Automated Analyses of IOT Event Monitoring Systems

Andrew Apicelli[1], Sam Bayless[1][0000−0002−0909−8986], Ankush
Das[1][0000−0003−2459−1258], Andrew Gacek[1][0000−0003−0321−8155], Dhiva
Jaganathan[1], Saswat Padhi[2][§], Vaibhav Sharma[3][0000−0001−9877−8926], Michael
W. Whalen[1][0000−0003−3824−1435], and Raveesh Yadav[1]

[1] Amazon Web Services, Inc.
{apicea,sabayles,daankus,gacek,dhivasj,mww,raveeshy}@amazon.com
[2] Google LLC
spadhi@google.com
[3] Amazon.com Services LLC
svaib@amazon.com

Abstract. AWS IoT Events is an AWS service that makes it easy to
respond to events from IoT sensors and applications. *Detector models* in
AWS IoT Events enable customers to monitor their equipment or device
fleets for failures or changes in operation and trigger actions when such
events occur. If these models are incorrect, they may become out-of-sync
with the actual state of the equipment causing customers to be unable
to respond to events occurring on it.
Working backwards from common mistakes made when creating de-
tector models, we have created a set of automated analyzers that al-
low customers to prove their models are free from six common mis-
takes. Our analyzers have been running in the AWS IoT Events pro-
duction service since December 2021. Our analyzers check six correct-
ness properties in the production service in real time. 93% of customers
of AWS IoT Events have run our analyzers without needing to have any
knowledge of them. Our analyzers have reported property violations in
22% of submitted detector models in the production service.

1   Introduction

AWS IoT Events is a managed service for managing fleets of IoT devices. Cus-
tomers use AWS IoT Events in diverse use cases such as monitoring self-driving
wheelchairs, monitoring a device's network connectivity, humidity, temperature,
pressure, oil level, and oil temperature sensing. Customers use AWS IoT Events by
creating a *detector model* that detects events occurring on IoT devices and no-
tifies an external service so that a corrective action can be taken. An example is
an industrial boiler which constantly reports its temperature to a detector. The
detector tracks the boiler's average temperature over the past 90 minutes and
notifies a human operator when it is running too hot.

[§] Work done while at Amazon Web Services, Inc.

Each detector model is defined as a finite state machine with dynamically typed variables and timers, where timers allow detectors to keep track of state over time. A model processes inputs from IoT devices to update internal state and to notify other AWS services when events are detected. Customers can use a single detector model to instantaneously detect events in thousands of devices. Ensuring well-formedness of a detector model is crucial as ill-formed detector models can miss events in *every* monitored device.

Starting from a survey that identified sources of well-formedness problems in customer models, we identified most common mistakes made by customers and detect them using type- and model-checking. To use a model-checker for checking well-formedness of a detector model, we formalize the execution semantics of a detector model and translate this semantics into the source-language notation of the JKind model checker [1]. Model checking [2–9] verifies desirable properties over the behavior of a system by performing the equivalent of an exhaustive enumeration of all the states reachable from its initial state. Most model checking tools use *symbolic encodings* and some form of *induction* [6] to prove properties of very large finite or even infinite state spaces.

We have implemented type-checking and model-checking as an analysis feature in the production AWS IoT Events service. Our analyzers have reported well-formedness property violations in 22% of submitted detector models. 93% of customers of AWS IoT Events have checked their detector models using our analyzers. Our analyzers report property violations to customers with an average latency of 5.6 seconds (see Section 4).

Our contributions are as follows:

1. We formalize the semantics of AWS IoT Events detector models.
2. We identify six well-formedness properties whose violations detect common customer mistakes.
3. We create fast, push-button analyzers that report property violations to customers.

2   Overview

Consider a user of AWS IoT Events who wants to monitor the temperature of an industrial boiler. If the industrial boiler overheats, it can cause fires and endanger human lives. To detect an early warning of an overheating event, they want to automatically identify two different alarming events on the boiler's temperature. They want their first alarm to be triggered if the boiler's reported temperature is outside the normal range for more than 1 minute. They want their second alarm to be triggered if the temperature is outside the normal range for another 5 minutes after the first alarm.

