

MAVERICK: An App-independent and Platform-agnostic Approach to Enforce Policies in IoT Systems at Runtime

M. Hammad Mazhar
Department of Computer Science
The University of Iowa
Iowa City, IA, USA
muhammadhammad-
mazhar@uiowa.edu

Li Li
Endadul Hoque
Department of EECS
Syracuse University
Syracuse, NY, USA
lli101@syr.edu
enhoque@syr.edu

Omar Chowdhury
Department of Computer Science
Stony Brook University
Stony Brook, NY, USA
omar@cs.stonybrook.edu

ABSTRACT

Many solutions have been proposed to curb unexpected behavior of automation apps installed on programmable IoT platforms by enforcing safety policies at runtime. However, all prior work addresses a weaker version of the actual problem due to a simpler, unrealistic threat model. These solutions are not general enough as they are heavily dependent on the installed apps and catered to specific IoT platforms. Here, we address a stronger version of the problem via a realistic threat model, where (i) undesired cyber actions can come from not only automation platform backends (e.g., SmartThings) but also close-sourced third-party services (e.g., IFTTT), and (ii) physical actions (e.g., user interactions) on devices can move the IoT system to an undesirable state. We propose a runtime mechanism, dubbed MAVERICK, which employs an app-independent, platform-agnostic mediator to enforce policies against all undesired cyber actions and applies corrective-actions to bring the IoT system back to a safe state from an unsafe state transition. MAVERICK is equipped with a policy language capable of expressing rich temporal invariants and an automated toolchain that includes a policy synthesizer and a policy analyzer for user assistance. We implemented MAVERICK in a prototype and showed its efficacy in both physical and virtual testbeds, incurring minimal overhead.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Security and privacy** → **Logic and verification**; **Access control**; **Security requirements**.

KEYWORDS

IoT systems, policy enforcement, policy analysis, policy synthesis

ACM Reference Format:

M. Hammad Mazhar, Li Li, Endadul Hoque, and Omar Chowdhury. 2023. MAVERICK: An App-independent and Platform-agnostic Approach to Enforce Policies in IoT Systems at Runtime. In *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '23)*, May 29-June 1, 2023, Guildford, United Kingdom. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3558482.3590188>



This work is licensed under a Creative Commons Attribution International 4.0 License.

WiSec '23, May 29-June 1, 2023, