

Automatic Verification of Application-Tailored OSEK Kernels

Hans-Peter Deifel, Merlin Göttlinger,
Stefan Milius and Lutz Schröder

Friedrich-Alexander-Universität (FAU) Erlangen-Nürnberg
Email: {hans-peter.deifel, merlin.goettlinger,
stefan.milius, lutz.schroeder}@fau.de

Christian Dietrich and Daniel Lohmann
Leibniz Universität Hannover

Email: {dietrich, lohmann}@sra.uni-hannover.de

Abstract—The OSEK industrial standard governs the design of embedded real-time operating systems in the automotive domain. We report on efforts to develop verification methods for OSEK-conformant compilers, specifically of a code generator that weaves system calls and application code using a static configuration file, producing a stand-alone application that incorporates the relevant parts of the kernel. Our methodology involves two verification steps: On the one hand, we extract an OS–application interaction graph during the compilation phase and verify that it conforms to the standard, in particular regarding prioritized scheduling and interrupt handling. To this end, we generate from the configuration file a temporal specification of standard-conformant behaviour and model check the arising formulas on a labelled transition system extracted from the interaction graph. On the other hand, we verify that the actual generated code conforms to the interaction graph; this is done by graph isomorphism checking of the interaction graph against a dynamically-explored state-transition graph of the generated system.

1. Introduction

Embedded real-time control systems are special-purpose systems dedicated to specific, predefined tasks [1], [2]. Already now, a typical (in particular, non-autonomous) car contains up to a hundred such systems. Hence, both the hardware and the system software of each embedded control system need to be tailored to its specific needs in order to keep per-unit hardware costs as low as possible [3]. The OSEK-OS standard [4] fulfils these demands for tailorability and has been (together with its superset AUTOSAR-OS [5]) the dominant industry standard for event-triggered automotive *real-time operating systems* (RTOSs) for the last decades. What sets OSEK apart from the common POSIX-like operating systems is that it is completely statically configured. For a specific automotive application, all system objects (tasks, interrupt-service routines, resources, etc.) and their configurations have to be declared at compile-time in a domain-specific language, the *OSEK Implementation Language* (OIL) [6]. From this specification, an application-specific, highly optimized RTOS instance is derived by means of a generator.

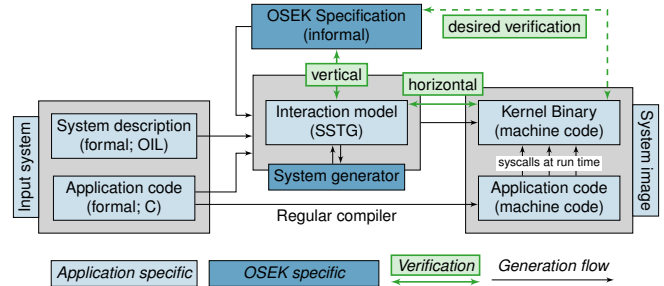


Figure 1: System generation and verification. We use the *static state-transition graph* (SSTG), a byproduct of the OSEK system generator, as an intermediate representation for our desired kernel verification.

However, with the advent of autonomous driving features, the industry is facing new challenges with respect to functional safety; the highest safety level ISO 26262 ASIL D demands the employment of a certified RTOS, such as *RTA-OS* (ETAS), *MICROSAR OS* (Vector), or *tresos Safety OS* (EB) (vendors named in parentheses). These certified operating systems offer significantly less tailorability and thus induce higher per-unit costs. The certification of the development process of the RTOS kernel is already extremely expensive, so vendors shy away from the even higher costs of certifying a kernel generator.

Taking a step back, we argue that application developers do not need a kernel that behaves correctly in all imaginable situations. However, they are very interested in a kernel that *always* behaves correctly for their specific application and all of their kernel-usage patterns. Therefore, we replace the isolated certification of the generator with a per-instance verification of the resulting kernel binary and thus formulate our verification goal: For a given application, the generated kernel binary must expose the specified behavior when executed together with our application. This bypass of the generator in the verification process allows all kinds of highly specialized system optimizations.

We achieve the desired verification (see Figure 1) by introducing an kernel–application interaction model: Our toolchain considers not only the OIL-specified system object instances but also how these system objects actually interact with each other via the syscall interface according to the

```

TASK(Low) {
  if (test()) {
    ActivateTask(High);
  }
  TerminateTask();
}

TASK(High) {
  TerminateTask();
}

void ISR() {
  ActivateTask(Med);
}

TASK(Med) {
  TerminateTask();
}

```

Figure 2: Example System – Source Code

OSEK semantics [7]. We enumerate the application-specific kernel’s state space, the *static state-transition graph* (SSTG), which is calculated and used by the system generator. The SSTG acts as an intermediate representation of kernel behavior and is the central data structure for our strategy. First, we statically verify (*vertically*, according to Figure 1) that the SSTG actually conforms to the OSEK standard. To this end, we formalize key aspects of the standard in CTL and model check the SSTG against this specification; this is feasible because the SSTG is of moderate size thanks to tailoring. Second, we dynamically verify the system against the SSTG (*horizontally*, according to Figure 1). To this end, we probe the system to explore a *dynamic state-transition graph* (DSTG) and check that it is isomorphic to the SSTG. We report on experiments with this methodology, both on systems from a standard test suite and on the control software of a quadrotor copter.

2. Background and Context

We give a brief overview of the single-core OSEK real-time operating system standard. Moreover, we describe the *static state-transition graph* (SSTG), which the dOSEK generator (dOSEK = dependable OSEK, our implementation of the OSEK standard) uses as an intermediate representation to model all possible interactions between application and kernel.

2.1. OSEK in a Nutshell. The tailored embedded systems that we are concerned with here are woven from application code and a tailored kernel instance. Kernel and application interact at runtime, typically subject to requirements on real-time performance. Depending on the current OS state, the kernel selects a control flow that is currently ready and dispatches it for execution. The application’s control flows, as the kernel’s counterpart, manipulate the OS state by invoking *system services* that influence the system behaviour (Table 1 gives a short overview). OSEK offers two main control flow abstractions: *interrupt-service routines* (ISRs) and *tasks* (i.e. threads). ISRs are activated by the hardware and fall into two classes: *category-1* ISRs, which are not allowed to call system services; and *category-2* ISRs, which are synchronized with the kernel. Tasks have a statically assigned priority, are allowed to use all system services, and are invoked according to a strict fixed-priority preemptive scheduling policy. On each new activation, tasks start from the very beginning and run until (self-)termination. Each task is configured to be either nonpreemptive (enforcing run-to-

completion semantics) or fully preemptive. Preemption points can be either *synchronous*, for example caused by an explicit activation of a higher priority task (ActivateTask), or *asynchronous*, if a higher priority task is activated inside an ISR. Periodically or aperiodically recurring task activations can be triggered by means of statically configured *alarms*, which are driven by a hardware timer.


Inter-task synchronization is realized by *resource* objects. Based on a *stack-based priority-ceiling protocol*, OSEK resources ensure mutual exclusion while preventing deadlocks and priority inversion. Through the acquisition of a resource, a task raises its *dynamic* priority to the *ceiling* priority of the resource – the highest static priority of all tasks that can obtain the resource, according to the OIL file.

Figure 2 shows a small example system that consists of three tasks and one ISR. These coordinate their execution with the help of the OS, which is activated through system service invocations (syscalls). Figure 3a depicts the same system as read in by the dOSEK generator.

