is modelled using broadcast synchronization. A broadcast action represents a specific message types. Let $\Sigma_i = I_i \oplus O_i$ be the action set of a composite model for any partition P_i . An action $a_k \in I_i \cap \Sigma^b$ (resp. $a_k \in O_i \cap \Sigma^b$) denotes that P_i receives (resp. sends) messages with the type msg_k from (resp. to) other partition(s). The symbol $j \triangleright i$ represents the condition that there exists a partition P_j sending messages to P_i via an action set $O_{j \rightarrow i} \subseteq I_i \cap O_j$.

Definition 7 (Message Interface). *Let O_i be the output action set of a stopwatch automaton $P_i \in \Omega$. For any output action $a_k \in O_i \cap \Sigma^b$, the timed automaton A_i^k with an action set $\Sigma_i^k = O_i^k = \{a_k\}$ is a message interface of P_i if and only if there exists a timed selection simulation relation \preceq on Ω such that*

$$P_i \preceq A_i^k. \quad (5)$$

We build the templates of message interfaces in accordance with the characteristics of message-sending actions. In practice, the structure of an interface can be designed straightforwardly from the task specification. The template in Fig.4 shows a message interface that sends messages periodically via the action array `pmsg`. Then we make an automatized binary search for the interface's parameters such as offset in the template and meanwhile check the satisfaction of timed selection simulation relation.

The message interfaces can serve as the assumptions of the communication environment of a partition. The composition $A_{i,j}$ of the message interfaces A_j^k for all $a_k \in O_{j \rightarrow i}$ provides P_i with a “complete” abstraction of P_j , which models the behavior of all the output actions from P_j to P_i . According to the abstraction compositionality (Theorem 3) of the preorder \preceq , we have

$$P_j \preceq A_{i,j}. \quad (6)$$

Considering all the partitions except P_i in the system, we describe the communication environment of P_i as the composite model $\parallel_{j=1, j \neq i}^n A_{i,j}$.

5.3 Model checking

In the third step, the local property φ_i of P_i under assumption $\parallel_{j=1, j \neq i}^n A_{i,j}$ can be verified by model checking. We denote these n subproblems by

$$P_i \parallel \left(\parallel_{j=1, j \neq i}^n A_{i,j} \right) \models \varphi_i \quad i \in \{1, 2, \dots, n\}. \quad (7)$$

Normally, $A_{i,j}$ in Eq.(7) has a much smaller model size than its corresponding partition model P_j in Eq.(4). Thus, the compositional approach allows us to verify a simpler abstract partition model instead of a complex concrete system model including the details about all the partitions.

In addition, we capture the computation time of each task as an interval between a best-case and worst-case execution time. When analyzing the schedulability of a partition, the model-checker explores all scheduling decisions that can be made in such an interval, and hence also examines possible cases of scheduling timing anomalies[17].

5.4 Deduction

We derive the global property φ by combining n local results in the last step. For any schedulable system, each property φ_i should be concluded from the satisfaction of Eq.(7) under assumptions and all the abstraction relations of Eq.(6). According to the compositionality (Theorem 4) and property preservation (Theorem 2) of timed selection simulation, we have the following assume-guarantee rule:

$$\frac{\bigwedge_{\{j|j>i\}} P_j \preceq A_{i,j} \quad P_i \parallel \left(\prod_{j=1, j \neq i}^n A_{i,j} \right) \models \varphi_i}{P_1 \parallel P_2 \parallel \dots \parallel P_n \models \varphi_i} \quad (8)$$

Note that this assume-guarantee rule only provides a sufficient schedulability condition, for abstract message interfaces might slightly over-approximate the external behavior of a partition.

A simplified DIMA system exemplifies the reasoning procedure. In the example, the system model is decomposed into three partitions $P_i, i \in \{1, 2, 3\}$. We divide the global property φ into three local properties $\varphi_i, i \in \{1, 2, 3\}$. Accordingly, the goal of the verification problem is to check

$$P_1 \parallel P_2 \parallel P_3 \models \varphi_1 \wedge \varphi_2 \wedge \varphi_3. \quad (9)$$

From Eq.(4), this problem can be replaced with three subproblems:

$$P_1 \parallel P_2 \parallel P_3 \models \varphi_i, i \in \{1, 2, 3\}. \quad (10)$$

Without loss of generality, we take the verification of φ_1 for example to show how the model-checking and deduction are carried out in the following steps.

Assume that P_2 sends P_1 two types of messages, msg_1 and msg_2 , via two actions a_1 and a_2 respectively, and P_3 sends P_1 only a msg_3 with action a_3 . We create one message interface $A_j^k, j \in \{2, 3\}$ (like Eq.(5)) for each message type $msg_k (k \in \{1, 2, 3\})$ received by P_1 in the system. The abstraction relations from Eq.(5) can be expressed as

$$P_2 \preceq A_2^1, P_2 \preceq A_2^2, P_3 \preceq A_3^3. \quad (11)$$

From abstraction compositionality of the preorder \preceq , we can obtain

$$P_2 \preceq A_2^1 \parallel A_2^2, P_3 \preceq A_3^3. \quad (12)$$

Then, from reflexivity and compositionality of the preorder \preceq , the composite model of the system satisfies

$$P_1 \parallel P_2 \parallel P_3 \preceq P_1 \parallel A_2^1 \parallel A_2^2 \parallel A_3^3. \quad (13)$$

Note that when we apply the compositionality to checking a partition P_i , any output actions sent to P_i will never be removed in abstraction relations (Eq.(12)), which satisfies the condition (2) of theorem 4.

With Eq.(13), we have from property preservation of the abstraction relation \preceq that if

$$P_1 \parallel A_2^1 \parallel A_2^2 \parallel A_3^3 \models \varphi_1, \text{ then} \quad (14)$$

$$P_1 \parallel P_2 \parallel P_3 \models \varphi_1. \quad (15)$$

Since Eq.(15) covering all three partitions in the system has a higher complexity than Eq.(14), the techniques of model checking can be adopted to verify the simpler problem Eq.(14) instead of the original goal Eq.(15). The same steps will be repeated for local properties φ_2 and φ_3 .

Consequently, we conclude all the local results of (10) according to the reasoning process from Eq.(11) to Eq.(15). When we analyze the partition P_1 and its communication environment, the local result of Eq.(15) can be deduced from Eq.(11) and Eq.(14) in the following assume-guarantee rule.

$$\frac{P_2 \preceq A_2^1 \wedge P_2 \preceq A_2^2 \wedge P_3 \preceq A_3^3 \quad P_1 \parallel A_2^1 \parallel A_2^2 \parallel A_3^3 \models \varphi_1}{P_1 \parallel P_2 \parallel P_3 \models \varphi_1} \quad (16)$$

The local results are then combined to constitute the global result of Eq.(9).

6 Case Study

In this section, we apply the compositional approach to an avionics system which combines the workload of [6] and the AFDX configuration of [13]. The workload consists of 5 partitions, and further divided into 18 periodic tasks and 4 sporadic tasks. Considering the inter-partition messages in the workload, we assign each message type $Msg_i, i = \{1, 2, 3, 4\}$ a separate VL with the same subscript. The messages of Msg_1 and Msg_2 are handled at the refresh period 50ms in sampling ports. Msg_3 and Msg_4 are configured to operate in queuing ports, each of which can accommodate a maximum of one message.

As shown in Fig.5, we consider the distributed architecture that comprises 3 ARINC-653 modules connected by an AFDX network. The module M_1 accommodates P_1 and P_2 , the module M_2 executes P_3 and P_5 , and the partition P_4 is allocated to M_3 . There are 4 VLs V_1-V_4 connecting 3 ESs across 2 switches S_1 and S_2 in the AFDX network. The arrows above VLs' names indicate the direction of message flow.

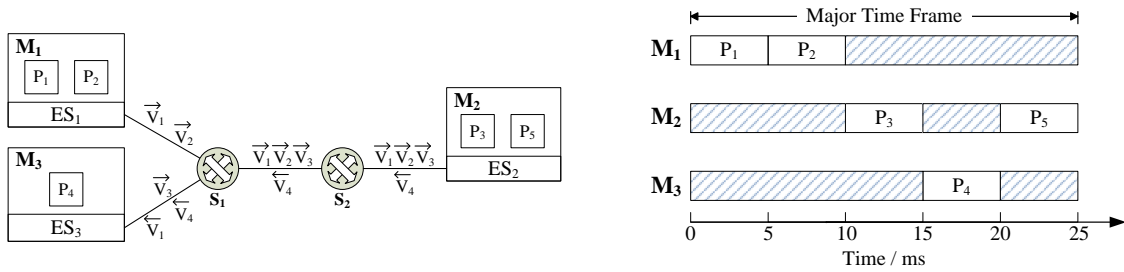


Figure 5: The Distributed Avionics Deployment and Partition Schedules (Times in Milliseconds)

The avionics system equips each of its processor cores with a partition schedule. Assume the modules in the experiment to be single-processor platforms. Fig.5 gives the partition schedules, which fix a common major time frame T_{mf} at 25ms and allocate 5ms to each partition within every T_{mf} . All the partition schedules are enabled at the same initial instant. The scheduling configuration keeps the temporal order of the partitions in [6]. Hence the partition schedules contain five disjoint windows $\langle P_1, 0, 5 \rangle$,

$\langle P_2, 5, 5 \rangle$, $\langle P_3, 10, 5 \rangle$, $\langle P_4, 15, 5 \rangle$, and $\langle P_5, 20, 5 \rangle$, where the second parameter is the offset from the start of T_{mf} and last the duration.

We analyze the schedulability of this avionics system following the procedure in section 5:

(1) *Decomposition*: The system is first decomposed into five sets of SWA template instances corresponding to five partitions. The schedulability of any partition $P_i, i = \{1, 2, 3, 4, 5\}$ is described as the UPPAAL query q_i :

$$A[] \text{ not perror}[i], \quad (17)$$

where the boolean variable `perror[i]` should be assigned to `True` once any error locations are reached in P_i . When analyzing the schedulability of P_i , we *only instantiate* the set of SWA template instances of P_i into UPPAAL processes. This set contains two scheduler models coming from `PartitionSupply` and `TaskScheduler`, all the `PeriodicTask` and `SporadicTask` models in P_i , and the communication layer models from which P_i receives messages.

(2) *Construction of message interfaces*: The message interfaces are constructed from the template depicted in Fig.4, for all the messages originate in periodic tasks. There are four unknown parameters `period`, `initOffset`, `offset`, and `jitter` in the template. Initially, the parameters of a message interface are set to the same values as these of the source task. Then we employ a binary search to heuristically refine `offset` and `jitter`, meanwhile guaranteeing timed selection simulation relation exists.

(3) *Model checking*: The schedulability of five partitions is checked individually. After combining the models of P_i and its message interfaces, we verify the property q_i by model checking in UPPAAL. The verification was repeated for each partition to evaluate the schedulability of a complete system. The experiment was executed on the UPPAAL 4.1.19 64-bit version and an Intel Core i7-5600U laptop processor.

(4) *Deduction*: According to the assume-guarantee rule described in Eq.(8), we conclude the schedulability of the complete system from the results of the verification of five partitions.

Results of the Analysis

The result in Table 1 shows that each partition is separately schedulable (The results “Yes” of Case 1) except the partition P_3 (The result “No”). From a global view, we cannot conclude directly that the system is non-schedulable, because the compositional approach described in section 5 only provides a sufficient condition for schedulability. Nevertheless, we find a counter-example by simulation in UPPAAL, and thus it can be concluded that the current system is not schedulable. The counter-example shows that P_3 violates the constraint of the refresh period of Msg_2 due to network latency.

Considering the effect of network latency on the scheduling configuration, we updated the partition schedules by performing a swap of time slots between P_1 and P_2 . The modified partition schedules provide five windows $\langle P_1, 5, 5 \rangle$, $\langle P_2, 0, 5 \rangle$, $\langle P_3, 10, 5 \rangle$, $\langle P_4, 15, 5 \rangle$, and $\langle P_5, 20, 5 \rangle$. The compositional analysis of the updated system was executed again. The result (Case 2 in Table 1) shows that all the partitions of the updated system are individually schedulable. Thus, the updated system finally achieves the schedulability at the global level.

Table 1 also shows the performance in terms of execution time and memory usage. In both cases, the partition P_3 contains more instantiated models (19 processes) than the other four partitions. As a result, model-checking runs evidently slower and requires more memory than the others. Nevertheless, the compositional analysis could be performed on ordinary computers within an acceptable time.

Compared with the compositional way, global analysis based on the same UPPAAL models would require 51 processes including all the 22 task models, whose state space is much more complex than the others. This causes UPPAAL to run out of memory within a few minutes, and thus makes the global

Table 1: The Experiment Results (Result), Execution Time (Time/sec.) and Memory Usage (Mem/MB)

No.	Case 1			Case 2		
	Result	Time	Mem	Result	Time	Mem
P_1	Yes	7.46	146	Yes	6.07	105
P_2	Yes	0.95	46	Yes	1.10	52
P_3	No	42.94	664	Yes	256.48	3041
P_4	Yes	0.69	43	Yes	0.68	43
P_5	Yes	19.41	509	Yes	128.56	2041

analysis infeasible. In contrast, the compositional approach only requires at most 5 task models when we perform model checking, offering effective state space reduction.

