Easy Modelling and Verification of Unpredictable and Preemptive Interrupt-driven Systems

Minxue Pan, Shouyu Chen, Yu Pei, Tian Zhang and Xuandong Li
State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China
mxp@nju.edu.cn, 1124688454@qq.com, max.yu.pei@polyu.edu.hk, {ztluck,lxd}@nju.edu.cn

***Abstract*—The widespread real-time and embedded systems are mostly interrupt-driven because their heavy interaction with environment is often initiated by interrupts. With the interrupt arrival being unpredictable and the interrupt handling being preemptive, a huge number of possible system behaviours are generated, which makes the correctness assurance of such systems a difficult and costly task. Model checking is considered to be one of the effective methods for exhausting behavioural state space for correctness, however, existing modelling approaches for interrupt-driven systems are either based on calculus or automata theory, with which designers in industry are not familiar. To address this problem, we propose a new modelling language called *interrupt sequence diagram* (ISD). By extending the popular UML sequence diagram notations, the ISD supports the modelling of interrupts' essential features visually and concisely, with a formal semantics interpreting the unpredictable and prioritised preemptive system behaviour caused by interrupts. For model checking purpose, we have devised algorithms that automatically translate an ISD to a subset of hybrid automata so as to leverage the abundant off-the-shelf checkers. Experiments on examples from both real-world and existing literature were conducted and the results demonstrate our approach's usability and effectiveness.**

I. Introduction

Interrupt-driven systems, where processing is initiated by interrupt requests, are gaining their popularity since interrupts are a key design primitive for software systems that actively make interactions among system components and closely interact with the environment. They are commonplace in all styles of computing platforms, including safety-critical embedded platforms, low-power mobile platforms and high-end information systems [1]. Particularly, most cyber-physical systems are interrupt-driven, since interrupts are an extremely common form of concurrency that the control software uses to obtain sensor data from its physical environment [2], and enable timely response to outside stimuli in a power-efficient way [1]. However, interrupts can cause problems, for many of them can happen at arbitrary time and preempt the running tasks, which adds non-determinism and concurrency to the systems. This poses challenges to develop reliable interrupt-driven systems, as designers have to explicitly handle unpredictable system behaviour caused by interrupts. As a consequence, interrupt-driven systems are error-prone [3], [4], [5], and need extensive efforts for quality assurance.

Testing is one of the primary ways to assure the quality of systems. However, existing testing techniques for sequential programs [6], [7] or even concurrent programs [8], [9] have not addressed the problems caused by interrupts adequately. They often cannot identify or capture the concurrency brought about by interrupts precisely. To address this limitation, researchers tried to find effective interrupt scheduling algorithms that fire interrupts at proper points of time [5], or suitable test adequacy criteria to guide and evaluate the testing process [4]. Nonetheless, testing of interrupt-driven systems can still be insufficient. The generation of interrupt requests is usually random and non-deterministic, and the interrupt handling is often preemptive and nested, which results in that the number of possible system behaviours grows exponentially in the number of occurred interrupts, while defects related to specific behaviour are difficult to detect by testing approach [5].

Model checking, on the other hand, can rigorously verify a system's behaviour by exhaustively exploring the state space of a software system. Model checking can be applied to programs with restricted grammar, or to assorted formal models. Recent studies have noticed the significance and uniqueness of interrupt-driven systems, and new modelling languages are proposed. For example, interrupt time automata (ITA) are proposed [10], [11] to model multi-task systems with interrupts. They form a subclass of stopwatch automata [12], where the real valued variables (with rate 0 or 1) are organised along priority levels. ITA are powerful in expressiveness, however, we argue that industrial designers may find them difficult to use. When modelling with ITA, designers have to consider all possible inter-leavings of states, as well as the clocks that specify the timing constraints. On the other hand, UML sequence diagrams [13] offer an intuitive and visual way of describing interactions among system components and the environment, and are widely used in industry. In a survey conducted in [14], sequence diagrams were recognised as one of the most frequently used diagrams, and system analysts and programmers admitted that they rely most on sequence diagrams along with class diagrams to capture requirements and exchange information. However, sequence diagrams are still inadequate to model interrupt-driven systems. Interrupts' arrival can be unpredictable, and their handling is preemptive and prioritised. Time could be a complex concept too, since execution time of tasks and interrupt service routines (ISRs)

This work was supported by the National Key R&D Program of China (No. 2017YFA0700604), National Natural Science Foundation of China (Nos. 61502228, 61632015, 61561146394, 61572249) and the Fundamental Research Funds for the Central Universities (020214380045).

can be interrupted and resumed. These interrupt-specific features are not supported by sequence diagrams. To address these limitations, we propose the *Interrupt Sequence Diagram (ISD)* which extends the sequence diagram with interrupt modelling mechanisms. It can model interrupt-driven systems easily and intuitively, while still reserve the original sequence diagram's usability.

The extension is conducted by following the UML standard notation for easy comprehension. We propose to introduce a new *CombinedFragment "int"* to specifically model interrupts. The behaviour of the CombinedFragment int can model an interrupt's unpredictable arrival and its handler's prioritised preemptive execution. Furthermore, we propose a new kind of timing constraints called *task constraints* to model the *actual* execution time of an ISR or an ordinary task. As it is mentioned, tasks and ISRs can often be preempted during execution, which means the actual execution time does not equal to the time duration from the start to the completion of the tasks/ISRs. The sequence diagram only supports the latter form of timing constraints, however, in interrupt-driven systems, having this simple form of timing constraints is not enough. The time duration of a task/ISR can be uncertain and infeasible to specify, since the time amount preempted by interrupts varies in different system behaviours.

To facilitate the formal verification, we also provide an automata based semantics for ISD. We present algorithms to translate the ISD to integration automata (IA) [15], [16] which are a subset of hybrid automata. Besides giving a rigorous semantics, another benefit of translating ISD to IA is that we can leverage the abundant off-the-shelf hybrid automata checkers [17], [18], [19], [20]. We conducted multiple case studies with the state-of-the-art tool SpaceEx [18], and the experiment results consistently showed that our approach is effective in defect detection and achieves good performance. Particularly, the best trait of ISD is that it is easier to learn and use compared with existing modelling languages such as ITA. Therefore, we conducted an experiment to compare the usability of ISD and ITA, and the results confirmed our conjecture. The main contributions of this paper are:

- We propose a novel modelling language ISD to specify the interrupt-driven systems. To our best knowledge, ISD is the first sequence diagram based language which supports the specification of interrupt's unpredictable arrival, prioritised preemptive handling and the consequent time suspension and resumption.
- We propose a translation algorithm that automatically translates ISD to IA, to provide rigorous semantics and to leverage existing checkers for correctness checking.
- We developed a tool named ISDChecker, which is, to our knowledge, the first available tool that checks graphical models for interrupt-driven systems. Evaluation on previous studied and real-world cases shows ISDChecker's effectiveness and usability.

The rest of the paper is organised as follows. In the next section, we introduce the UML sequence diagram and a modelling example. Section 3 proposes the ISD, and Section 4 presents the algorithm that translates ISD to IA. Section 5 conducts a series of experiments to evaluate our approach's usability, effectiveness and efficiency. Section 6 discusses the related work and Section 7 draws the conclusion.