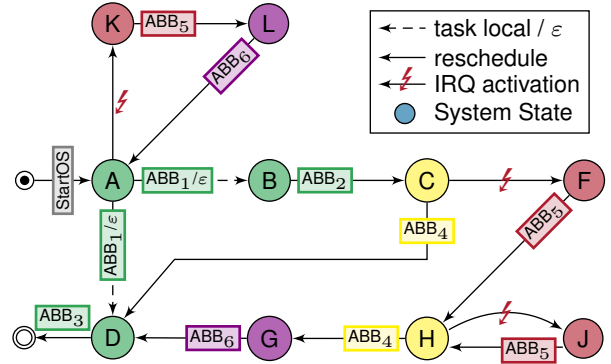
2.2. Generating a System. A dOSEK kernel instance is generated from two inputs (see also Figure 1): The application’s OIL file specifies the employed RTOS objects (i.e., the tasks *Low*, *Med*, *High* and the ISR *ISR* in our example). The application’s source code (Figure 2) specifies how these system objects interact according to the OSEK semantics.

Internally, we structure the application code into a set of *control-flow graphs* (CFGs) consisting of *atomic basic blocks* (ABBs) (Figure 3a). An ABB [7], [8] is a control-flow superstructure that subsumes one or more traditional *basic blocks* (BBs) forming a single-entry single-exit region; it has *exactly one* distinguished entry BB and one exit BB. The construction (see [7] for a detailed description) results in one ABB-graph for each task within the application code. By construction, every ABB either contains a single syscall or only computation code that does not interact with the OS (no syscalls). From the kernel’s point of view, an ABB executes atomically, but can be interrupted by ISRs.

At build time, the dOSEK generator computes a *state transition graph* (STG) from the ABB graphs of the individual tasks and the system configuration (OIL). (We formally define STGs in Section 2.3.) This STG is the *static state-transition graph* (SSTG) already mentioned in the introduction. Starting with an initial global system state, derived from the OIL file, the generator enumerates all reachable system states explicitly (Figure 3b). Every state carries the currently running task, a block that is executed in this state (e.g. state A executes ABB_1), and other relevant scheduling data, like the list of activated tasks. A state transition is caused by either a computation block, a system call block, or an interrupt request. For the associated post-state, the currently running task and the currently executed block are calculated according to the OSEK scheduling semantics. The SSTG subsumes the interwoven application–kernel behaviour and includes all possible scheduling sequences an OSEK kernel exposes for the given application. Since *computation* blocks do not perform syscalls, their execution does not influence future scheduling decisions, and is therefore represented



(a) The application is structured into three tasks with high, medium, and low priority, respectively, and one interrupt service routine. The application code is partitioned into *atomic basic blocks* (ABBs). In ABB_2 , the low-priority task activates the high-priority task via `ActivateTask`. The ISR activates the medium-priority task. (To simplify the example, the IRQ may happen only in ABB_1 and ABB_4 , and the tasks issue `TerminateTask()` implicitly.)



(b) This *static state-transition graph* (SSTG) captures all possible state transitions from the start to the termination of the low-priority task according to the OSEK semantics. Every state carries and executes exactly one ABB.

Figure 3: Example system

TABLE 1: Excerpt of system services provided by the OSEK-OS API.

System Service	Arguments	Brief Description
<code>ActivateTask</code>	<code>TaskID</code>	<code>Task TaskID</code> is activated. If the current task is preemptible, immediate rescheduling takes place.
<code>TerminateTask</code>	—	The current task terminates itself and immediate rescheduling takes place.
<code>GetResource</code>	<code>ResID</code>	Acquires the resource identified by <code>ResID</code> .
<code>ReleaseResource</code>	<code>ResID</code>	Leaves the critical region associated with the resource <code>ResID</code> . The dynamic priority of the calling task is changed and a reschedule takes place for preemptible tasks.

by ϵ -transitions. For instance, ABB_1 is computational and, therefore, the transition between states A and D becomes an ϵ -transition.

Thus computed, the SSTG has node identities representing the current global system state, while the edges of the SSTG are either ϵ -transitions or labelled with the system calls triggering the respective state transitions. System call labels can be either system calls in the proper sense, interrupts (which are triggered by the hardware outside the system), interrupt returns (irets), or the idle system call. The idle system call is executed from an idle state when no other task is ready for execution. This ensures that all maximal paths in the SSTG are infinite. The latter three types of system call labels are artificially added by the generator and only model the implicit state transitions of an OSEK system; they are not explicitly named in the OSEK specification.

From the SSTG information, the dOSEK system generator (Figure 1) produces a kernel binary that is optimized for the actual application usage patterns. Optimizations include, for example, the avoidance of scheduler invocations for system call sites that have a known scheduling outcome.

2.3. State Transition Graphs. Our verification method is concerned with properties of and relations between STGs, which in process-theoretic parlance are essentially deterministic labelled transition systems. That is, given a set \mathcal{A} of *labels* a state transition graph consists of a set S of *states* and a set $T \subseteq S \times \mathcal{A} \times S$ of *labelled transitions*. We write $s \xrightarrow{a} t$ for $(s, a, t) \in T$. The transition relation is required to be *deterministic* (but may be, and typically is, *partial*), that

is, whenever $s \xrightarrow{a} t$ and $s \xrightarrow{a} t'$ then $t = t'$. We occasionally consider state transition graphs with ϵ -*transitions*; these additionally admit transitions of the form $s \xrightarrow{\epsilon} t$ where ϵ is a special label not contained in \mathcal{A} . For ϵ -transitions, we do not require determinism. For example, the graph from Figure 3b is nondeterministic at state A w.r.t. ABB_1 -transitions, which are replaced by ϵ -transitions as explained in Section 2.2.

To STGs with ϵ -transitions, we apply the usual process of ϵ -*elimination*, i.e. we insert an a -transition from state s to state t whenever t is reachable from s by first performing any number of ϵ -transitions (possibly none) and then an a -transition. We then remove all unlabelled transitions, and all states that become unreachable as a result. In general, this will produce a nondeterministic STG; we will explain in Section 3.2 why the particular STGs that appear in our verification framework do remain deterministic after ϵ -elimination.

3. The Formal Verification Method

Our formal verification methodology comprises a *vertical* and a *horizontal* verification process (cf. Figure 1). The central data structure for our verification is the SSTG (Section 2.2). In the *vertical verification*, we check that the SSTG adheres to key aspects of the behaviour specified by the OSEK standard, in particular regarding prioritized scheduling and interrupt handling. To this end, we formalize the corresponding parts of the OSEK standard in CTL and model check the SSTG against the arising temporal specification (Section 3.1). In the horizontal verification, we then ensure that the actual

generated code conforms to the originally projected system-wide control flow; this is achieved by graph isomorphism checking of the SSTG against a judicious abstraction of the code, viz. another STG called the *dynamic state-transition graph* (DSTG) (Section 3.2).

3.1. Vertical Verification. We next describe how we formally verify that the static state transition graph (SSTG) complies with the OSEK specification. To this end, we generate a NuSMV-model and CTL formulas that formalize (parts of) the OSEK specification. Our formalization currently covers the standard roughly up to conformance class ECC1, with the exception of alarms and resource management. As input for our generator we use the (ϵ -eliminated) SSTG and the OIL specification, which specifies all tasks and interrupts with their respective priorities, as well as all events and resources of the system.