A user might try to implement these requirements by creating the (flawed) detector model shown in Figure 1. This detector receives temperature data from the boiler and responds by sending a text message to the user. The detector model contains four states:
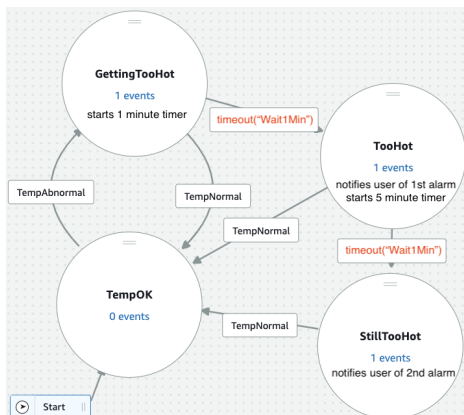
Fig. 1: AWS IoT Events detector model with two alarms (buggy version)

Fig. 2: An action in the detector model from Figure 1

- TempOK: starting state of the detector model. The detector stays in this state as long as the boiler's temperature lies in a normal range. The detector transitions from TempOK to GettingTooHot on detecting a temperature outside normal range, indicated by TempAbnormal.
- GettingTooHot: detector starts a 1 minute timer and transitions back to TempOK if the boiler cools down. When the timer expires, it transitions to TooHot.
- TooHot: detector first notifies the user of the 1st alarm. It then starts a 5 minute timer and transitions back to TempOK if the boiler cools down. When the 5 minute timer expires, it transitions to StillTooHot.
- StillTooHot: detector notifies user of the 2nd alarm.

Understanding the bug: Every state in the detector model consists of *actions*. An action changes the internal state of a detector or triggers an external service. For example, the GettingTooHot state consists of an action that starts a timer. The user can edit these actions with an interface shown in Figure 2. This action starts a one minute timer named Wait1Min. Note that timers are accessible from every state in the detector model. Even though the Wait1Min timer is created in the GettingTooHot state of Figure 1, it can be checked for expiration in all the four states of Figure 1.

The detector model in Figure 1 has a fatal flaw based on a typo. The user has written timeout("Wait1Min") instead of timeout("Wait5Min") when transitioning out of TooHot. This is allowed as timers are globally referenceable. However, it is a bug because each global timer has a unique name and the Wait1Min timer has already been used and expired. This makes StillTooHot unreachable, meaning the 2nd alarm won't ever fire, since a timer can expire at most once.

Related Work Languages such as IOTA [10], SIFT [11], and the system from Garcia et. al [12] use *trigger-condition-action rules* [13] to control the behavior of

internet of things applications. These languages have the benefit of being largely declarative, allowing users to specify desired actions under different environmental stimuli. Similar to our approach, SIFT [11] automatically removes common user mistakes as well as compiles specifications into controller implementations without user interaction, and IOTA [10] is a reasoning calculus that allows custom specifications to be written both about why something *should* or *should not* occur. AWS IoT Events is designed explicitly for monitoring, rather than control, and our approach is imperative, rather than declarative: detector models do not have the same inconsistencies as rule sets, as they are disambiguated using explicit priorities on transitions. On the other hand, customers may still construct machines that do not match their intentions, motivating the analyses described in this paper.

3   Technique

In this section, we present a formal execution semantics of an AWS IoT Events detector model and describe specifications for the correctness properties.

Formalization of Detector Models Defining the alphabet and the transition relation for the state machine is perhaps the most interesting aspect of our formalization. Since detector models may contain global timers, *timed automata* [14] might seem like an apt candidate abstraction. However, AWS IoT Events users are not allowed to change the clock frequency of timers, nor specify arbitrary *clock constraints*. These observations allow us to formalize the detector models as a regular state machine, with timeout durations as additional state variables.

Formally, we represent the state machine for a detector model $\mathbf{M}$ as a tuple $\langle \mathbf{S}, \mathbf{S}_0, \mathbf{I}, \mathbf{G}, \mathbf{T}, \mathcal{E}_E, \mathcal{E}_X, \mathcal{E}_I \rangle$, where:

- $\mathbf{S}$: finite set of states in the FSM,
- $\mathbf{S}_0 \subseteq \mathbf{S}$: set of *initial* state(s),
- $\mathbf{I}$: set of input variables assigned by the environment
- $\mathbf{G}$: set of global variables assigned by the state machine
- $\mathbf{T}$: set of timer variables that are reset by the model and updated as time evolves in the environment
- $\mathcal{E}_E : \mathbf{S} \to \kappa$ `list`: mapping from states to a (possibly empty) list of entry events to be performed when entering a state. κ describes an *event*, further explained in the description of the grammar.
- $\mathcal{E}_X : \mathbf{S} \to \kappa$ `list` is a mapping from states to a list of exit events to be performed when exiting a state.
- $\mathcal{E}_I : \mathbf{S} \to (\kappa$ `list` $\times \mu$ `list`): mapping from states to a list of input events, including transitions to other states.

It is assumed that the sets $\mathbf{I}$, $\mathbf{G}$, and $\mathbf{T}$ are pairwise disjoint, and we define the set $\mathbf{V} \triangleq \mathbf{I} \cup \mathbf{G}$ to represent input and global variables in the model.

We denote by $\mathbb{V}$ the set of values for global ($\mathbf{G}$) and input ($\mathbf{I}$) variables; $\mathbb{V}$ ranges over the values of primitive types: integers, decimals (rationals), booleans,

$$\tau ::= \texttt{int} \mid \texttt{dec} \mid \texttt{str} \mid \texttt{bool}$$

$$\epsilon ::= e_0 \; bop \; e_1 \mid uop \; e_0 \mid l \mid v \mid \texttt{timeout}(t) \mid \texttt{isundefined}(v) \mid \ldots$$

$$\alpha ::= \texttt{setTimer}(t, e) \mid \texttt{resetTimer}(t)$$
$$\mid \texttt{clearTimer}(t) \mid \texttt{setGlobal}(g, e)$$

$$\kappa ::= \texttt{event}(e, a*)$$

$$\mu ::= \texttt{transition}(e, a*, s)$$

$$\iota ::= \texttt{message}(i, v) \mid \texttt{timeout}(t)$$

Fig. 3: Types, expressions, actions, and events in IoT Events Detector Models

and strings. Integers and rationals are assumed to be unbounded, and rationals are arbitrarily precise. We use $\mathbb{N}$ as the domain for time and timeout values. Sets $\mathbb{V}^{\perp}$ and $\mathbb{N}^{\perp}$ are extended with the value $\perp$ to represent an *uninitialized* variable.

The grammar for types (τ), expressions (ϵ), actions (α), events (κ), transitions (μ) and input triggers (ι) is shown in Figure 3. In the grammar, metavariable e stands for an expression, l stands for a literal value in $\mathbb{V}$, v stands for any variable in $\mathbf{V}$, t is a timer variable in $\mathbf{T}$, a is an action, and i is an input in $\mathbf{I}$. The unary and binary operators include standard arithmetic, Boolean, and relational operators. The `timeout` expression is true at the instant timer t expires, and the `isundefined` expression returns true if the variable or timer in question has not been assigned. Actions (α) describe changes to the system state: `setTimer` starts a timer and sets the periodicity of the timer, while the `resetTimer` and `clearTimer` reset and clear a timer (without changing the periodicity of the timer). The `setGlobal` action assigns a global variable. Events (κ) describe conditions under which a sequence of actions occur.

We define configurations $\mathbf{C}$ for the state machine as:

$$\mathbf{C} \triangleq \mathbf{S} \times (\mathbf{I} \to \mathbb{V}^{\perp}) \times (\mathbf{T} \to (\mathbb{N}^{\perp} \times \mathbb{N}^{\perp})) \times (\mathbf{G} \to \mathbb{V}^{\perp})$$

Each configuration $C = \langle s, i, t, g \rangle$ tracks the following:

- a state $s \in \mathbf{S}$ in the detector model,
- the input valuation $i \in (\mathbf{I} \to \mathbb{V}^{\perp})$ containing the values of inputs,
- the timer valuation $t \in (\mathbf{T} \to (\mathbb{N}^{\perp} \times \mathbb{N}^{\perp}))$ for user-defined timers. Each timer has both a periodicity and (if active) a time remaining, and
- the global valuation $g \in (\mathbf{G} \to \mathbb{V}^{\perp})$ for global variables in the detector model.

Example 1. Consider a corrected version of our example detector model from Figure 1 which has two timers, `Wait1Min` and `Wait5Min`, and no global variables. Some examples of configurations for this model are:

- $\langle \texttt{TempOK}, \{\texttt{temp} : \perp\}, \{\texttt{Wait1Min} : (\perp, \perp), \texttt{Wait5Min} : (\perp, \perp)\}, \{\} \rangle$ is the initial configuration. The model contains input `temp`, timers `Wait1Min` and `Wait5Min`, and no global variables. As no variables or timers have been assigned, all variables have value undefined ($\perp$).

$$\boxed{\begin{aligned}
&C \vdash_\epsilon e \to v \\
&C \vdash_\alpha a \to C' \qquad C \vdash_{\alpha*} al \to C' \\
&C \vdash_\kappa k \to C' \qquad C \vdash_{\kappa*} kl \to C' \\
&C \vdash_{\mu*} ml \to C' \qquad C \vdash_{\mathcal{E}_I} \mathcal{E}_I \to C' \\
&\vdash_\iota C \xmapsto{i} C'
\end{aligned}}$$

$$\frac{\langle s,i,t,g \rangle \vdash_\epsilon e \to v}{\langle s,i,t,g \rangle \vdash_\alpha \texttt{setTimer}(tr,e) \to \langle s,i,t[tr \leftarrow (v,v)],g \rangle}$$

$$\frac{t(tr)=(p,v)}{\langle s,i,t,g \rangle \vdash_\alpha \texttt{resetTimer}(tr) \to \langle s,i,t[tr \leftarrow (p,p)],g \rangle}$$

$$\frac{t(tr)=(p,v)}{\langle s,i,t,g \rangle \vdash_\alpha \texttt{clearTimer}(tr) \to \langle s,i,t[tr \leftarrow (p,\bot)],g \rangle}$$

$$\frac{\langle s,i,t,g \rangle \vdash_\epsilon e \to v}{\langle s,i,t,g \rangle \vdash_\alpha \texttt{setGlobal}(gv,e) \to \langle s,i,t,g[gv \leftarrow v] \rangle}$$

$$\frac{C \vdash_\epsilon e \to \texttt{false}}{C \vdash_\kappa \texttt{event}(e,al) \to C}$$

$$\frac{C \vdash_\epsilon e \to \texttt{true} \qquad C \vdash_{\alpha*} al \to C'}{C \vdash_\kappa \texttt{event}(e,al) \to C'}$$

$$\frac{}{C \vdash_{\mu*} \texttt{nil} \to C}$$

$$\frac{C \vdash_\epsilon e \to \texttt{false} \qquad C \vdash_{\mu*} tl \to C'}{C \vdash_{\mu*} \texttt{transition}(e,al,s') :: tl \to C'}$$

$$\frac{\begin{aligned}&C \vdash_\epsilon e \to \texttt{true} \qquad C \vdash_{\alpha*} al \to C' \\ &C' \vdash_{\kappa*} \mathcal{E}_X(C.s) \to C'' \\ &C''[s \leftarrow s'], ti \vdash_{\kappa*} \mathcal{E}_E(s') \to C'''\end{aligned}}{C \vdash_{\mu*} \texttt{transition}(e,al,s') :: tl \to C'''}$$

$$\frac{C \vdash_{\kappa*} kl \to C' \qquad C' \vdash_{\mu*} ml \to C''}{C \vdash_{\mathcal{E}_I} (kl,ml) \to C''}$$

$$\frac{\begin{aligned}&\texttt{matchesEarliest}(C.t, ti) \wedge \texttt{subtractTimers}(C, ti) \to C' \\ &C' \vdash_{\mathcal{E}_I} \mathcal{E}_I(C'.s) \to C'' \wedge \texttt{clearTimers}(C'') \to C'''\end{aligned}}{\vdash_\iota C \xmapsto{timeout(ti)} C'''}$$

$$\frac{\langle s,i[iv \leftarrow v],t,g \rangle \vdash_{\mathcal{E}_I} \mathcal{E}_I(C.s) \to C'}{\vdash_\iota \langle s,i,t,g \rangle \xmapsto{message(iv,v)} C'}$$

Fig. 4: Rules describing behavior of the system

— $\langle \texttt{TooHot}, \{\texttt{temp}:300\}, \{\texttt{Wait1Min}:(60,\bot), \texttt{Wait5Min}:(300,260)\}, \{\} \rangle$ is the configuration at global time t if the temperature is still beyond the normal range and we transition to the `TooHot` detector model state. Note the `Wait1Min` timer is no longer set whereas the `Wait5Min` timer has a periodicity of 300 and is set to expire at $t+260$.

To define the execution semantics, we create a structural operational semantics for each of the grammar rules and for the interaction with the external environment, as shown in Figure 4. We distinguish semantic rules by decorating the turnstiles with the grammar type that they operate over ($\epsilon, \alpha, \kappa, \mu, \mathcal{E}_I$, and ι). The variables e, a, k, m, i stand in for elements of the appropriate syntactic class defined by the turnstile. For lists of elements, we decorate the syntactic class with * (e.g. $\vdash_{\alpha*}$), and the variables with 'l' (e.g. al). We use the following notation conventions: Given $C = \langle s,i,t,g \rangle$, we say $C.s = s$, and similarly with the other components of C. We also say $C[s \leftarrow s']$ is equivalent to $\langle s',i,t,g \rangle$, and similarly with the other components of C.

Expressions ($\vdash_\epsilon$) evaluate to values, given a configuration. We do not present expression rules (they are simple), but illustrate the other rule types in Figure 4. For actions ($\vdash_\alpha$), the *setTimer* rule establishes the periodicity of a timer and also starts it. The *resetTimer* and *clearTimer* rules restart an existing timer given a periodicity p or clear it, respectively, and the *setGlobal* rule updates

the value of a global variable. *Events* (κ) are used by entry and exit events for states. The list rules for actions ($\alpha*$) and events ($\kappa*$) are not presented but are straightforward: they apply the relevant rule to the head of the list and pass the updated configuration to the remainder of the list, or return the configuration unchanged for `nil`. *Transition event lists* ($\mu*$) cause the system to change state, executing (only) the first transition from the list whose guard e evaluates to true. Finally, the top-level rule $\vdash_\iota$ describes how the system evolves according to external stimuli.

A *run* of the machine is any valid sequence of configurations produced by repeated applications of the $\vdash_\iota$ rule. Timeout inputs increment the time to the earliest active timeout as described by the *matchesEarliest* predicate:

$$\mathtt{matchesEarliest}(t, x) \equiv \exists t_i, p_i.(p_i, x) = t(t_i) \wedge$$
$$\forall t_j, p_j, y.((p_j, y) = t(t_j) \implies y = \bot \vee y \geq x)$$

The `subtractTimers` function subtracts t_i from each timer in C, and the `clearTimers` function, for any timers whose time remaining is equal to zero, calls the `clearTimer` action[4].

3.1   Well-formedness Properties

To find common issues with detector models, we surveyed (i) detector models across customer tickets submitted to AWS IoT Events, (ii) questions posted on internal forums like the AWS re:Post forum [15], and (iii) feedback submitted via the web-based console for AWS IoT Events. Based on this survey, we determined that the following correctness properties should hold over all detector models. For more details about this survey, please refer to Appendix A.

The model does not contain type errors: The AWS IoT Events expression language is *untyped*, and thus, may contain ill-typed expressions, e.g., performing arithmetic operations on Booleans. A large class of such bugs can be readily detected and prevented using a *type inference* algorithm. The algorithm follows the standard Hindley-Milner type unification approach [16–18] and generates (and solves) a set of type constraints or reports an error if no valid typing is possible. We use this type inference algorithm to detect type errors in the detector model. Every type error is reported as a warning to the customer. When our type inference successfully infers types for expressions, we use them to construct a well-typed abstract state machine using the formalization reported in Section 3.

For the remaining well-formedness properties we use model checking. We introduce one or more *indicator variables* in our global abstract state to track certain kinds of updates in the state machine, and then we assert temporal properties on these indicator variables. Because we use a model checker that

---

[4] In the interests of space, we do not cover the *batch* execution mode, where all variables used in expressions maintain their "pre-state" value until the step is completed; it is a straightforward extension.

checks only *safety properties*, in many cases we invert the property of interest and check that its negation is falsifiable, using the same mechanism often used for test-case generation [19].

Every Detector Model State is Reachable and Every Detector Model Transition and Event can be Executed: For each state $s \in \mathbf{S}$, we add a new Boolean *reachability indicator* variable v_{reached}^s to our abstract state that is initially `false` and assigned `true` when the state is entered (similarly for transitions and events). To encode the property in a safety property checker, we encode the following *unreachability* property expressed in LTL and check it is falsifiable. If it is provable, the tool warns the user.

$$\mathrm{Unreachable}(s) \triangleq \Box \ (\neg \ v_{\mathrm{reached}}^s)$$

Every Variable is Set Before Use: In order to test that variables are properly initialized, first we identify the places where variables are assigned and used. In detector models, there are three places where variables are used: in the evaluation of conditions for events and transitions, and in the `setGlobal` action (which occurs because of an event or transition). We want to demonstrate that the variables used within these contexts are never equal to $\bot$ during evaluation. In this case, we can reuse the reachability variables that we have created for events and transitions to encode that variables should always have defined values when they are used.

We first define some functions to extract the set of variables used in expressions and action lists. The function $Vars(e) : \epsilon \to v$ `set` simply extracts the variables in the expression. For action lists, it is slightly more complex, because variables are both defined and used:

$$Vars(\mathtt{nil}) = \{\}$$
$$Vars(\mathtt{setTimer}(t,e) :: tl) = Vars(e) \cup Vars(tl)$$
$$Vars(\mathtt{resetTimer}(t) :: tl) = Vars(tl)$$
$$Vars(\mathtt{clearTimer}(t) :: tl) = Vars(tl)$$
$$Vars(\mathtt{setGlobal}(g,e) :: tl) = Vars(e) \cup (Vars(tl) - \{g\})$$
$$Vars(\mathtt{event}(e,al)) = Vars(e) \cup Vars(al)$$
$$Vars(\mathtt{transition}(e,al,s')) = Vars(e) \cup Vars(al)$$

Every event or transition can be executed at most once during a computation step, so we can use the execution indicator variables to determine when a variable might be used.

$$\forall a_i, v_j \in Vars(a_i) \ .$$
$$\mathrm{SetBeforeUse}(a_i, v_j) \triangleq \Box \, (v_{\mathrm{exec}}^{a_i} \implies v_j \neq \bot)$$

Input Read Only on Message Trigger: This property is covered in the previous property, with one small change. To enforce it, we modify the translation of the semantics slightly so that at the beginning of each step, prior to processing the input message, all input variables are assigned $\bot$.

Message Triggered Between Consecutive Timeouts: We conservatively approximate a liveness property (no infinite path consisting of only timeout events) with a safety property: the same timer should not timeout twice without an input message occurring in between the timeouts. This formulation may flag models that do not have infinite paths with no input events, but our customers consider it a reasonable indicator.

We begin by defining an indicator variable for each timer t_i (of type integer rather than Boolean): v^i_{timeout} and initialize it to zero. We modify the translation of `updateTimers` to increment this variable when its timer variable equals zero, and modify the translation of the *message* rule to reset all v^i_{timeout} variables to zero. The property of interest is then:

$$\text{NoConsecutiveTimeouts}(t_i) \triangleq \square \ \left(v^i_{\text{timeout}} < 2 \right)$$

4   Experiments

In this section, we evaluate the performance of model-checking safety properties on detector models, with a focus on model checking latency. Low analysis latency is crucial because our tool warns customers of property violations while they are editing their detector model. Our type inference implementation runs with an average latency of 10 milliseconds on all the detector models in our experiments. Since type inference is much faster than model checking and can be successfully run on all detector models, we do not evaluate it in this section.

AWS IoT Events has a commercial feature [20] which uses the type checking and model checking described in Section 3. The feature's implementation first infers types using the type inference algorithm. Next, it translates the detector model into the Lustre language [21]. The translation of IoT Events into Lustre is straightforward and directly follows from the semantics presented in Section 3. The safety properties described in Section 3.1 are attached to the model, along with location information. Then the feature analyzes the model using the JKind [1] tool suite, an open-source industrial model-checker. If JKind invalidates a safety property, the feature decodes the location from the safety property and includes it in the warning.

To evaluate this implementation, we randomly selected 210 detector models previously analyzed by the commercial feature. We checked the five properties described in Section 3.1 in parallel on a c4.8xlarge EC2 instance running Amazon Linux 2 x86_64 using JKind version 4.4.1, with a timeout of 60 seconds.

Of the safety properties that we were able to translate to Lustre, JKind resolved 96% within our timeout of 60 seconds, with 80% completing in less than 10 seconds.

Table 1 shows that checking the *no-unreachable-action* safety property requires the most time to complete. The detector models analyzed in the evaluation include models for monitoring self-driving wheel chairs, monitoring device connectivity, humidity, temperature, pressure, oil level, oil temperature, doors, motion, refrigerator temperature, dough fermentation, and vehicle speed-sensing. They consisted of between 1-7 states and from 0-14 state changes. The

Table 1: Performance of our model-checking tool against 210 detector models

safety property	avg. latency (milliseconds)	# completed	# translation failed	# timeout
no-unreachable-state	3544	176	28	6
no-unreachable-action	5586	171	28	11
var-always-set-before-use	2968	179	28	3
no-infinite-timer-expiration	2875	174	28	8
no-input-read-with-timer-expiration	5477	177	30	3

no-unreachable-action safety property is checked on every action, generating an average of 17 safety properties per detector model, the most of any kind of safety property. This large number of properties to be checked on every detector model caused checking the *no-unreachable-action* safety property to have the highest average latency (5.6 seconds per analysis).

Table 1 shows that about 13% of the properties could not be translated to Lustre. In 2% of the detector models, translation failures arose due to type errors or incorrect use of the AWS IoT Events expression language in the detector model. The remaining translation failures occurred due to either: (1) use of operations not supported by Lustre, (2) no types being inferred for inputs or variables in the detector model, or (3) use of non-linear arithmetic, which is unsupported in JKind. Bitwise functions, strings, and array data types are supported in the AWS IoT Events expression language but not in Lustre. This language gap prevented us from translating 19 of the 210 detector models. Failing to infer a type for a variable in the detector model prevented translation of 6 of the 210 detector models. JKind's lack of support for non-linear arithmetic prevented model-checking 2 of the 210 detector models. We are actively working to support more functions, string and array data types, type annotations, and non-linear arithmetic in our model-checking of detector models.

5   Conclusion

Our analyzers have been running in the AWS IoT Events production service since December 2021. Since then, 93% of AWS IoT Events customers have used our implementation to check their detector models for well-formedness, without needing to have any knowledge of the underlying type checking and model checking. Our analyzers successfully complete for 85% of real-world detector models and we are actively working on improving this support as explained in Section 4. Overall, our implementation has reported well-formedness property violations in 22% of submitted detector models in the production service, with an average latency of 5.6 seconds. We find giving customers push-button access to fast verification without requiring any knowledge of the underlying techniques enables adoption of automated reasoning-based tools.

References

1. A. Gacek, J. Backes, M. Whalen, L. Wagner, and E. Ghassabani, "The JKind model checker," in *Computer Aided Verification*, H. Chockler and G. Weissenbacher, Eds. Cham: Springer International Publishing, 2018, pp. 20–27.

2. C. Y. Cho, V. D'Silva, and D. Song, "Blitz: Compositional bounded model checking for real-world programs," in *2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2013, pp. 136–146.

3. Z. Baranová, J. Barnat, K. Kejstová, T. Kučera, H. Lauko, J. Mrázek, P. Ročkai, and V. Štill, "Model checking of C and C++ with DIVINE 4," in *Automated Technology for Verification and Analysis*, D. D'Souza and K. Narayan Kumar, Eds. Cham: Springer International Publishing, 2017, pp. 201–207.

4. A. Gargantini and C. Heitmeyer, "Using model checking to generate tests from requirements specifications," in *Software Engineering — ESEC/FSE '99*, O. Nierstrasz and M. Lemoine, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 146–162.

5. C. Karamanolis, D. Giannakopoulou, J. Magee, and S. Wheater, "Model checking of workflow schemas," in *Proceedings Fourth International Enterprise Distributed Objects Computing Conference. EDOC2000*, 2000, pp. 170–179.

6. E. M. Clarke, T. A. Henzinger, H. Veith, R. Bloem *et al.*, *Handbook of model checking.* Springer, 2018, vol. 10.

7. E. M. Clarke Jr, O. Grumberg, D. Kroening, D. Peled, and H. Veith, "Model checking. cyber physical systems series," 2018.

8. A. R. Bradley, "Incremental, inductive model checking," in *2013 20th International Symposium on Temporal Representation and Reasoning*, 2013, pp. 5–6.

9. L. de Moura, H. Rueß, and M. Sorea, "Bounded model checking and induction: From refutation to verification," in *Computer Aided Verification*, W. A. Hunt and F. Somenzi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 14–26.