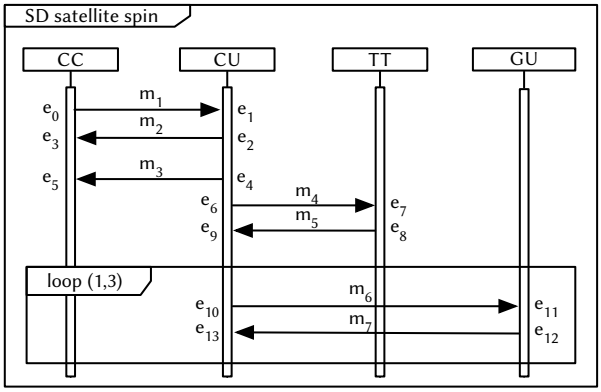

Fig. 1. A sequence diagram example

II. Sequence Diagram and Motivating Example

UML sequence diagrams form a class of important UML interaction models. Each of them describes a set of interactive scenarios, as Fig. 1 shows an example. There are two dimensions in a sequence diagram. The vertical dimension represents time, and the horizontal dimension consists of different lifelines representing participating entities. Information exchange between lifelines is carried out by messages represented by arrows. In the simplest form, a sequence diagram depicts the desired exchange of messages, and corresponds to a single execution of the system. To specify complex scenarios conveniently, sequence diagram enables operations such as choice and iteration through *CombinedFragments* (or *fragments* for short). A fragment is defined by an interaction operator and one or more interaction operands. The notation for a fragment is a solid-outline rectangle, and the operator is shown in a pentagon in the upper left corner of the rectangle. Each operand is composed of a subset of messages in the diagram. In this paper, we consider three most frequently used interaction operators, which are *loop*, *alt* and *opt*. The operator loop designates that the fragment represents a loop, so its operand will be repeated a number of times. A guard that may include a lower and an upper number of iterations of the loop can be associated to the fragment. Fig. 1 shows an example of the loop fragment, where the iteration times is restricted from 1 to 3 times. The operator alt designates that the fragment represents a choice of behaviour. At most one of the operands will be chosen, and the chosen operand must have a guard expression that evaluates to true at this point in the interaction. The operator opt has a similar meaning as the operator alt, except that it has only one operand. The fragment opt specifies the behaviour where either the operand happens or nothing happens. In this paper, we enforce strict sequencing on the fragments. Consequently, a fragment will cover all lifelines, so that when the execution control of flow enters a fragment, all lifelines enter the fragment.

Sequence diagrams are popular in industry, because they help designers focus on the most frequent or critical scenarios. However, they are still insufficient to model interrupt-driven systems. To support this claim, we present a real-world case of designing the spin action in a satellite controlling system. The original sequence diagram provided by the designers was an informal sketch, which we revised to conform to the UML standard and to exclude sensitive information, as shown in Fig. 1. The spin action involves four participating entities performing three tasks. In the first task, the Command Centre (CC) sends an inquiry to the satellite's Computing Unit (CU) about the status of the satellite (Message m_1 and m_2). Then, at some point in time, the second task starts that CU informs CC that it is going to compute the instructions for the satellite to spin an angle to better absorb sunlight and in the meantime will not response to CC commands (Message m_3). Then CU sends the instructions to the thruster (TT) to be executed (Message m_4) and TT acknowledges CU when the spin action is completed (Message m_5). In the third task, the satellite communicates with the Ground Unit (GU) periodically to exchange information (Message m_6 and m_7). The communication cannot be interrupted by other operations, or the communication link would be lost. If the interpretation of the diagram in Figure 1 strictly followed the UML specification, then the system should be running without problems, because in a UML sequence diagram, events corresponding to message sending and receiving are subject to predefined partial orders deduced from the visual order of the diagram (see [13] for more details about the partial orders among events). Thus, in Figure 1, e_4 would be considered to happen after the occurrence of e_2, and e_{10} after e_9, which means the three operations happen in a sequential order. However, this sequence does not cover all the actual system behaviours. In reality, the task of CC querying CU and the task of CU communicating with GU are both interrupts, so their arrivals are unpredictable. Even worse, the former has a higher priority than the latter, which means that the inquiry from CC to CU is not obliged to happen before CU and GU's communication, but can interrupt the communication process and cause a problematic behaviour. The engineers had not noticed the problem and implemented the system as designed, which resulted in a costly failure during the system integration test. To prevent this kind of problems to happen again, it requires that the problematic behaviour be captured by the designs so that they can be found either by humans or automatic checkers. Unfortunately, because of the partial orders enforced on the events, it is impossible for a UML sequence diagram to specify unpredictable interrupts. In the following section, we propose a smooth extension to the UML sequence diagrams to allow the specification of unpredictable and preemptive interrupt-driven system behaviours.

III. Interrupt Sequence Diagram

The Interrupt sequence diagram (ISD) is an extension of UML sequence diagram to model interrupt-driven systems. We notice that the most unique feature of interrupts is that in most occasions, their arrivals are spontaneous and unpredictable, which needs a particular modelling mechanism. We propose a new fragment called *int*, which consists of one operator and one operand. The operator int declares that the fragment represents an interrupt behaviour, whose syntax is:
'int' ⟨priority_exp⟩ [⟨occur_bound⟩] ['['⟨mask_cond⟩']'].*
priority_exp models the interrupt's priority, *occur_bound* specifies the times the interrupt can occur and the separation in time between two consecutive interrupt arrivals, and *mask_cond* models the interrupt mask condition. In the following, we give detailed explanations about different parts of the int fragment.

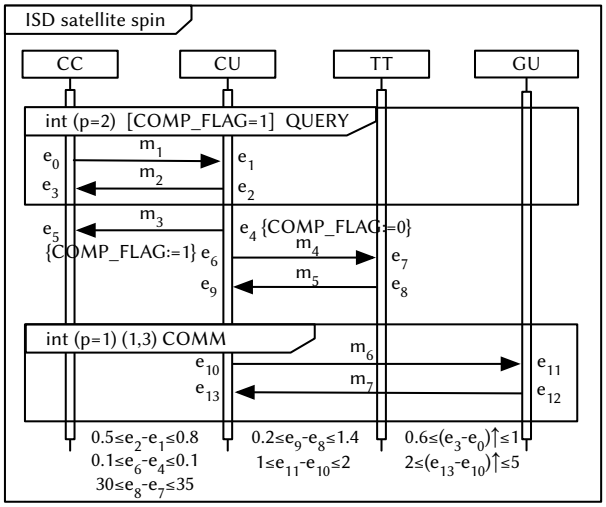

Fig. 2. An interrupt sequence diagram example

A. Order of Events

The events in the operand of the int fragment specify the execution process of the interrupt arrival and handling. As an event will not belong to two interrupts, it is required no overlapping between two int fragments. To depict interrupt's unpredictability, it is required that there should be no partial orders deduced from the visual orderings of messages between the events from the inside and outside an int fragment, respectively. For example, in Fig. 2, e_2 and e_4 are two events inside and outside an int fragment, respectively, and thus there is no ordering require that e_2 occur before e_4.

To specify the prioritised preemption of interrupts, a *priority expression* is proposed, defined as *⟨priority_exp⟩::= '(p='⟨priority⟩')'*, where *⟨priority⟩::=non-negative natural*. The *priority* of the int fragment comes from the priority of the interrupt and applies to all the events in the fragment. To be consistent, we assign a default priority to the entire diagram, which applies to events not in any int fragment. In fact, excluding the int fragments, the entire diagram with its enclosed loop, alt and opt fragments can also be viewed as a fragment, except its priority is the lowest. From now on, we use the term *interaction fragment* indiscriminately to represent the int fragment or the diagram excluding all int fragments. The priority of the interaction fragments introduces a new kind of orderings among events: it is required that when the execution control of flow is in an interaction fragment of

priority p_1, it can move to another interaction fragment of priority p_2 if $p_2 > p_1$, but not vice versa. Formal semantics of the int fragment is presented in Sec. IV-B4.

Let's revisit the example of the satellite spin and re-model it example with ISD as shown in Fig. 2. There are two int fragments: QUERY and COMM. QUERY has a priority of 1, and COMM has a priority of 2. So, events in int COMM can occur after the occurrence of e_{10} and before the occurrence of e_{12}, causing the communication between CU and GU to fail. With the help of int fragments, a possible defect of the design can be revealed.

B. Interrupt Mask

To prevent defects caused by unpredictable interrupts, designers often use interrupt masks to prevent some less important interrupts to interfere with a crucial task, while still be able to respond to more important interrupts. The ISD supports modelling the interrupt masks by using *mask variable* updates and tests. Each interrupt mask is modelled by a mask variable, with a value from $\{0, 1\}$. The enablement or disablement of the interrupt mask is achieved by setting the variable's value to 1 or 0, respectively. In UML sequence diagrams, the values of variables are updated via actions associated with messages, whereas in ISD we need a more accurate mechanism to model mask variables, since the enablement/disablement of interrupts always happen at the beginning or the end of uninterruptible operations. Thus, the ISD supports variable updating actions to be associated with events, since it is the events that represent the beginning and the ending of message processing.

An int fragment can have *mask conditions*, which are defined as ⟨*mask_cond*⟩*::=* ⟨*mask_variable*⟩*'=1'*, where ⟨*mask_variable*⟩*::=letter tokens*. The events in the int fragment can happen when all mask conditions are evaluated as true, or no mask condition is provided. For example, the int fragment COMM in Fig. 2 has one mask condition "COMP_FLAG=1". Therefore, the events in this fragment can happen when variable COMP_FLAG is set to 1.