The SSTG already resembles a NuSMV-model, except that NuSMV-models do not have edge labels. Thus, we need to convert the edge labels into state labels, and our generator does that by pushing labels along the arrows into the next state. This leads to a multiplication of states by the number of different labels of incoming edges. Of course, this conversion produces a nondeterministic model, however with a unique edge between two states as every syscall has a unique effect on the current state. In more detail, our NuSMV-model has as global state variables (a) the variable `syscall` that contains the name of the system call (i.e. the label of the edge) that took the system into the current state, (b) a variable `state` carrying the current state in form of the node id from the SSTG, and (c) variables for all other parts of the global system state. Note that the values of the latter variables are already determined uniquely by the node id; we included them merely to make the CTL formulas and counterexample traces more readable.

All OS objects, i.e. tasks, ISRs, events, and resources are realized by means of NuSMV-modules, which are instantiated as specified in the OIL specification of the system. They do not carry internal state but are merely used to group related variables for readability in formulas and error traces. For example, the state and (dynamic) priority of a task t are referred to by $t.state$ and $t.priority$. The actual state is fed into the instances through global variables that form parameters of the modules. As the SSTG contains an artificial first state where nothing is running and the only outgoing edge is labelled with the `StartOS` syscall, we eliminate this state and initialize `syscall` with `StartOS` and `state` with the (node id of the) target node of the eliminated edge labelled `StartOS` in the SSTG. This is sound because nothing is specified to happen in the system before the `StartOS` transition. The next value of state at any state is chosen nondeterministically as one of the successor nodes of the current SSTG node (given by the current value of state). The next(`syscall`) is then uniquely determined by the current value of state and the value of next(`state`).

Some variables, for example the resource priorities, are not explicitly contained in the OIL file and thus have to be calculated according to the OSEK specification. The latter

demands that resource priorities are at least the maximum priority of all tasks using the resource but are lower than the priority of every task not using the resource but having higher priority than the ones using the resource. To avoid priority collisions between tasks due to resource occupation, we scale all priorities by a factor of two and calculate the resource priorities as the maximum of all the priorities of the tasks using the resource plus one. This does not affect the scheduling behaviour and ensures that there can never be a situation where multiple tasks of the same dynamic priority are ready to run at the same time.

Figure 4 shows a slice of the NuSMV-model generated from the SSTG for the example system from Figure 2 (depicted in simplified form in Figure 3(b)). We include only the variables `state`, `syscall` and `running` and only the transitions from the starting state `ABB_67_0`, which corresponds to state A. The transition to `ABB_4_0` corresponds to the transition $A \xrightarrow{\epsilon} B \xrightarrow{ABB_2} C$ and the one to `ABB_63_0` corresponds to the transition to the final state with label `ABB_3`. The remaining three transitions essentially correspond to $A \rightarrow K$ but take into account that in reality an interrupt in the program of Figure 2 may happen before `test()` is executed or at the two points after the branch (i.e. just before `ActivateTask(High)` or just before `TerminateTask()`).

We verify that the SSTG adheres to the OSEK specification by model checking the NuSMV-model generated from the SSTG against the CTL formulas generated from the OIL file according to the OSEK specification. The latter arise by instantiating formula patterns that are parametric in the OIL configuration. Example properties are shown in Table 3. We frequently need to quantify over finite sets that are obtained from the OIL specification or are specified by OSEK (Table 2); such quantifiers are just expanded into finite conjunctions or disjunctions.

Formula (1) ensures that the event state changes in the specified way. The first conjunct states that if task o is newly started (rather than resumed), i.e. its state changes from `Suspended` (rather than `Ready`) to `Ready` or `Running`, the event state of event e owned by o is cleared. The second conjunct specifies that the owner, if running, can clear the event and that it will otherwise remain set. The last conjunct states that the event can only be set with the specified syscalls.

Formula (2) formalizes how the state of a resource has to behave in reaction to the observed system calls. A normal resource changes from being unoccupied to being occupied iff a `GetResource` system call is observed. The `ReleaseResource` system call has the opposite effect. If the resource is internal, then “it is automatically taken when the task enters the running state, [and] at points of rescheduling [...] it is automatically released.” [4, p. 34]. Possible rescheduling points are *waiting* for an event and task *termination*. Note that we use the internal resource mechanism with a resource owned by all tasks to model non-preemptible tasks.

Formula (3) states that the execution state (comparable to the instruction pointer of a program) of a control flow r can only change when r is currently running. This ensures

```

MODULE main()
VAR
  ...
  syscall : { Start, ..., TerminateTask, ...,
             ActivateTask_High, ..., interrupt_37, ... };
  running : { Idle, Low, Med, High, ISR };
  state : { ABB_67_0, ..., ABB_4_0, ..., ABB_23_0, ABB_24_0,
           ABB_25_0, ..., ABB_63_0, ... };
  ...
ASSIGN
  init(syscall) := Start;
  next(syscall) := case
    ((state = ABB_67_0) & next((state = ABB_4_0))) :
      ActivateTask_High;
    ((state = ABB_67_0) & next((state = ABB_23_0))) :
      interrupt_37;
    ((state = ABB_67_0) & next((state = ABB_24_0))) :
      interrupt_37;
    ((state = ABB_67_0) & next((state = ABB_25_0))) :
      interrupt_37;
    ((state = ABB_67_0) & next((state = ABB_63_0))) :
      TerminateTask;
    ...
  next(TRUE) : syscall;
esac;
init(running) := Low;
next(running) := case
  next(... | (state = ABB_4_0)) : High;
  next(... | (state = ABB_23_0) | (state = ABB_24_0) |
         (state = ABB_25_0) | ...) : ISR;
  next((state = ABB_63_0)) : Idle;
  ...
  next(TRUE) : running;
esac;
init(state) := ABB_67_0;
next(state) := case
  ...
  (state = ABB_67_0) : {ABB_4_0, ABB_23_0, ABB_24_0,
                      ABB_25_0, ABB_63_0};
  ...
esac;
...
CTLSPEC ...

```

Figure 4: Slice of the NuSMV-model for the graph in Fig. 3

that an interrupted or otherwise preempted control flow will continue executing where it was stopped earlier.

Formula (4) is a trivial property stating that automatically started tasks do indeed start at the beginning.

Formula (5) expresses that the global interrupt state is respected: either all interrupts are enabled, or only category 2 interrupts are disabled, or all interrupts are disabled, and only enabled interrupts are allowed to start. Note that one cannot simply ask that for a disabled ISR x one has $x.state \neq \text{Running}$ because x might change the global interrupt state (e.g. disable all interrupts) while it is running.

Formula (6) states that the global interrupt state does not make any forbidden transitions.

Formula (7) formalizes that interrupts are only allowed to make certain system calls. Note here that $A[\varphi W \psi]$ is the usual weak until connective, read ‘ φ holds forever or until ψ ’. The set of allowed calls differs between category 1 and 2 interrupts. This is captured by $AC(i)$ (see Table 2).