10. J. L. Newcomb, S. Chandra, J.-B. Jeannin, C. Schlesinger, and M. Sridharan, "IOTA: A calculus for internet of things automation," in *Proceedings of the 2017 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, ser. Onward! 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 119–133. [Online]. Available: https://doi.org/10.1145/3133850.3133860

11. C.-J. M. Liang, B. F. Karlsson, N. D. Lane, F. Zhao, J. Zhang, Z. Pan, Z. Li, and Y. Yu, "SIFT: Building an internet of safe things," in *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, ser. IPSN '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 298–309. [Online]. Available: https://doi.org/10.1145/2737095.2737115

12. M. García-Herranz del Olmo, P. A. Haya, and X. Alamán, "Towards a ubiquitous end-user programming system for smart spaces," *Journal of Universal Computer Science*, 2010.

13. B. Ur, E. McManus, M. Pak Yong Ho, and M. L. Littman, "Practical trigger-action programming in the smart home," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 803–812. [Online]. Available: https://doi.org/10.1145/2556288.2557420

14. R. Alur and D. L. Dill, "A theory of timed automata," *Theoretical Computer Science*, vol. 126, pp. 183–235, 1994.

15. Amazon Web Services, Inc., "Find answers to AWS questions about AWS IoT Events — AWS," https://repost.aws/tags/TANsxSwnClQ_Wfh-uklXi7hQ, 2021.
16. L. Damas and R. Milner, "Principal type-schemes for functional programs," in *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL '82. New York, NY, USA: Association for Computing Machinery, 1982, p. 207–212. [Online]. Available: https://doi.org/10.1145/582153.582176
17. R. Milner, "A theory of type polymorphism in programming," *Journal of Computer and System Sciences*, vol. 17, no. 3, pp. 348–375, 1978. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0022000078900144
18. R. Hindley, "The principal type-scheme of an object in combinatory logic," *Transactions of the American Mathematical Society*, vol. 146, pp. 29–60, 1969. [Online]. Available: http://www.jstor.org/stable/1995158
19. G. Gay, M. Staats, M. Whalen, and M. P. E. Heimdahl, "The risks of coverage-directed test case generation," *IEEE Transactions on Software Engineering*, vol. 41, no. 8, pp. 803–819, 2015.
20. AWS IoT Events, "Troubleshooting a detector model by running analyses," https://docs.aws.amazon.com/iotevents/latest/developerguide/iotevents-analyze-api.html, 2021.
21. N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud, "The synchronous data flow programming language LUSTRE," *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1305–1320, 1991.

A   Common Issues with Detector models

Table 2: Issues seen in detector models from customer questions

#	Issue	# of instances
1	incorrectly scaling detector model	1
2	unreachable action	2
3	infinite loop	3
4	variable-used-before-set	3
5	input read on timer expiration	3
6	insufficient logging permissions	3
7	incorrectly typed expression	5
8	incorrect cross-service setup	8
9	missing simplifications	8
		36

As mentioned in Section 3.1, we surveyed customer detector models for generic correctness problems. We present the root causes of the problems from this study in Table 2. Incorrect scaling (#1) occurs when the customer does not set up their detector model to be instantiated correctly for every IoT device in their fleet. Infinite loop (#3) occurs when the detector model has an infinite execution path involving only timeout events and no external input messages. IoT models should be eventually quiescent if no external inputs occur.

Variable-used-before-set (#4) occurs when a variable in the detector's state is read from before being set to an initial value. AWS IoT Events does not require variables in detector models to be initialized.

A step through a detector can be triggered due to both a timer expiration or a new value being sent to the detector by the outside world. Input read on timer expiration (#5) occurs when a step, triggered by timer expiration, causes the detector to read from its input(s). This is a problem because customers often do not realize that such a read will return the last value sent to the detector by the outside world. Insufficient logging permissions (#6) occurs when a detector is not given sufficient permissions to produce logging output. Incorrect cross-service setup (#8) occurs when customers do not correctly set up data flow across services in AWS IoT. While unnecessarily complex detector models (#9) is not a correctness problem, it poses a significant difficulty to customers in maintaining their detector models, and so, we include it in Table 2.

Of these 9 root causes, we identified that type checking and model checking detected 5 root causes highlighted in green in Table 2. These 5 root causes were responsible for 44% of issues in our survey. Based on Table 2, we determined that the following correctness properties should hold over all detector models:

1. *Detector models must be well-typed*
2. *Every detector model state must be reachable*
3. *Every detector model action must be executable*
4. *Every variable must be set before being used*
5. *Input reads shall not happen on timer expiration*
6. *Detector model must not have infinite timer expirations*

We explain these properties further s in Section 3.1.