C. Modelling of Time

The other mechanism to reduce the unpredictability imposed by interrupts is timing specification. Timing specification can specify the timing properties of the system. Designers can designate the time point or interval in which an interrupt could happen, or the duration of the interrupt handler, to restrict interrupts' behaviour and reduce their uncertainty. For example, suppose that a system consists of one uninterruptible task and one interrupt. If there were some timing constraints specifying that the interrupt would not occur during the execution of the task, then no interrupt mask would need to be disabled, which could save the computing resource and time. In UML sequence diagrams, the timing mechanisms are all about the time duration between two events, which specifies how much time has passed from the occurrence of one event till the next. By using event names to represent the occurrence time of events, the timing constraint can be defined as $a \leq e - e' \leq b$ (a, b are real numbers, b may be ∞), which requires that the time duration from the occurrence of e' till the occurrence of e be within the range $[a, b]$. Although this kind of timing constraints is sufficient for most systems, they cannot meet the need to specify interrupt-driven systems. The designers of interrupt-driven systems often need to specify the actual execution time of a task, however, it is not the time duration between the task's start and completion events, since its execution can be interrupted by unpredictable interrupt occurrences. For example, in the satellite spin example in Fig. 2, although one may know that executing the interrupt COMM (from the occurrence of e_{10} to that of e_{13}) requires 2 to 5 time units, it is incorrect to model this time information using $2 \leq e_{13} - e_{10} \leq 5$, because the execution of COMM can be suspended by QUERY. We provide a new mechanism called *task constraint*. A task constraint, denoted as $a \leq (e - e') \uparrow \leq b$, is about two events e and e' satisfying that both events are from the same interaction fragment f. The value of $(e - e') \uparrow$ is computed as the time duration from the occurrence of e' till the occurrence of e, subtracts the time when the diagram's execution control of flow is not in f. For example, the task constraint $2 \leq (e_{13} - e_{10}) \uparrow \leq 5$ in Fig. 2 specifies the time that the execution control of flow stays in the fragment COMM is between 2 to 5 time units. Formal semantics of timing and task constraints are given in Section IV-B5.

Furthermore, it is possible to specify the minimum and maximum times an interrupt can occur, and for interrupt that can occur multiple times, the minimum separation in time between two consecutive interrupt arrivals. We use the expression *occur_bound* to specify these timing requirements. The expression is defined as ⟨*occur_bound*⟩*::='('*⟨*min*⟩*','*⟨*max*⟩ *[','*⟨*separation*⟩*]')'*, where ⟨*min*⟩*::=positive natural,* ⟨*max*⟩ *::= positive natural (greater than or equal to* ⟨*min*⟩*)* | *'*∞*'*, ⟨*separation*⟩*::=positive real*. Without an explicit *occur_bound*, the interrupt is supposed to occur exactly once by default, and if there is no *separation* field provided, arbitrary separation time (positive real) is assumed. For example, for the int fragment COMM, the *occur_bound* $(1, 3)$ specifies that the fragment can occur 1 to 3 times, and the minimum separation time can be arbitrary.

D. Syntax Definition of ISD

Now we can formally define the syntax of ISD as follows.

Definition 1: An interrupt sequence diagram (ISD) is a tuple $D = (L, E, M, R, V, U, F, C)$, where

- L is a finite set of lifelines;
- E is a finite set of prioritised events whose elements are pairs (e, p), where e is the event and p is its priority.
- M is a finite set of messages. For each $m \in M$, $m = (e, e')$, where $e, e' \in E$ correspond to the sending and the receiving of m, respectively;
- $R : E \to L$ is a labelling function which maps each event $e \in E$ to a its sending (receiving) lifeline $R(e) \in L$;
- V is a finite set of message orderings whose elements are a pair (m, m') $(m, m' \in M)$ such that m visually precedes m';

- U is a finite set of mask variable updates, whose element is a tuple (var, val, e) where var is a variable, $val \in \{0, 1\}$ is the updated value and $e \in E$ is the event associated with the variable update;
- $F = F_{loop} \cup F_{alt} \cup F_{opt} \cup F_{int}$ is a finite set of fragments. Each element $f \in F$ is a tuple (o, g, M_f) where $o \in \{loop, alt, opt, int\}$ is the operator, g is the guard expression of f, and $M_f \in 2^M$ is a subset of M;
- C is a finite set of timing constraints and task constraints.

IV. Translation of ISD to IA

We define an automata based semantics for ISD to facilitate the formal verification. The semantics of ISD is interpreted by a translation to integration automata (IA), which are a special case of hybrid automata. In this section, we will first visit the concept of IA, followed by the presentation of the translation process.

A. Integration Automata

Hybrid automata are finite automata extended with a finite set of real typed variables whose values change continuously at each location. The change rate of the variables are designated by the *flow conditions* associated with locations. There are also *invariants* in each location indicate that the conditions need to be satisfied when the location is active. Transitions between locations are guarded by *jump conditions* on the variables and their executions may reset some of the variables by the *reset actions*. If the invariants, jump conditions and reset actions are all linear expressions over the variables, and the flow conditions specifying the allowed values of the first derivatives of the variables are in the form of a rate polytope, then the hybrid automata are called linear hybrid automata. The integration automata strengthen the restrictions by stipulating the values of the first derivatives of the variables to be either 0 or 1, which is suitable to model time suspension and resumption. Formal definitions of the integration automata can be found in [16].

B. Translation Algorithm

We translate the ISDs into integration automata following the generally agreed semantics of basic interactions and the alt, opt and loop fragments. Furthermore, we handle the int fragment, timing constraints and mask variables which are used in modelling interrupts.

A common approach to translating sequence-based diagrams into automata is to generate one automaton for every object, and use their parallel composition as the complete model [21], [22]. The acquisition of the complete model requires the synchronisation between different object automata, which is achieved by using synchronisation labels corresponding to the message names. Since transitions with the same synchronised labels must happen simultaneously, it is impossible to distinguish the message sending and receiving events over time. In interrupt-driven systems, time duration between events plays a vital role in correct system functions. Thus, in contrast to this approach, we propose to translate ISD using fragment as a basic unit, similar to [23].

Algorithm 1: Translation algorithm for a basic fragment

```
1  generate the initial location q0;
2  for each non-final location q that has no outgoing edges do
3      Acquire the set L of events in the path from q0 to q;
4      for any event e ∈ (E − L) do
5          if any event e' satisfying (e' ≺ e) ∈ O is in L then
6              generate a location q' and a transition (q, e, q');
7          for any location q''(q'' ≠ q') do
8              L' ← set of events in the path from q0 to q';
9              L'' ← set of events in the path from q0 to q'';
10             if L' = L'' then
11                 q'' ← merge q' and q'';
12                 change (q, e, q') to (q, e, q'');
13         end
14     end
15     if q has no outgoing edges then
16         mark q as final;
17 end
```

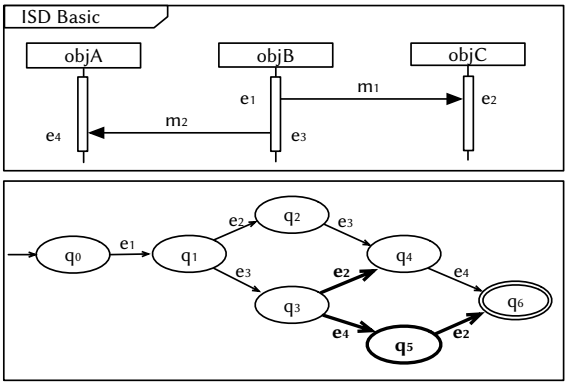

Fig. 3. From a basic fragment to an automaton

1) Translation of Basic Fragment: The simplest form of ISD is a basic fragment that does not have any nested fragments. The translation of a basic fragment is to generate an automaton satisfying that its label sequence set is equivalent to the event sequence set of the fragment. For the basic fragment, its events have a partial order relation which can be deduced from the message visual orderings. Let E be its event set and O be the set of event orderings where $e \prec e'$ such that e must occur before e'. Algorithm 1 takes E and O as its input can outputs an automaton. It constructs automaton paths where the events on transitions satisfy the events orderings. Additionally, it merges locations whose preceding paths contain the same set of events, so as to reduce the automaton's complexity.

Algorithm 1 focuses only on the event sequences in the basic fragment. The timing and task constraints are handled later in this section. We use an example to explain the translation algorithm. For the basic fragment in Figure 3, we have the set of events $E = \{e_1, e_2, e_3, e_4\}$ and the set of event orderings $O = \{e_1 \prec e_2, e_1 \prec e_3, e_3 \prec e_4\}$. Suppose we follow Algorithm 1 and has generated a part of the automaton, which is drawn in thin lines. Now for location q_3, we acquire the set $L = \{e_1, e_3\}$ which comprises events in the path from q_0 to q_3 *(line 3)*. Following *line 4* we pick event e_2 because it is in $E - L$. Since e_1 is already in L, the condition in *line 5* is satisfied, and a location q' and a transition (q_3, e_2, q') is generated at *line 6*. Location q' is not shown in Figure 3 because when executing *line 7-10*, we find that the set of

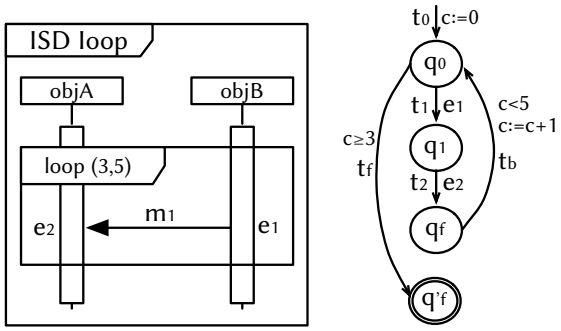

Fig. 4. Translation of a loop fragment

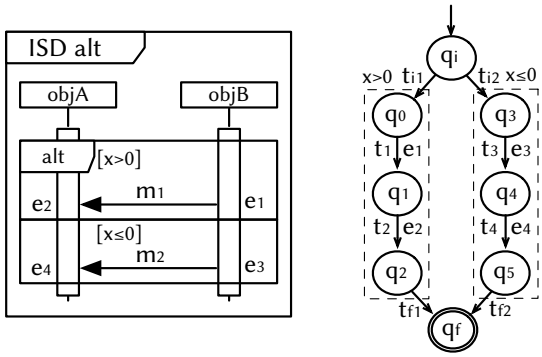

Fig. 5. Translation of an alt fragment

events in the path from q_0 to q' is the same as from q_0 to q_4. Since q_4 is an existing location, we merge q' into q_4 *(line 11)* and change (q_3, e_2, q') to (q_3, e_2, q_4) *(line 12)*. Then we loop back to *line 4*, pick event e_4, and generate new locations and transitions in the same manner. When q_6 is picked, since there are no events left for q_6 to grow outgoing edges, we mark it as final. Obviously for the generated automaton, there is only one initial location, and it is easy to prove that there is only one final location as well. Suppose there were more than one final location. Since the events along each path to the final locations are all the same, the final locations can be merged as one. Guaranteeing that there is only one initial location and one final location is useful when the fragment is nested in other fragments so that the translated automaton needs to be connected with other automaton.

2) Translation of loop Fragment: When it comes to the translation of fragments loop, alt and opt, Algorithm 1 for translating a basic interaction can be reused. We require a strict sequencing for loop fragment, which means the events in one iteration can happen only when all the events in previous iterations have happened. If the loop fragment contains just one basic fragment, we can use Algorithm 1 to generate an automaton A first. Let q_0 be the initial location, and q_f be the final location of A. The automaton A_L translating the loop fragment is constructed based on A as follows:

- add a new transition t_b from q_f to q_0 to form loops;
- generate a new location q'_f, and add a new transition t_f from q_0 to q'_f. Make q'_f as the final location and q_f as the non-final location;
- generate a variable c to count the iteration times. Initialise c to 0 on the transition to q_0, and increase c by 1 on t_b;
- for the guard (a, b) restricting the iteration times, assign a constraint $c < b$ to t_b, and a constraint $c \geq a$ to t_f.

Figure 4 shows an example. For the ISD in Figure 4, locations q_0, q_1 and q_f and transitions t_0, t_1 and t_2 constitute the automaton obtained by applying Algorithm 1 for translating a basic fragment. Transition t_b is added to form the loop, and its condition $c < 5$ ensures that the iteration times will not exceed 5 (this condition is checked before the increase of c). Location q'_f and transition t_f are added so that when exiting the loop, there is a condition $c \geq 3$ ensuring that the iteration times will be at least 3.

If there are other fragments nested, the translation can be done from inside to outside recursively. Specifically, for an enclosing fragment that has nested fragments, we first translate the nested ones. Without the nested fragments, the rest of the enclosing fragment is divided into separate parts which are later translated individually. Then the automaton corresponding to the entire fragment can be obtained by connecting one automaton's final location to the other's initial location, following the visual orderings of the nested fragments and the parts they separated in the enclosing fragment. The only exception is that the nested fragment is an int fragment, which will be discussed in IV-B4.

3) Translation of alt and opt Fragments: In the alt fragment, each operand has a guard condition. The events in an operand can happen only if the guard of this operand is true when the control of flow reaches the fragment, similar to the `if-else` structure in programming languages. To translate an alt fragment, we first generate two automata A_1 and A_2 for the two operands $o1$ and $o2$, respectively, and then compose them into a single one as follows:

- generate a location q_i and add transitions t_{i1} and t_{i2} from q_i to the initial locations of A_1 and A_2, respectively;
- generate a location q_f and add transitions t_{f_1} and t_{f_2} from the final locations of A_1 and A_2 to q_f, respectively;
- add the guard condition of $o1$ to transition t_{i1}, and that of $o2$ to transition t_{i2};
- make q_i the initial location and q_f the final location of the composed automaton.

An example is shown in Figure 5. The parts in dotted box are the two automata translated from the two operands of the alt fragment. The reason we use two new location q_i and q_f is to ensure that every automaton translated from one fragment has only one initial location and one final location, to simplify the translation of nested fragments.

The opt fragment can be viewed as a special case of the alt fragment with only one operand. Its translation is more straightforward: the operand is translated first, and the guard condition is added to the initial transition to the initial location.

4) Translation of int Fragment: Since the occurrence of interrupts is unpredictable, it is required that the events in the int fragment have no partial orders with events outside the fragment. In other words, the events in the int fragment and in the enclosing fragment form two independent event sets,

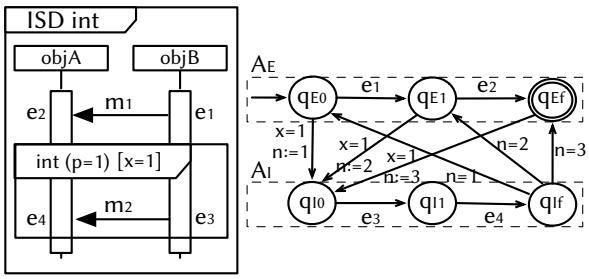

Fig. 6. Translation of an int fragment

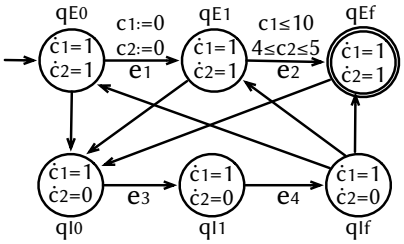

Fig. 7. Translation of task and timing constraints

and therefore, the int fragment and its enclosing fragment can be translated separately.

In interrupt-driven systems, when an interrupt request occurs, it is accepted when its priority is higher than the current task and the interrupt mask (if there is any) is enabled. When the preemptive interrupt completes, it returns the control of execution to the preempted task if there is no higher priority interrupt request. We refer this interrupt handling mechanism to interpret the relationship between the two automata translated from an int fragment and its enclosing fragment. Let A_I (q_{I0} and q_{If} are the initial and final locations, respectively) be the IA translated from an int fragment and A_E (q_{E0} and q_{Ef} are the initial and final locations, respectively) be the IA translated from its enclosing fragment. It is assumed that all events in A_E have lower priorities than those in A_I, which can be guaranteed by composing the automata translated from int fragments in the ascending order of priorities. Then the IA A having the equivalent behaviour as the ISD consisting both the int and its enclosing fragments can be obtained by composing A_I and A_E as follows:

- for any location q_i of A_E, add an transition from q_i to q_{I0} of A_I and associate it with an variable assignment $n := i$ and the interrupt mask condition of the int fragment if there is any;
- for any location q_i of A_E, add an transition from q_{If} to q_i and associate it with a guard condition $n = i$;
- make q_{E0} and q_{Ef} as the initial and final locations of A, respectively.

In the example shown in Fig. 6, the two parts in dotted box labelled with A_I and A_E are the automata translated from the int fragment and its enclosing fragment, respectively. In a location of A_E, e.g., q_{E1}, the control of flow can transfer to A_I, provided that the interrupt mask x equals to 1. The assignment $n := 2$ on the transition from q_{E1} to q_{I0} ensures that when exiting from A_I, the control of flow would return to q_{E1} by taking the transition with guard condition $n = 2$.

5) Translation of Timing and Task Constraints: Besides the common timing constraints, the ISD offers the new notion of task constraints. Whereas timing constraints can be easily translated to clock constraints in timed automata [24], the task constraints are beyond the expressiveness of the classic timed automata. Suppose that we translate a task constraint $a \leq (e - e') \uparrow\leq b$ to a clock constraint. When the control of flow is not in the fragment to which e and e' belong, the clock needs to freeze, which means a change in the clock flow rate. Fortunately, the IA supports this change of rate for variables.

For each timing constraint $a \leq e - e' \leq b$, we generate one variable c representing a clock. We make the initialisation $c := 0$ at the transition with label e', and set a flow condition $\dot{c} := 1$ in every location. At the transition with label e, the value of c equals to the time duration $e - e'$, so we add the jump condition $a \leq c \leq b$ to the transition. Fig. 7 shows the same automaton as the one in Fig. 6 with the addition of the translation of timing and task constraints (assignments and conditions irrelevant to time are omitted). Suppose for the ISD in Fig. 6, there is a timing constraint $e_2 - e_1 \leq 10$ specifying that the deadline for the completion of transferring message m_1 should within 10 time units, and a task constraint $4 \leq (e_2 - e_1) \uparrow\leq 5$ specifying that the actual time for transferring m_1 takes 4 to 5 time units. Then for the timing constraint $e_2 - e_1 \leq 10$, we generate a variable c_1, initialise c_1 on the edge labelled with e_1, and add a constraint $c_1 \leq 10$ on the edge labelled with e_2. The flow condition of c_1 is set to $\dot{c}_1 := 1$ in all locations.

For each task constraint, we also generate a variable c' and initialise it to 0 at the transition labelled e'. Different than the timing constraints, the variable c' does not increase in all the locations but only in those translated from the interaction fragment to which e and e' belong (recall that the term interaction fragment can represent either the int fragment or the diagram excluding all int fragments). Therefore, locations translated from the interaction fragment to which e and e' belong are equipped with a flow condition $\dot{c}' := 1$; while other locations are equipped with a flow condition $\dot{c}' := 0$. In this way, when reaching the transition with label e, the value of the variable c' would represent the actual execution time that the control of flow stays in the same interaction fragment between the occurrences of e and e', so the jump condition $a \leq c' \uparrow\leq b$ can be added to this transition. For example, for the task constraint $4 \leq (e_2 - e_1) \uparrow\leq 5$, we generate a variable c_2, initialise c_2 on the edge labelled with e_1, and add a constraint $4 \leq c_2 \leq 5$ on the edge labelled with e_2. The flow conditions in locations q_{I0}, q_{I1} and q_{If} translated from the int fragment are set to $\dot{c}_2 := 0$, and in locations q_{E0}, q_{E1} and q_{Ef} translated from the enclosing fragment are set to $\dot{c}_2 := 1$.

C. Specification and Verification of Properties

Whereas the ISD models how the system behaves, the property to be checked is captured by property specification. For ISD, the property needs to specify both temporal orderings and time durations between events. Motivated by our goal of easy modelling

and verification, we propose a simple specification language defined as ⟨*spec_clause*⟩*::=*⟨e_1⟩'$\prec$'⟨e_2⟩|⟨*min*⟩'$\leq$'⟨e_2⟩'-'⟨e_1⟩'$\leq$'⟨*max*⟩, where ⟨e_1⟩*::=event name,* ⟨e_2⟩*::=event name* ($e_2 \neq e_1$)*,* ⟨*min*⟩*::=non-negative real,* ⟨*max*⟩ *::= non-negative real (greater than or equal to* ⟨*min*⟩*)* | '∞'. A specification is composed of one or more spec clauses. $e_1 \prec e_2$ specifies that e_1 occurs before e_2 in temporal order, and $min \leq e_2 - e_1 \leq max$ specifies that the time duration from the occurrence of e_1 to that of e_2 is between $[min, max]$ time units. The specification language is expressive in specifying various properties in interrupt-driven systems. For example, for the property of timeout freeness, one can use $0 \leq e_2 - e_1 \leq bound$ to specify that for a task starting by event e_1 and completing by event e_2, it shall not exceed the given time bound.

The verification is to check whether there is a behaviour in the IA can reach the final location and satisfies the negation of the properties, and if so, the properties are not satisfied and a counterexample represented by the behaviour is reported. For a property of the form $e_1 \prec e_2$, its negation is $e_2 \prec e_1$ and we first find all paths in which e_2 occurs before e_1 and then check if any behaviour of such paths can reach the final location. For a property of the form $min \leq e_2 - e_1 \leq max$, its negation is $e_2 - e_1 < min$ and $e_2 - e_1 > max$ (if $b \neq \infty$). We translate the negation to clock constraints, add them to the IA, and check if there is a behaviour reaches the final location. The checking of the existence of behaviours reaching the final location is essentially a reachability analysis problem, and can be solved by exploiting the existing hybrid automata checkers.

V. Experimental Evaluation

We implemented our approach as a prototype called *ISDChecker*, which supports the graphical modelling of ISD, ISD to IA translation and IA verification, and is available online[1]. The graphical modelling interface is based on UMLet [25] which is a free, open-source UML tool. The generated IA are passed to SpaceEx [18] which is a state-of-the-art tool for verifying safety properties of hybrid systems. Note that there is no restriction of the choice of hybrid automata checkers so long as the verification of IA is supported.

Our approach of modelling and verifying interrupt-driven systems with ISDs aims to guarantee the correctness of the systems. Ideally, it should be effective in finding counterexamples, and at the same time, easy to use. Therefore, our evaluation addresses the two following research questions:

- ***RQ1:*** *Can our approach effectively detect problems in interrupt-driven system models?*
- ***RQ2:*** *Is ISD easy to use compared with existing modelling languages for interrupt-driven systems?*

A. RQ1: Approach Efficacy

To conduct the experiments, we searched papers published after year 2010 using keywords "interrupt driven", "interrupt program", "interrupt softrware", "interrupt system", "embedded system" or "real-time system" combined with keywords "verification", "testing", "analysis" or "model checking", and collected 18 closely related to interrupt-driven system papers, from whose references we snowballed 7 more related papers published after year 2000. We studied these papers to collect cases that satisfy the following criteria: (1) they are not toy cases without realistic settings, and (2) they have been detailedly presented using models, programs and/or textual descriptions. As most papers just use 1 or 2 cases and many of them are toy cases, in the end 5 cases were collected from the existing literature, and we modelled these cases with ISD, which are shown in Column *case name* with references. To evaluate the effectiveness of ISDChecker, we expect the cases to contain problems so that we can check whether ISDChecker can find these problems. The case *ADC_Bug* itself has a data race problem. For the rest 4 cases, *medical_monitor* and *car_controller* have execution time bound requirements, and we modified the time values in ISD to make the requirements violated. The timeout problems in cases *attitude_display* and *fridge_controller* were manually inserted, as the original examples did not mention any temporal or time properties. We also consulted the experts in aerospace area and acquired 6 flawed designing cases. These cases were modelled with UML sequence diagram and we revised them with ISD. In total, 11 cases were studied, as shown in Table I. Case *satellite_spin* is the example in Sec. 2. Due to space limit, the ISDs for the rest cases are not presented here, but can be found online[1].