7 Conclusion

In this paper, we present a compositional approach for schedulability analysis of DIMA systems, which are modeled as a set of stopwatch automata in UPPAAL, describing schedulability as safety properties of models. We check each ARINC-653 partition including its communication environment individually, thereby reducing the complexity of model-checking. The techniques presented in this paper are applicable to the design of DIMA scheduling systems. We have applied the compositional approach to a concrete DIMA system. As future work, we plan to develop a model-based approach to the automatic optimization and generation of the partition schedules of a DIMA system.

References

- [1] AEEC (2010): *Avionics application software standard interface: part 1 - required services*. ARINC Specification 653P1-3, Aeronautical Radio Inc.
- [2] Tobias Amnell, Elena Fersman, Leonid Mokrushin, Paul Pettersson & Wang Yi: *TIMES: a tool for schedulability analysis and code generation of real-time systems*. In: *FORMATS 2003*, doi:10.1007/978-3-540-40903-8_6.
- [3] Björn Annighöfer & Frank Thielecke (2014): *A systems architecting framework for distributed integrated modular avionics*. DGLR, doi:10.1007/s13272-015-0156-1.
- [4] Jalil Boudjadar, Kim Guldstrand Larsen, Jin Hyun Kim & Ulrik Nyman: *Compositional schedulability analysis of an avionics system using UPPAAL*. In: *AASE 2014*.
- [5] Laura Carnevali, Giuseppe Lipari, Alessandro Pinzuti & Enrico Vicario: *A formal approach to design and verification of two-level hierarchical scheduling systems*. In: *RST 2011*, doi:10.1007/BF00360340.
- [6] Laura Carnevali, Alessandro Pinzuti & Enrico Vicario (2013): *Compositional verification for hierarchical scheduling of real-time systems*. *IEEE Transactions on Software Engineering* 39(5), pp. 638–657, doi:10.1109/TSE.2012.54.
- [7] Franck Cassez & Kim Larsen: *The impressive power of stopwatches*. In: *CONCUR 2000*, doi:10.1007/3-540-44618-4_12.
- [8] Franco Cicirelli, Angelo Furfaro, Libero Nigro & Francesco Pupo: *Development of a schedulability analysis framework based on pTPN and UPPAAL with stopwatches*. In: *DSRA 2012*, doi:10.1109/DS-RT.2012.16.

- [9] Alexandre David, Kim G Larsen, Axel Legay, Ulrik Nyman & Andrzej Wasowski: *Timed I/O automata: a complete specification theory for real-time systems*. In: *HSCC 2010*, doi:10.1145/1755952.1755967.
- [10] RB Dodd (2006): *Coloured petri net modelling of a generic avionics mission computer*. Technical Report, DTIC.
- [11] Arvind Easwaran, Insup Lee, Oleg Sokolsky & Steve Vestal: *A compositional scheduling framework for digital avionics systems*. In: *RTCSA 2009*, doi:10.1109/RTCSA.2009.46.
- [12] Orna Grumberg & David Long (1994): *Model checking and modular verification*. *Toplas* 16(3), pp. 843–871, doi:10.1145/177492.177725.
- [13] J Javier Gutiérrez, J Carlos Palencia & Michael González Harbour (2014): *Holistic schedulability analysis for multipacket messages in AFDX networks*. *Real-Time Systems* 50(2), doi:10.1007/s11241-013-9192-2.
- [14] Henrik Jensen (1999): *Abstraction-based verification of distributed systems*. Ph.D. thesis, Aalborg university.
- [15] Henrik Jensen, Kim Larsen & Arne Skou: *Scaling up UPPAAL*. In: *FTRFS 2000*, doi:10.1007/3-540-45352-0_4.
- [16] Marius Mikučionis, Kim Larsen, Jacob Rasmussen, Brian Nielsen, Arne Skou, Steen Palm, Jan Pedersen & Poul Hougaard: *Schedulability analysis using UPPAAL: Herschel-Planck case study*. In: *ISoLA 2010*, doi:10.1007/978-3-642-16561-0_21.
- [17] Jan Reineke, Björn Wachter & Stefan Thesing et al.: *A definition and classification of timing anomalies*. In: *WCET 2006*.
- [18] Youcheng Sun, Giuseppe Lipari, Romain Soulat, Laurent Fribourg & Nicolas Markey: *Component-based analysis of hierarchical scheduling using linear hybrid automata*. In: *RTCSA 2014*, doi:10.1109/RTCSA.2014.6910502.
- [19] Guoqing Wang & Qingfan Gu: *Research on distributed integrated modular avionics system architecture design and implementation*. In: *DASC 2013*, doi:10.1109/dasc.2013.6712647.