Finally, Formula (8) expresses that the scheduling for a task t is correct, i.e. each transition made in the SSTG is legal, and every transition required by the specification is in fact made by the system. Note that the transitions from Waiting and Suspended directly to Running are not technically legal according the OSEK specification; but

the required intermediate state where the task in question would be Ready is not observable to the application, and we therefore opt to allow the direct transitions. The formula is instantiated for every task, and, in a slightly modified version, for every ISR, as interrupt scheduling is to some extent left up to the implementation. The formula is structured into four subformulas ψ_1, \dots, ψ_4 , each handling the allowed state transitions from one of the four possible starting states. To make this rather large formula readable we employ the following abbreviations for subformulas and properties:

- $t.isHighestPriority$ specifies that control flow t wants to run, e.g. has state Ready or Running, and has higher priority than every other control flow that currently wants to run.
- $t.isWaitingFor(e)$ states that the task t is waiting for the event e to occur.
- $allOthersPreemptible(t)$ states that all tasks currently ready, other than t , are preemptible. Note that this formula needs to be used together with $t.isHighestPriority$ to ensure that no interrupt is currently running and is needed altogether because when checking the starting transition of a non-preemptible task t , knowing $t.isHighestPriority$ alone is not enough. For example, before the transition, $t.isHighestPriority$ would always be false, since t was suspended, and after the transition, t would have already received the ceiling priority of the internal scheduler resource due to being non-preemptible.
- $othersWillPreempt(t)$ denotes that task t will be preempted by another control flow due to witnessing a scheduling system call. This happens only if another control flow s currently has the highest priority, and if s is not an ISR, then t also has to be preemptible.
- $waitSc(t)$ formalizes that task t is waiting for at least one event that is not set.

Note that the parameter e in $t.isWaitingFor(e)$ is realized in our NuSMV-model as follows: For every task, an array containing the states of all its events is generated. The parameter e then simply is an index into this array.

Formula ψ_1 handles the case where t was previously suspended. Here, the only way for t to start up again is by a system call from the set Act (with a transition to Ready or Running depending on what else is currently running). In the case where t was previously Running, ψ_2 captures the possible transitions:

- (1) There is another control flow that will preempt t ($othersWillPreempt(t)$) due to a system call that causes rescheduling, and t will become Ready.
- (2) t issues a system call from the set $TCS(s)$ signalling its termination, and becomes Suspended.
- (3) t waits for one of its events. If the event is not set, t becomes Waiting.

In the case where t was previously Ready, ψ_3 formulates that t will either remain Ready, or become Running if it has the highest priority. Finally, in the case where t was previously Waiting, ψ_4 expresses that t keeps Waiting until one of the events it was waiting for is set and then becomes

TABLE 2: Finite sets used in the OSEK formalization.

Set	Name/Description
SC	Scheduling calls, i.e. system calls that lead to a (re-)scheduling of tasks
Int	All (category 1 and category 2) interrupts
Int _{<i>i</i>}	Category <i>i</i> interrupts where <i>i</i> = 1 or 2
AC(<i>i</i>)	Allowed calls, i.e. system calls that are allowed for interrupt <i>i</i> to make (determined by its category)
Act(<i>r</i>)	System calls that activate the control flow (i.e. task or ISR) <i>r</i>
E(<i>t</i>)	Events of task <i>t</i>
clearCalls(<i>e</i>)	Syscalls that lead to the event <i>e</i> being cleared
setCalls(<i>e</i>)	Syscalls that lead to the event <i>e</i> being set
ABBs(<i>r</i>)	ABBs, i.e. states of execution, of the control flow <i>r</i>
TSC(<i>r</i>)	Terminating system calls, i.e. system calls that would lead to the termination of <i>r</i>
Tasks	The set of all tasks
NPTasks	Non-preemptible tasks, i.e. tasks that cannot be preempted by higher priority tasks
RPreempt(<i>r</i>)	Control flows that could preempt <i>r</i>
WaitCalls	WaitEvent System calls
TIntRes(<i>r</i>)	Tasks that have <i>r</i> as internal resource

```

TASK(Low) {
  print_state_hash(at: ABB1, next: interrupt);
  trigger_interrupt();
  if (read_decision(0)) {
    print_state_hash(at: ABB2, next: ActivateTask);
    ActivateTask(High);
  }
  print_state_hash(at: ABB3, next: TerminateTask);
  TerminateTask();
}

```

Figure 5: Example system – generated mockup for task Low

either Ready or Running, depending on whether it currently has the highest priority.

3.2. Horizontal Verification. To increase trust in the correctness of the actual generated code, we complement the verification that the SSTG complies with the OSEK specification (Section 3.1) with a verification procedure ensuring that the the generated code agrees with the SSTG. To this end, we extract a normalized STG, the *dynamic state transition graph (DSTG)*, from the actual binary. Interestingly, while one might expect the notion of agreement of the DSTG with the SSTG to be based on classical process-algebraic notions of equivalence such as bisimilarity [9], it turns out that the normalization process in fact guarantees agreement of the two STGs up to isomorphism. We therefore base our verification procedure on isomorphism checking, not only because we thus obtain stronger correctness guarantees but also because isomorphism checking of deterministic systems is computationally cheap, and in fact can be performed on-the-fly.

In more detail, we extract the DSTG by executing and probing the generated system binary with all possible syscall sequences that can originate from the given application. In order to facilitate exploration of the state space, we transform the tasks’ ABB graphs (see Section 2.2) into a *mock-up*, a C-program that omits the processing logic of the application and retains only the control flow, and then run an external search procedure on the mock-up that traverses the state space depth-first. We generate the mock-up (see Figure 5 for a partial

mock-up of the running example) in two steps: (1) We use tools from the dOSEK framework to generate function-local ABB graphs from the generated LLVM code. (2) From these ABB graphs, we generate C-code that emulates the control flow, i.e. performs function and system calls as specified. Additionally, the mock-up:

- outputs node identifiers containing the identifier of the current ABB, as well as a hash of the actual current operating system state, where the latter includes the program counter;
- outputs identifiers for system calls containing the name of the system routine as well as the call site;
- optionally triggers any enabled interrupts;
- reads decisions on branching (including whether to trigger an interrupt) from standard input.

The mock-up is linked with the specialized kernel produced for the actual application by the dOSEK generator. It is then used by the *Dynamic State Explorer (DSE)*, a search procedure that generates the space of reachable states depth-first, steering the mock-up through the state space by feeding input to it in order to determine branching. The result of the search procedure is an STG. More precisely, some transitions in the STG are labelled with system calls as indicated above, and some are unlabelled, i.e. the STG includes ε -transitions; the latter correspond to internal transitions between computational ABBs. The ε -transitions are nondeterministic, since we omit the processing logic in the mock-up, so any form of conditional branching in the original application turns into nondeterminism; also, we cannot foresee external input. This STG is then subjected to ε -elimination as described in Section 2.3. We thus generate an STG with only labelled transitions; it is this STG that we refer to as the dynamic state transition graph (DSTG). The DSTG, as well as the SSTG, is deterministic, since every non- ε label is a tuple (*call site, syscall type*) and as such deterministically changes the system state. The same holds for interrupts, even though these are nondeterministic with regard to the activation time: Every interrupt transition label contains the interrupt number and, therefore, exactly describes its influence on the OS state, just like every other syscall-induced transition.

We check the DSTG for isomorphism with the SSTG. This is computationally unproblematic: since both LTS are deterministic after ε -elimination, we only need to check that both sides allow for the same transition labels, and then propagate this property to the states reached by the corresponding transitions. The reason that both graphs are isomorphic is that the states of the DSTG are identified by the hash value over the OS state, which is also reflected in the fields of every SSTG node.

4. Experiments

4.1. Positive Tests. To evaluate our verification method, we have run experiments on a number of OSEK systems. These systems stem from a test suite originally designed for the dOSEK implementation. In total, we have fully verified 58 test systems. We have selected eight systems that highlight

TABLE 3: Formalization of OSEK requirements.

Property	Formula	
Event state changes of event e owned by owner o	$\text{AG}((o.\text{state} = \text{Suspended} \rightarrow \text{AX}((o.\text{state} = \text{Ready} \vee o.\text{state} = \text{Running}) \rightarrow \neg e.\text{set}))$ $\wedge ((o.\text{state} = \text{Running} \wedge e.\text{set}) \rightarrow \text{AX}(\neg e.\text{set} \leftrightarrow \text{syscall} \in \text{clearCalls}(e)))$ $\wedge (\neg e.\text{set} \rightarrow \text{AX}(\text{syscall} \in \text{setCalls}(e) \leftrightarrow e.\text{set}))$	(1)
Resource state changes of r	$(\neg r.\text{isInternal} \wedge \text{AG}((r.\text{occupied} \rightarrow \text{AX}(\neg r.\text{occupied} \leftrightarrow \text{syscall} = \text{ReleaseResource}(r))))$ $\wedge (\neg r.\text{occupied} \rightarrow \text{AX}(r.\text{occupied} \leftrightarrow \text{syscall} = \text{GetResource}(r))))$ $\vee (r.\text{isInternal} \wedge \text{AG}((r.\text{occupied} \rightarrow \bigwedge_{t \in \text{TIntRes}(r)} \text{AX}(\neg r.\text{occupied} \leftrightarrow$ $((t.\text{state} = \text{Suspended} \vee t.\text{state} = \text{Waiting}) \wedge \bigwedge_{ot \in \text{TIntRes}(r) \setminus \{t\}} \neg ot.\text{state} = \text{Running})))$ $\wedge (\neg r.\text{occupied} \rightarrow \text{AX}(r.\text{occupied} \leftrightarrow \bigvee_{t \in \text{TIntRes}(r)} t.\text{state} = \text{Running})))$	(2)
Execution context of control flow r only changes when executing	$\text{AG} \bigwedge_{a \in \text{ABBs}(r)} (r.\text{abb} = a \rightarrow \text{A} [r.\text{abb} = a \text{ W } r.\text{state} = \text{Running}])$	(3)
Autostart of task t is respected	$t.\text{autostart} \leftrightarrow t.\text{state} \in \{\text{Ready}, \text{Running}\}$	(4)
Interrupt state is respected	$\text{AG} \left((\text{irsIS} = \text{Disabled} \rightarrow \bigwedge_{x \in \text{Int}} \neg \text{startable}(x)) \wedge (\text{irsIS} = \text{Cat2Disabled} \rightarrow \bigwedge_{x \in \text{Int}_2} \neg \text{startable}(x)) \right)$	(5)
Interrupt state transitions	<p>where $\text{startable}(x) \equiv x.\text{state} \neq \text{Running} \wedge \text{EX } x.\text{state} = \text{Running}$</p> $\text{AG } \text{irsStays}(\text{Disabled}, \{\text{EnableIR}, \text{ResumeIR}, \text{ResumeOSIR}\})$ $\wedge \text{irsStays}(\text{Enabled}, \{\text{DisableIR}, \text{SuspendIR}, \text{SuspendOSIR}\})$ $\wedge \text{irsStays}(\text{Cat2Disabled}, \{\text{ResumeOSIR}\})$ <p>where $\text{irsStays}(s, \text{sys}) \equiv \text{isrIS} = s \rightarrow \text{A} [\text{isrIS} = s \text{ W } \text{syscall} \in \text{sys}]$</p>	(6)
ISR i allowed calls	$\text{AG } i.\text{state} = \text{Running} \rightarrow \text{A} \left[\bigvee_{s \in \text{AC}(i)} \text{AX } \text{syscall} = s \text{ W } (i.\text{state} \neq \text{Running}) \right]$	(7)
State transitions of task t	$\text{AG}((t.\text{state} = \text{Suspended} \rightarrow \psi_1) \wedge (t.\text{state} = \text{Running} \rightarrow \psi_2)$ $\wedge (t.\text{state} = \text{Ready} \rightarrow \psi_3) \wedge (t.\text{state} = \text{Waiting} \rightarrow \psi_4))$ <p>where</p> $\psi_1 \equiv (\text{allOthersPreemptible}(t) \rightarrow \text{AX}(\text{syscall} \in \text{Act}(t) \rightarrow (t.\text{state} = \text{Running} \leftrightarrow t.\text{isHighestPriority})))$ $\wedge (\neg \text{allOthersPreemptible}(t) \rightarrow \text{AX}(\text{syscall} \in \text{Act}(t) \leftrightarrow t.\text{state} = \text{Ready}))$ $\wedge \text{AX}((t.\text{state} = \text{Suspended} \rightarrow \neg \text{syscall} \in \text{Act}(t)))$ $\wedge \text{AX}((t.\text{state} = \text{Ready} \vee t.\text{state} = \text{Running}) \rightarrow \text{syscall} \in \text{Act}(t)))$ $\psi_2 \equiv \text{AX}((t.\text{state} = \text{Ready} \leftrightarrow \text{othersWillPreempt}(t)) \wedge (t.\text{state} = \text{Suspended} \leftrightarrow \text{syscall} \in \text{TSC}(t))$ $\wedge (t.\text{state} = \text{Waiting} \leftrightarrow \text{waitSc}(t)))$ $\psi_3 \equiv (\text{allOthersPreemptible}(t) \rightarrow \text{AX}(\text{syscall} \in \text{SC} \rightarrow (t.\text{state} = \text{Running} \leftrightarrow t.\text{isHighestPriority})))$ $\wedge \text{AX}(t.\text{state} = \text{Running} \vee t.\text{state} = \text{Ready})$ $\psi_4 \equiv \bigwedge_{e \in \text{E}(t)} (t.\text{isWaitingFor}(e) \rightarrow$ $((\text{allOthersPreemptible}(t) \rightarrow \text{AX}(e.\text{set} \rightarrow (t.\text{state} = \text{Running} \leftrightarrow t.\text{isHighestPriority})))$ $\wedge (\neg \text{allOthersPreemptible}(t) \rightarrow \text{AX}(e.\text{set} \rightarrow t.\text{state} = \text{Ready}))$ $\wedge \text{AX}((t.\text{state} = \text{Waiting} \wedge \neg e.\text{set}) \vee t.\text{state} = \text{Running} \vee t.\text{state} = \text{Ready})))$ <p> $\text{allOthersPreemptible}(t) \equiv \bigwedge_{ot \in \text{NPTasks} \setminus \{t\}} ot.\text{state} \neq \text{Running}$ $\text{othersWillPreempt}(t) \equiv \text{syscall} \in \text{SC} \setminus (\text{TSC}(t) \cup \text{WaitCalls}) \wedge \bigvee_{or \in \text{RPreempt}(t) \setminus \{t\}} or.\text{isHighestPriority}$ $\text{waitSc}(t) \equiv \bigvee_{e \in \text{E}(t)} \neg e.\text{set} \wedge t.\text{isWaitingFor}(e)$ </p>	(8)

key properties targeted in the verification (Table 4). The test cases “copter” and “copter-small” are the control software of a quadrotor copter and a simplified version thereof that arises by removing one asynchronous signal. For each of the systems, we list the number of system objects of the relevant types, i.e. interrupts, tasks (including the idle task), events, and resources specified in the OIL file.

Table 4 shows key parameters and the performance of the model checking tool on those systems. To qualify the generated NuSMV-model, we give the number of reachable states as well as its diameter (i.e. the length of the longest loop-free path). On a 2.4 GHz Intel Core i7-5500U machine with 8 GB of memory, the smaller experiments were completed within a fraction of a second. Only the verification of the two copter examples took longer, but finished in under

three minutes.

For the horizontal verification, we probed the same 58 OSEK systems in two generator configurations (with and without system call site specialization) and established isomorphism with the respective SSTG in all cases. For most systems, the horizontal verification took less than 1 second. Only for the copter, the probing took 2.13s (0.81s for copter-small) and the isomorphism checking 0.17s (0.04s). When developing the hash function, we have probed examples with over 400.000 states in under 4 minutes.

4.2. Negative Tests and Fault Injection. For the vertical verification, we have performed *negative tests* by introducing faults into systems to check that these are correctly identified by our verification tool chain. To this end we have implemented a modified dOSEK generator that injects various

TABLE 4: Performance of vertical verification.

Name	ISRs/Tasks/ Events/Res	Reachable states	Diameter	User time (sec)	Memory (MB)
bcc1-resource1j	0/6/0/4	26	16	0.10	27
bcc1-sse1c	0/6/0/4	24	14	0.08	28
ecc1-bt1g	0/6/2/2	10	10	0.05	24
ecc1-event1e	0/4/4/2	13	13	0.11	29
bcc1-isr2d	1/4/0/2	21	10	0.07	23
timing-abcomp	1/3/2/2	77	13	0.14	27
copter-small	3/11/0/3	1366	29	19.44	250
copter	4/12/0/3	4458	32	147.15	829

types of faults into the input for our verification method. One can select between (a) mutations in the SSTG and (b) mutations in the input OIL specification. For (a), there is a random choice of either adding an edge or merging two states, and for (b), there is a random choice of either exchanging the priorities of two tasks or toggling the preemptability or the auto-start flag.

It should be noted here that not all faults introduced in this random way will necessarily lead to actual errors, e.g. when the change to the configuration does not influence the actual interaction between application and OS, or when additional transitions are actually valid. In our experiments, additional edges mostly did lead to errors and were detected by the formal verification; in some cases, additional edges produced legal transitions or violated parts of the specification not currently reflected in our formalization, e.g. that a task is not allowed to release resources it has not currently reserved. Merging graph nodes almost always produced errors caught by the verification, except in cases where the net effect would have been produced by ε -elimination anyway.

For the vertical verification, we injected 188 different faults into the test cases; this did not change the performance of the formal verification significantly compared to the unmodified test cases (Table 4). 177 faults lead to errors and were detected by the vertical verification. The other 11 faults were manually verified to be benign in the given usage pattern (a more detailed discussion can be found in the full version [10]). For the horizontal verification, we inserted 81 OIL-level faults: 61 lead to detected errors, while the other faults were manually checked to be benign.

4.3. Lessons Learned. Summing up the experience gained, it seems possible to achieve a fair degree of coverage in the verification of key aspects of the OSEK specification in tailored systems. (Unbounded) CTL model checking is feasible as the reachable part of the abstracted state space that we use remains within tractable range even for fairly large systems such as the quadrotor copter controller (with fewer than 4.5k states after abstraction); as stated above we attribute this fact to OS tailoring. Without wishing to get involved in the long-lasting linear-vs-branching-time war (e.g. [11]), we note that at the scale of our examples, LTL model checking does appear to reach the frontier of feasibility. For example, while in CTL, the copter-small example (Table 4) was discharged in under twenty seconds using less than 300 MB of memory, on the LTL correspondent of essentially the same specification we stopped the model

checker after 30 minutes at 3.5 GB of allocated memory. This may be due to the higher formula complexity of LTL model checking (PSPACE instead of PTIME). A clear disadvantage of CTL, on the other hand, is that counterexamples are less informative, and typically stop at the first nested path quantifier.

Another somewhat surprising aspect is the fact that the notion of correspondence between the SSTG and the DSTG has turned out to be isomorphism of STGs. Actually getting this insight to work in full has required a somewhat laborious tuning process regarding the hashing of the OS state in the exploration of the DSTG, and in fact maintaining the tool chain in the future might be easier if one replaces isomorphism with strong bisimilarity.

5. Related Work

Our work is set in the highly active area of software model checking; see [12] for an (admittedly dated) overview. Our method exploits the high degree of predictability of scheduling afforded by OSEK, and in particular avoids the state space explosion caused by thread interleaving [13]. The static generation of state transition graphs from code in a somewhat similar style as featured in our approach has been used, in combination with LTL model checking, in the verification of event-condition-action systems [14].

The OSEK standard has been the subject of formal verification efforts to some degree. Waszniowski [15] modelled OSEK using timed automata within the UPPAAL model checker, and performed schedulability analyses. Huang et al. [16] modelled OSEK in CSP to verify various properties such as deadlock freedom. In this model, the internal application structure is not considered, and interrupts are excluded entirely. Vu et al. [17], [18] formalize the OSEK standard in Event-B and then verify designs of full OSs against the formalization. Where applications are considered [19], [20], these are verified in connection with an OS model rather than the actual OS implementation. In contrast, our approach avoids the verification gap between OS model and OS implementation by verifying the entire system composed of application and OS. This is made possible by focusing on the part of the OS behaviour actually relevant for the application at hand, instead of attempting to verify the full OS. We are thus able to a) work on the actual implementation, and b) fully model check the entire application/OS system including interrupts (expressly not covered in cited work on verifying OSEK applications).

Zhang et al. [21] formalize the OSEK standard in the K framework, along with the OIL and the programming language for applications. This is then used for test case generation and to verify applications by symbolic execution within the model. Interrupts are not considered.

Tigori et al. [22] use reachability checking of *extended finite automata* to remove dead code in tailored OSEK systems; the automata models involved are produced manually, while we generate STGs during code generation and from the actual code, respectively. Also, verification of OSEK

adherence in [22] is by standardized testing, while we model-check the formalized standard.

On an entirely different scale, Klein et al. [23] formally verified the seL4 microkernel for functional correctness, in a project of 25 person years, and Sewell et al. [24] extended this verification from the C-Code level to the binary.

6. Conclusions

We have presented a framework for the fully automatic lightweight verification of tailored embedded systems following the OSEK industrial standard. Specifically, we have introduced a *vertical* verification process whereby the statically generated control flow of the tailored system is checked for conformance with the standard, and a *horizontal* verification method that ensures agreement between the static control and the actual generated code. Initial experiments run on a benchmark suite and on the control software of a quadrotor copter show promising results regarding the feasibility of full verification of key aspects of task interaction in OSEK systems, and in particular show that even substantial examples generate moderate-sized control flow graphs that allow fully-fledged model checking. The key to keeping state spaces small was to exploit OS tailoring as well as the particularities of scheduling in the OSEK standard. While our experiments indicate that dynamic exploration of control flow graphs scales up well, static methods that reconstruct the control flow from the compiled binary [24] may serve as a complementary approach in the future.

In further work, we plan to build more comprehensive coverage of the OSEK standard and to develop methods for validating our formalization of the standard against the informal specification, possibly building on previous work in this direction [17]. Also, we will apply our approach of model checking compiler-generated control flow graphs to application-specific verification goals beyond standard conformance.