The experiments were conducted on a DELL PC with 3.4GHz Quad-Core CPU, 16GB RAM and OS of Ubuntu 16.04. The version of SpaceEx is 0.9.8f. The results are shown in Table I. For each of the 11 cases, ISDChecker was able to find a counterexample, whose error type is shown in Column *type*. In total, 3 *race* counterexamples and 8 *timeout* counterexamples are found, which are consistent with the ones that identified by manual inspection. This confirms that ISDChecker can effectively find problems in the design models by exhausting the state space. We did not compare our approach with other approaches such as model checking with ITA, since our approach is the only one targeting sequence diagram based models, and moreover, there were no tools for model checking interrupt-driven systems available at the time of this writing. For the evaluation of efficiency, we have recorded the sizes of the ISDs and the corresponding IA, and the translation time and checking time. The size of the ISD includes the number of entities (*# entity*), messages (*# msg.*), constraints (*# cons.*), int fragments (*# int.*) and different priorities of the int fragments (*# prior.*). The size of the IA includes the number of the total locations (*# loc.*), transition (*# trans.*), and variables (*# var.*). The verification time consists of two parts: one from the translation of models (*translation*), and the other from the execution of the checker (*check*). The cases are arranged from top to bottom in the table in the ascending order of first *# int.* and then *#msg.*. The number of int fragments in an ISD has a dominant impact on the number of transitions in the corresponding IA, as it is possible for each location translated from a non-int fragment to have transitions connecting locations translated from an int fragment. The

[1]https://github.com/isdchecker

TABLE I
EXPERIMENTAL RESULTS OF VERIFYING INTERRUPT-DRIVEN CASES

case name	type	#entity	# msg.	# cons.	# int.	# prior.	# loc.	# trans.	# var.	translate (ms)	check (s)
ADC_Bug[5]	race	4	5	6	1	1	12	28	10	12	0.05
fridge_controller[26]	timeout	7	8	8	1	1	18	38	12	15	0.40
altitude_display[27]	timeout	6	9	5	1	1	20	36	10	16	0.16
medical_monitor[28]	timeout	7	9	11	1	1	20	44	15	16	3.78
time_sync	race	4	11	4	1	1	30	80	11	73	1.24
orbit_upload	race	5	6	6	2	1	21	82	16	58	19.7
backup_computing	timeout	4	6	7	2	2	16	56	14	62	1.70
system_tick	timeout	4	7	6	2	2	17	56	12	16	1.62
satellite_spin	timeout	4	7	7	2	2	21	76	14	24	5.01
car_controller[2]	timeout	8	11	9	3	2	31	170	18	96	1879
task_rotate	timeout	11	14	7	3	1	80	560	15	123	2751

TABLE II
COMPARISON OF FAMILIARITY DEGREES WITH ISD AND ITA

familiar with	none of SD	simple SD	fragments
none of FSM	$2(2,0,0)$/I	$6(5,1,0)$/II	$5(4,1,0)$/III
FSM	$0(0,0,0)$/IV	$5(2,3,0)$/V	$7(1,5,1)$/VI
HA	$0(0,0,0)$/VII	$0(0,0,0)$/VIII	$1(0,0,1)$/IX

TABLE III
COMPARISON OF USABILITY DEGREES OF ISD AND ITA

Group	I		II		III		V		VI		IX		Total	
no./time(s)	n	t	n	t	n	t	n	t	n	t	n	t	n	t
1st:ISD	0	$*$	2	16	4	12	4	15	5	13	1	10	16	13
1st:ITA	0	$*$	0	$*$	0	$*$	0	$*$	1	47	1	36	2	42
2nd:ISD	1	14	4	13	4	11	3	13	7	12	1	10	20	12
2nd:ITA	0	$*$	0	$*$	0	$*$	2	46	2	41	1	27	5	40

* : the average time is not computed since no subjects have completed the task.

number of messages has a dominant impact on the number of locations, since events and their partial orders, which decides the number of locations, are deduced from messages. As we can see from Table I, the number of each kind of elements in a case's ISD is all below or around 10, which is clearly manageable for human designers, however, a small increase in the scale of an ISD can result in a significant increase of the corresponding IA. For example, a moderate sized ISD for the *satellite_spin* example in Fig. 2 has been translated to an automaton with 21 locations and 76 transitions; and with just 1 more int fragment and 7 more messages, the *task_rotate* case is translated to an automaton with 80 locations and 560 transitions. Nevertheless, our translation process only takes 12-123 ms, which indicates that the translation cost can almost be negligible. By exploiting the advanced hybrid automata checker, the checking time is also acceptable for taking several seconds for moderate sized cases and 30-40 minutes for large sized ones to exhaust all state space.

From the above discussion, we can derive the answer to RQ1: *Our approach can effectively detect design defects for interrupt-driven systems in an efficient way.*

B. RQ2: Approach Usability

In addressing RQ1 we have seen that the IA translated from the ISD is much larger in size. It is natural to think that modelling with ITA would be more difficult, as larger models are often less manageable. In order to have a more convincing result about the usability of our approach, we conducted a controlled experiment. It is known that before users can use formal methods, they must be well trained [29]. The training cost not only affects the users' decision upon whether to use the method, but also reflects the usability. Thus, one of our goals is to evaluate the training cost of ISD compared with ITA. The other goal is to compare the time cost to correctly develop a specification with ISD and ITA. The subjects of the experiments are students in our department consisting of 14 undergrads, 10 graduate students and 2 Ph.D. students. First, we profiled the subjects for their prior knowledge. We asked the subjects to first answer two questions: Are you familiar with the basic knowledge of sequence diagram (without fragments)/ finite state machines? If any answer is yes, then following up questions are asked: Are you familiar with the concepts of fragments/ hybrid automata? Based on the answers, each subject was assigned to one of the nine groups, as shown in Table II (subjects familiar with fragments are also familiar with simple SD, and those familiar with HA are also familiar with FSM). In each cell, the first number is the total number of subjects, and the three numbers in the parentheses are the number of undergrads, master students and Ph.D. students, respectively. The Roman number after the slash is the label assigned to each group. As shown in Table II, among the 26 subjects, 24 subjects are familiar with sequence diagrams, and 13 subjects have learned the usage of fragments; whereas only half subjects know how to model with finite state machine, and only 1 subject has used hybrid automata.

To evaluate the training cost, we prepared 2 training sessions. Each session lasts 2 hours: the first 1 hour for ISD and the second 1 hour for ITA. We were aware that both modelling language target the interrupt-driven systems and share some common background knowledge, which makes it possible for the language trained later to benefit from the earlier training of the other one. In view of that, we made the schedule to have ISD trained first, to reduce the preferences that the result may have towards ISD. After the first session, we gave the subjects a description on the example in Fig. 2, and asked them to model the system. Again, since it was the same modelling task, we ask the subjects to model with ISD first to reduce the results' preference towards ISD. We carefully manual-checked the models and collected the number of subjects who correctly designed the models (Column *n*) in each group (Group IV, VII and VIII have no subjects and are not considered). We also recorded the modelling time of the subjects who gave correct models in each group, and present the average time

(in seconds) in Table III (Column *t*). After the first training session, 16 subjects can have correct models with ISD, while with ITA only 2 subjects can complete the modelling correctly. Then we gave a more elaborative training session, and assigned the subjects a new modelling task of Case *car_control* used in the experiment for RQ1. This time, 20 out of 24 subjects can model with ISD, but still only 5 subjects can model with ITA. From these data we can conclude that one needs less training time to master ISD than ITA, to which two factors contribute most. The first one is that most people have experience using sequence diagrams in system modelling, so little training is needed for the basics. The second factor is our ISDs follow the standard sequence diagram notations, and therefore do not require much effort to learn. On the other hand, ITA is a special kind of hybrid automata, which can be difficult to comprehend in a short time. We also compared the modelling time with ISD and ITA. In Table III, the average time of using ISD is about 12 to 13 minutes, while the average time of using ITA needs more than 40 minutes. The reason, as the subjects reported, is that they can design the ISD model by following a frequent scenario, and do not have to worry about the interrupt nesting and unpredictable occurrence.

From the above discussion, we can derive the answer to RQ2: *ISD is easier to learn and use compared to the automata theory based modelling language.*

VI. Related Work

We discuss some closely related work concerning the modelling and correctness checking for interrupt-driven systems.

Modelling. Prioritised preemption and interrupt mask register (IMR) are the most unique features in interrupt-driven systems, and most modelling approaches attempt to support them. Some of them are variants of flow graphs with extensions to model prioritised preemption and IMR [30], [31], [32]. In [30], a directed graph called interrupt preemption graph is proposed where each edge corresponds to a potential preemption by an interrupt handler, and IMR is exploited to remove unreachable branches in the graph. Time is not considered in these studies, whereas the behaviour of interrupt-driven systems largely depend on the timing properties.

Timed automata [24], as a mature model for specifying timing requirements though, lack the feature of time suspension which is critical to model the executions preempted by interrupts. Hybrid automata, or more specifically the subclasses such as integration automata, stopwatch automata [33] or suspension automata [34] are sufficiently expressive to model time suspension and resumption, but lack the mechanism to model interrupt priority. Therefore, in [10], [11], Interrupt Timed Automata (ITA) is proposed, where the states are organised according to interrupt priorities, ranging from 1 to n, with one active clock that can be suspended for one priority.

Whereas these models are visual and graphic, some others are based on calculus. In [35], the algebra of communicating processes (ACP) is augmented with priorities and non-deterministic choice to describe the working of interrupts. Work [36] studies the "interrupt driven round robin system" where tasks run in round robin scheduling and interrupt service routines perform urgent actions, and proposes to model the system with a variant of Event B. Work [37] provides a calculus for reasoning about interrupt-driven systems and a type system for checking stack boundedness. Compared with these existing modelling methods using automata theory or calculus, the ISD is more friendly to users. Furthermore, the simplicity of ISD does not sacrifice its expressiveness. With the int fragment, mask variable and task constraint, the unpredictable and prioritised preemptive behaviour and the time suspension and resumption can be easily modelled.

Correctness assurance. Testing has been widely used to create reliable embedded software [38], [5]. In [5], an interrupt scheduler is proposed to fire interrupts at specific time points and prohibit the firings of an interrupt at a time when the system cannot handle it properly during the testing. However, interrupt-driven software is hard to be thoroughly tested since it usually contains a very large number of executable paths. Therefore, in [4], test adequacy criteria are introduced to measure the quality of test suites that test interrupt-driven applications. Both studies target nesC applications in TinyOS, as the simple scheduling policy adopted by TinyOS makes the interleaving between tasks more tractable [4].

Different from testing, static analysis and verification of program codes focus on specific code problems, and most of them study the subject of stack size analysis [39], [3], [30], since the execution of too many interrupt handlers could use up stack space. Data inconsistency is of interest as well, due to the concurrency induced by interrupts. Work [40] analyses data race and transactional behaviour of procedures for interrupt-driven programs synchronised via the priority ceiling protocol. Work [41] proposes to sequentialise interrupt-driven programs into sequential programs, and using existing numerical analysis tools. In [42] nesC programs are transformed to POSIX threads programs for checking race conditions. Timing analysis is also considered, although most studies are restricted to worst case execution time analysis [2], such as maximum interrupt latency [39]. Experiments in [39] show that static deadline analysis cannot work without information about the behaviour of external devices, so in [43] static analysis is combined with testing to analyse whether every interrupt can be handled before the deadline. Static analysis or verification of program codes need to work on specific languages, such as codes for the Z86 architecture [39], [43], [3], or for some Atmel's architectures [32], [30]. Differently, our work chooses the model checking approach in which systems are modelled before verification. The closest work to ours is [10], [11], where the interrupt-driven systems are modelled by ITA. Differently, we propose ISDs for easy modelling and translate ISDs to IA for checking.

VII. Conclusion

In this paper, we introduce a novel approach to modelling and verifying interrupt-driven systems. We propose the ISD by extending the UML sequence diagram with interrupt fragment, mask variable and task constraint, to model the unpredictable

and preemptive system behaviour. Following our translation algorithms, the ISD can be automatically translated to IA, which can be checked by existing checkers for various properties. Experiments on previous studied cases and real-world aerospace applications consistently confirm our approach's effectiveness. In another experiment, the performance of the participants shows that both the training cost and usage cost of ISD are lower than the automata-based modelling language ITA, indicating ISD's good usability. We have implemented the proposed approach as a prototype tool, in the hope that it could become a powerful assistant to system designers.

REFERENCES

[1] D. Kroening, L. Liang, T. Melham, P. Schrammel, and M. Tautschnig, "Effective verification of low-level software with nested interrupts," in *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, ser. DATE '15. San Jose, CA, USA: EDA Consortium, 2015, pp. 229–234. [Online]. Available: http://dl.acm.org/citation.cfm?id=2755753.2755803

[2] J. Kotker, D. Sadigh, and S. A. Seshia, "Timing analysis of interrupt-driven programs under context bounds," in *Proceedings of the International Conference on Formal Methods in Computer-Aided Design*, ser. FMCAD '11. Austin, TX: FMCAD Inc, 2011, pp. 81–90. [Online]. Available: http://dl.acm.org/citation.cfm?id=2157654.2157670

[3] K. Chatterjee, D. Ma, R. Majumdar, T. Zhao, T. A. Henzinger, and J. Palsberg, "Stack size analysis for interrupt-driven programs," *Information and Computation*, vol. 194, no. 2, pp. 144 – 174, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0890540104001063

[4] Z. Lai, S. C. Cheung, and W. K. Chan, "Inter-context control-flow and data-flow test adequacy criteria for nesc applications," in *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. SIGSOFT '08/FSE-16. New York, NY, USA: ACM, 2008, pp. 94–104. [Online]. Available: http://doi.acm.org/10.1145/1453101.1453115

[5] J. Regehr, "Random testing of interrupt-driven software," in *Proceedings of the 5th ACM International Conference on Embedded Software*, ser. EMSOFT '05. New York, NY, USA: ACM, 2005, pp. 290–298. [Online]. Available: http://doi.acm.org/10.1145/1086228.1086282

[6] P. G. Frankl and E. J. Weyuker, "An applicable family of data flow testing criteria," *IEEE Trans. Softw. Eng.*, vol. 14, no. 10, pp. 1483–1498, Oct. 1988. [Online]. Available: http://dx.doi.org/10.1109/32.6194

[7] M. J. Harrold and M. L. Soffa, "Efficient computation of interprocedural definition-use chains," *ACM Trans. Program. Lang. Syst.*, vol. 16, no. 2, pp. 175–204, Mar. 1994. [Online]. Available: http://doi.acm.org/10.1145/174662.174663

[8] Y. Lei and R. H. Carver, "Reachability testing of concurrent programs," *IEEE Trans. Softw. Eng.*, vol. 32, no. 6, pp. 382–403, Jun. 2006. [Online]. Available: http://dx.doi.org/10.1109/TSE.2006.56

[9] C.-S. D. Yang, A. L. Souter, and L. L. Pollock, "All-du-path coverage for parallel programs," in *Proceedings of the 1998 ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA '98. New York, NY, USA: ACM, 1998, pp. 153–162. [Online]. Available: http://doi.acm.org/10.1145/271771.271804

[10] B. Bérard, S. Haddad, and M. Sassolas , "Real time properties for interrupt timed automata," in *Proceedings of the 2010 17th International Symposium on Temporal Representation and Reasoning*, ser. TIME '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 69–76. [Online]. Available: http://dx.doi.org/10.1109/TIME.2010.11

[11] B. Bérard, S. Haddad, and M. Sassolas, "Interrupt timed automata: Verification and expressiveness," *Form. Methods Syst. Des.*, vol. 40, no. 1, pp. 41–87, Feb. 2012. [Online]. Available: http://dx.doi.org/10.1007/s10703-011-0140-2

[12] F. Cassez and K. G. Larsen, "The impressive power of stopwatches," in *Proceedings of the 11th International Conference on Concurrency Theory*, ser. CONCUR '00. London, UK, UK: Springer-Verlag, 2000, pp. 138–152. [Online]. Available: http://dl.acm.org/citation.cfm?id=646735.701625

[13] OMG, "Uml2.0 superstructure specification," *Available at http://www.uml.org*, 2005.

[14] B. Dobing and J. Parsons, "How uml is used," *Commun. ACM*, vol. 49, no. 5, pp. 109–113, May 2006. [Online]. Available: http://doi.acm.org/10.1145/1125944.1125949

[15] R. Alur, C. Courcoubetis, T. A. Henzinger, and P.-H. Ho, "Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems," in *Hybrid Systems*. London, UK, UK: Springer-Verlag, 1993, pp. 209–229. [Online]. Available: http://dl.acm.org/citation.cfm?id=646874.709849

[16] Y. Kesten, A. Pnueli, J. Sifakis, and S. Yovine, "Integration graphs: A class of decidable hybrid systems," in *Hybrid Systems*. London, UK, UK: Springer-Verlag, 1993, pp. 179–208. [Online]. Available: http://dl.acm.org/citation.cfm?id=646874.709977

[17] B. I. Silva, K. Richeson, B. Krogh, and A. Chutinan, "Modeling and verifying hybrid dynamic systems using checkmate," in *Proceedings of 4th International Conference on Automation of Mixed Processes*. Shaker Publisher, 2000, pp. 323–328.