Source code and data are available at
<https://gitlab.cs.fau.de/dosek-verification>

References

- [1] P. Marwedel, *Embedded System Design*. Springer, 2006.
- [2] J. Cooling, *Software Engineering for Real-Time Systems*. Addison-Wesley, 2003.
- [3] M. Broy, “Challenges in automotive software engineering,” in *Proc. ICSE’06*. ACM Press, 2006, pp. 33–42.
- [4] OSEK/VDX Group, “Operating system specification 2.2.3,” Tech. Rep., Feb. 2005.
- [5] AUTOSAR, “Specification of operating system (version 5.1.0),” Automotive Open System Architecture GbR, Tech. Rep., Feb. 2013.
- [6] OSEK/VDX Group, “OSEK implementation language specification 2.5,” Tech. Rep., 2004.
- [7] C. Dietrich, M. Hoffmann, and D. Lohmann, “Global optimization of fixed-priority real-time systems by RTOS-aware control-flow analysis,” *ACM Trans. Embed. Comp. Sys.*, vol. 16, pp. 35:1–35:25, 2017.
- [8] F. Scheler and W. Schröder-Preikschat, “The RTSC: Leveraging the migration from event-triggered to time-triggered systems,” in *Proc. ISORC’10*. IEEE Computer Society Press, 2010, pp. 34–41.
- [9] R. Milner, *A Calculus of Communicating Systems*. Springer, 1980.
- [10] H.-P. Deifel, C. Dietrich, M. Göttlinger, D. Lohmann, S. Milius, and L. Schröder, “Automatic verification of application-tailored OSEK kernels,” full version; available at ???
- [11] M. Vardi, “Branching vs. linear time: Final showdown,” in *Proc. TACAS 2001*, ser. LNCS, vol. 2031. Springer, 2001, pp. 1–22.
- [12] R. Jhala and R. Majumdar, “Software model checking,” *ACM Comput. Surv.*, vol. 41, pp. 21:1–21:54, 2009.
- [13] L. Cordeiro and B. Fischer, “Verifying multi-threaded software using SMT-based context-bounded model checking,” in *Proc. ICSE’11*. ACM Press, 2011, pp. 331–340.
- [14] M. Schordan and A. Prantl, “Combining static analysis and state transition graphs for verification of event-condition-action systems in the RERS 2012 and 2013 challenges,” *STTT*, vol. 16, pp. 493–505, 2014.
- [15] L. Waszniewski and Z. Hanzálek, “Formal verification of multitasking applications based on timed automata model,” *Real-Time Systems*, vol. 38, no. 1, pp. 39–65, Jan. 2008.
- [16] Y. Huang, Y. Zhao, L. Zhu, Q. Li, H. Zhu, and J. Shi, “Modeling and verifying the code-level OSEK/VDX operating system with CSP,” in *Proc. TASE’11*. IEEE Computer Society Press, 2011, pp. 142–149.
- [17] D. Vu and T. Aoki, “Faithfully formalizing OSEK/VDX operating system specification,” in *Proc. SoICT’12*. ACM, 2012, pp. 13–20.
- [18] D. Vu, Y. Chiba, K. Yatake, and T. Aoki, “Verifying OSEK/VDX OS design using its formal specification,” in *Proc. TASE’16*. IEEE Computer Society, 2016, pp. 81–88.
- [19] H. Zhang, T. Aoki, and Y. Chiba, “Verifying OSEK/VDX applications: A sequentialization-based model checking approach,” *IEICE Transactions*, vol. 98-D, no. 10, pp. 1765–1776, 2015.
- [20] H. Zhang, T. Aoki, H. Lin, M. Zhang, Y. Chiba, and K. Yatake, “SMT-based bounded model checking for OSEK/VDX applications,” in *Proc. APSEC’13*. IEEE Computer Society, 2013, pp. 307–314.
- [21] M. Zhang, Y. Choi, and K. Ogata, “A formal semantics of the OSEK/VDX standard in K framework and its applications,” in *Proc. WRLA’14*. Springer, 2014, pp. 280–296.
- [22] K. Tigori, J.-L. Béchenec, S. Faucou, and O. Roux, “Formal model-based synthesis of application-specific static RTOS,” *ACM Trans. Embed. Comp. Sys.*, vol. 16, pp. 97:1–97:25, 2017.
- [23] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood, “seL4: formal verification of an OS kernel,” in *Proc. SOSP’09*. ACM, 2009, pp. 207–220.
- [24] T. Sewell, M. Myreen, and G. Klein, “Translation validation for a verified OS kernel,” in *Proc. PLDI’13*. ACM, 2013, pp. 471–482.

Appendix

This appendix is not part of the official publication

We present additional details concerning our formal verification experiments. Table 5 lists names, priorities and preemptibility of the tasks and ISRs for the examples from Table 4.

In the following we discuss more in detail our results on negative tests obtained by fault injection.

Details for Section 4.2. The verification results for the systems with changes introduced by the modified generator are shown in Table 6. In each case it is indicated whether the verification has discovered an error (+) or not (-). In the cases of added or merged states the node ids of the adjacent nodes are given. For added edges the corresponding system call is given by abbreviation as follows:

Abbrev.	System call
CT	chain task
T	terminate task
e	ε -transition
EI	enable all interrupts
DI	disable all interrupts
GR	get resource
RR	release resource
AT	activate task
WE	wait event
SE	set event
CE	clear event

Note that chain task, get resource, release resource and activate task all have a parameter and the last three systems calls concerning events take a set of events as parameters.

All the injected errors that have not been found (see Table 6) can be explained. In the example `bcc-resource1j` there was an injected error that merged nodes 20 and 22 that went undetected because the relevant edge would have been eliminated during ε -elimination. The undetected toggle of the preemptible flag of task H5 was not detected because we have no formula specifying that only non-preemptible tasks can have the internal scheduler resource; the SSTG marks it as occupied and therefore the scheduling behaviour remains correct. In the example `bcc1-sse1c` an added `DisableAllInterrupts` edge goes undetected because we have not formalized that this `syscall` does disable the interrupts. We did however formalize that interrupts cannot become disabled without witnessing this `syscall`. Similarly,

TABLE 5: Examples used for testing. A “P” after the priority denotes that the task is preemptible, and parentheses indicate that a task never runs in that specific scenario.

Name	Control flow priorities
<code>bcc1-resource1j</code>	(H2: 5 P), H1: 6 P, H3: 4 P, H4: 3 P, H5: 2 P
<code>bcc1-sse1c</code>	(H1: 6 P), H2: 5 P, H3: 4 P, H4: 3 P, H5: 2 P
<code>ecc1-bt1g</code>	(BT2: 12 P), (ET2: 11 P), BT1: 4 P, Coord: 2, ET1: 3 P
<code>ecc1-event1e</code>	H1: 5 P, H2: 4 P, H3: 6 P
<code>bcc1-isr2d</code>	H1: 6 P, H2: 5 P, H3: 4 P, ISR1: 10
<code>timing-abcomp</code>	Abort: 10, Computation: 2 P, Control: 6 P
<code>copter-small</code>	AlarmFlightControl: 100, AlarmSignalGatherInitiate: 100, CopterControl: 6 P, FlightControlActuate: 14 P, FlightControlAttitude: 13 P, FlightControl: 12, MavlinkRecvHandler: 100, MavlinkSend: 11 P, SignalGatherFinished: 26 P, SignalGatherInitiate: 25 P, SignalGatherTimeout: 24 P, SignalProcessingActuate: 23 P, SignalProcessingAttitude: 22 P
<code>copter</code>	AlarmCopterControlWatchdog: 100, AlarmFlightControl: 100, AlarmSignalGatherInitiate: 100, CopterControl: 6 P, CopterControlWatchdog: 2, FlightControlActuate: 14 P, FlightControlAttitude: 13 P, FlightControl: 12, MavlinkRecvHandler: 100, MavlinkSend: 11 P, SignalGatherFinished: 26 P, SignalGatherInitiate: 25 P, SignalGatherTimeout: 24 P, SignalProcessingActuate: 23 P, SignalProcessingAttitude: 22 P

the task H5 receives the internal resource despite not being preemptible in reality. The preemptible toggle of ET2 is not found in the example `ecc1-bt1g` because the task simply never runs during the lifetime of the application. The undetected added ε -edge in this example is parallel to an existing ε -edge and therefore does not cause a violation of any property. All other undetected errors in this example are due to priority changes with tasks that would not have run anyway and do not affect the ordering of priorities of the tasks that actually run. The merging of nodes 25 and 19 in example `bcc1-isr2d` only eliminates an ε -edge and is therefore undetected. The other undetected errors in that example are caused due to the fact that only H3 can ever be ready simultaneously to any other task. In the larger examples `timing-abcomp`, `copter-small` and `copter` all injected errors are found due to the fact that it is simply much less likely that our 30 injected errors accidentally produce one of the very few changes (relative to the total number of possible changes) that are legal to make.

TABLE 6: Negative test results for vertical verification. For every example, we introduced up to 15 faults into the SSTG and into the OIL configuration read in by the generator. The detection of the fault by our verification is indicated by a (+) if it lead to a violation of the OSEK standard detected by the tool chain. For the **undetected faults (-)**, we manually confirmed that the fault had no influence on that specific test case. In total, 177 faults (out of 188 injections) lead to an error and where detected by our toolchain.

Example	SSTG Fault					Fault in OIL configuration file				
	Transition added	States merged	Preemptible toggled	Autostart toggled	Priorities switched					
bee1-resource1j	+T60,50 +DI61,33 +AT(H2)45,68 +RR(SCHEDULER)27,36 +e2,54 +EI15,62	+50,39 +63,27 +35,58 +41,48 +4,16 +7,16 +10,43 +43,50 -20,22	-H5	+H4 +H5 +H2	+H4,H3 +H4,H2 +H5,H3 +H1,H5 +H1,H3					
bee1-sselc	+EI68,60 -D161,33 +AT(H3)45,66 +RR(R2)34)27,36 +e2,54 +e72,15	+50,39 +63,27 +35,58 +68,41 +4,16 +7,67 +10,43 +43,50 +20,22	-H5	+H4 +H5 +H2	+H3,H4 +H4,H2 +H5,H3 +H1,H5 +H1,H3					
ecc1-bf1g	-e17,18 +RR(SCHEDULER)21,15 +WE(ET1-E1)11,18 +RR(SCHEDULER)6,9 +EI18,0 +CT(ET2)18,17	+12,9 +15,6 +8,14 +17,10 +1,19 +1,16 +2,10 +10,24 +21,5	-ET2	+ET1 +ET2	-ET1,Coord +ET1,BT2 -ET2,Coord -BT1,ET2 +BT1,Coord					
ecc1-event1e	+SE(H3-E1)17,26 +e23,21 +WE(H2-E1)11,18 +RR(SCHEDULER)24,6 +EI18,0 +WE(H3-E1)18,31	+28,12 +15,6 +8,14 +17,10 +1,19 +1,16 +2,25 +10,25 +21,5	+H3	+H2 +H3 +H1	+H2,H1 +H3,H2 +H1,H3					
bee1-istr2d	+e34,37 +DI30,16 +AT(H1)22,37 +RR(SCHEDULER)13,18 +EI37,1 +CT(H2)36,35	-25,19 +31,13 +17,29 +34,20 +2,39 +3,33 +5,21 +21,25 +10,11	+H3	+H2 +H3 +H1	-H2,H1 +H3,H2 +H1,H3					
timing-abcomp	+e137,151 +WE(Control-Computation)Finished)122,66 +WE(Control-Computation)Finished)91,149 +RR(SCHEDULER)54,72 +EI150,4 +CT(Computation)145,140	+101,79 +126,54 +71,116 +137,82 +9,32 +15,134 +20,86 +86,101 +40,44	+Control +Computation	+Control +Computation	+Control,Computation					
copter-small	+e2195,2427 +RR(SCHEDULER)2731,1955 +AT(SignalProcessingActuate)1458,2391 +RR(SPIBus)868,1156 +EI2403,78 +CT(FlightControl)2336,2250	+1622,1269 +2018,870 +1142,1869 +2195,1318 +152,2549 +246,2161 +32,1393 +1378,1624 +2714,648	+SignalProcessingAttitude	+SignalGatherInitiate +SignalGatherTimeout +SignalProcessingActuate +FlightControl	+SignalGatherInitiate,MavlinkSend +SignalGatherTimeout,FlightControl +MavlinkSend,SignalGatherTimeout +SignalProcessingActuate,SignalGatherFinished +CopterControl,SignalProcessingActuate +FlightControlActuate,SignalGatherFinished					
copter	+EI18775,7702 +DJ7817,4286 +AT(FlightControl)Attitude)5826,8488 +RR(SPIBus)3465,4619 +e313,7024 +EI1987,7943	+6485,5071 +8067,3475 +4560,7472 +8774,5268 +611,2085 +985,8641 +1299,5567 +5509,6491 +2593,2865	+SignalProcessingActuate	+SignalGatherFinished +SignalGatherInitiate +SignalProcessingActuate +SignalGatherTimeout +SignalProcessingAttitude +FlightControlAttitude	+SignalGatherFinished,FlightControl +SignalGatherInitiate,FlightControl +SignalGatherTimeout,MavlinkSend +CopterControl,SignalGatherTimeout +CopterControlWatchdog,MavlinkSend					

TABLE 7: Negative test results for horizontal verification. For every example, we introduced up to 15 faults into the generated OS kernel. A (+) marks faults detected by our verification, i.e the DSTG was non-isomorphic to the SSTG. For the **undetected faults (-)**, we manually confirmed that the fault had no influence on that specific test case. In total, 61 faults (out of 81 injections) lead to an error and were detected by our toolchain.

Example	Priorities switched	Preemptable toggleed	Autostart toggleed
bcc1-resource1j	+H2,H5 +H1,H3 +H2,H3	+H5,-H3,-H2,+H4	+H2,+H5,+H4,+H3,+H1
bcc1-sse1c	+H2,H5 +H1,H3 +H2,H3	+H5,-H3,-H2,+H4	+H2,+H5,+H4,+H3,+H1
ecc1-bt1g	+BT1,Coord -BT2,ET2 -BT2,Coord	+Coord,-ET2,-BT2,+ET1	+BT2,+ET2,+ET1,+Coord,+BT1
ecc1-event1e	+H1,H3 +H1,H2	+H1,+H2,-H3	+H1,+H3,-H2
bcc1-isr2d	+H1,H3 -H1,H2	-H3,-H2,-H1	+H1,+H3,+H2
timing-abcomp	+Control, Computation	+Computation, -Control	+Control, -Computation
copter-small	+FlightControlAttitude,SignalProcessingAttitude +CopterControl,FlightControlActuate +FlightControl,MavlinkSend	+FlightControl, +SignalGatherTimeout, -SignalProcessingAttitude, -SignalGatherFinished	+FlightControl, +SignalProcessingAttitude, +SignalGatherTimeout, +MavlinkSend, +FlightControlActuate
copter	+FlightControlActuate,SignalProcessingActuate +CopterControl,CopterControlWatchdog -FlightControlAttitude,FlightControl	+MavlinkSend, -SignalProcessingActuate, +FlightControlAttitude, +SignalGatherInitiate	+FlightControlAttitude, +SignalProcessingActuate, +SignalGatherInitiate, +FlightControl, +CopterControlWatchdog