[18] G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, "Spaceex: Scalable verification of hybrid systems," in *Proceedings of the 23rd International Conference on Computer Aided Verification*, ser. CAV'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 379–395. [Online]. Available: http://dl.acm.org/citation.cfm?id=2032305.2032335

[19] G. Audemard, M. Bozzano, A. Cimatti, and R. Sebastiani, "Verifying industrial hybrid systems with mathsat," *Electron. Notes Theor. Comput. Sci.*, vol. 119, no. 2, pp. 17–32, Mar. 2005. [Online]. Available: http://dx.doi.org/10.1016/j.entcs.2004.12.022

[20] S. Gao, S. Kong, and E. M. Clarke, "dreal: An smt solver for nonlinear theories over the reals," in *Proceedings of the 24th International Conference on Automated Deduction*, ser. CADE'13. Berlin, Heidelberg: Springer-Verlag, 2013, pp. 208–214. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-38574-2_14

[21] T. Firley, M. Huhn, K. Diethers, T. Gehrke, and U. Goltz, "Timed sequence diagrams and tool-based analysis: A case study," in *Proceedings of the 2Nd International Conference on The Unified Modeling Language: Beyond the Standard*, ser. UML'99. Berlin, Heidelberg: Springer-Verlag, 1999, pp. 645–660. [Online]. Available: http://dl.acm.org/citation.cfm?id=1767297.1767363

[22] S. Uchitel, J. Kramer, and J. Magee, "Incremental elaboration of scenario-based specifications and behavior models using implied scenarios," *ACM Trans. Softw. Eng. Methodol.*, vol. 13, no. 1, pp. 37–85, Jan. 2004. [Online]. Available: http://doi.acm.org/10.1145/1005561.1005563

[23] A. Knapp and J. Wuttke, "Model checking of uml 2.0 interactions," in *Proceedings of the 2006 International Conference on Models in Software Engineering*, ser. MoDELS'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 42–51. [Online]. Available: http://dl.acm.org/citation.cfm?id=1762828.1762836

[24] R. Alur and D. L. Dill, "A theory of timed automata," *Theor. Comput. Sci.*, vol. 126, no. 2, pp. 183–235, Apr. 1994. [Online]. Available: http://dx.doi.org/10.1016/0304-3975(94)90010-8

[25] M. Auer, T. Tschurtschenthaler, and S. Biffl, "A flyweight uml modelling tool for software development in heterogeneous environments," in *Proceedings of the 29th Conference on EUROMICRO*, ser. EUROMICRO '03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 267–272. [Online]. Available: http://dl.acm.org/citation.cfm?id=942796.943259

[26] F. Pereira, F. Moutinho, and L. Gomes, "Model-checking framework for embedded systems controllers development using iopt petri nets," in *2012 IEEE International Symposium on Industrial Electronics*, May 2012, pp. 1399–1404.

[27] C. Fidge and P. Cook, "Model checking interrupt-dependent software," in *Proceedings of the 12th Asia-Pacific Software Engineering Conference*, ser. APSEC '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 51–58. [Online]. Available: http://dx.doi.org/10.1109/APSEC.2005.80

[28] L. A. Cortes, P. Eles, and Z. Peng, "Formal coverification of embedded systems using model checking," in *Proceedings of the 26th Euromicro Conference. EUROMICRO 2000. Informatics: Inventing the Future*, vol. 1, Sept 2000, pp. 106–113 vol.1.

[29] A. Hall, "Seven myths of formal methods," *IEEE Softw.*, vol. 7, no. 5, pp. 11–19, Sep. 1990. [Online]. Available: http://dx.doi.org/10.1109/52.57887

[30] J. Regehr, A. Reid, and K. Webb, "Eliminating stack overflow by abstract interpretation," *ACM Trans. Embed. Comput. Syst.*, vol. 4, no. 4, pp. 751–778, Nov. 2005. [Online]. Available: http://doi.acm.org/10.1145/1113830.1113833

[31] W. Le, J. Yang, M. L. Soffa, and K. Whitehouse, "Lazy preemption to enable path-based analysis of interrupt-driven code," in *Proceedings of the 2Nd Workshop on Software Engineering for Sensor Network Applications*, ser. SESENA '11. New York, NY, USA: ACM, 2011, pp. 43–48. [Online]. Available: http://doi.acm.org/10.1145/1988051.1988060

[32] B. Schlich, T. Noll, J. Brauer, and L. Brutschy, "Reduction of interrupt handler executions for model checking embedded software," in *Proceedings of the 5th International Haifa Verification Conference on Hardware and Software: Verification and Testing*, ser. HVC'09. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 5–20. [Online]. Available: http://dl.acm.org/citation.cfm?id=1965974.1965981

[33] T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya, "What's decidable about hybrid automata?" in *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing*, ser. STOC '95. New York, NY, USA: ACM, 1995, pp. 373–382. [Online]. Available: http://doi.acm.org/10.1145/225058.225162

[34] J. McManis and P. Varaiya, "Suspension automata: A decidable class of hybrid automata," in *Proceedings of the 6th International Conference on Computer Aided Verification*, ser. CAV '94. London, UK, UK: Springer-Verlag, 1994, pp. 105–117. [Online]. Available: http://dl.acm.org/citation.cfm?id=647763.735660

[35] J. Bergstra, J. Baeten, and J. W. Klop, "Syntax and defining equations for an interrupt mechanism in process algebra," *Fundamenta informaticae: quarterly*, vol. 9, pp. 127–167, 1986.

[36] B. Stoddart, D. Cansell, and F. Zeyda, "Modelling and proof analysis of interrupt driven scheduling," in *Proceedings of the 7th International Conference on Formal Specification and Development in B*, ser. B'07. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 155–170. [Online]. Available: http://dx.doi.org/10.1007/11955757_14

[37] J. Palsberg and D. Ma, "A typed interrupt calculus," in *Proceedings of the 7th International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems: Co-sponsored by IFIP WG 2.2*, ser. FTRTFT '02. London, UK, UK: Springer-Verlag, 2002, pp. 291–310. [Online]. Available: http://dl.acm.org/citation.cfm?id=646847.707120

[38] B. M. Broekman, *Testing Enbredded Software*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2002.

[39] D. Brylow, N. Damgaard, and J. Palsberg, "Static checking of interrupt-driven software," in *Proceedings of the 23rd International Conference on Software Engineering*, ser. ICSE '01. Washington, DC, USA: IEEE Computer Society, 2001, pp. 47–56. [Online]. Available: http://dl.acm.org/citation.cfm?id=381473.381478

[40] M. D. Schwarz, H. Seidl, V. Vojdani, P. Lammich, and M. Müller-Olm, "Static analysis of interrupt-driven programs synchronized via the priority ceiling protocol," in *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL '11. New York, NY, USA: ACM, 2011, pp. 93–104. [Online]. Available: http://doi.acm.org/10.1145/1926385.1926398

[41] X. Wu, L. Chen, A. Miné, W. Dong, and J. Wang, "Numerical static analysis of interrupt-driven programs via sequentialization," in *Proceedings of the 12th International Conference on Embedded Software*, ser. EMSOFT '15. Piscataway, NJ, USA: IEEE Press, 2015, pp. 55–64. [Online]. Available: http://dl.acm.org/citation.cfm?id=2830865.2830872

[42] J. Regehr and N. Cooprider, "Interrupt verification via thread verification," *Electron. Notes Theor. Comput. Sci.*, vol. 174, no. 9, pp. 139–150, Jun. 2007. [Online]. Available: http://dx.doi.org/10.1016/j.entcs.2007.04.002

[43] D. Brylow and J. Palsberg, "Deadline analysis of interrupt-driven software," in *Proceedings of the 9th European Software Engineering Conference Held Jointly with 11th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. ESEC/FSE-11. New York, NY, USA: ACM, 2003, pp. 198–207. [Online]. Available: http://doi.acm.org/10.1145/940071.940098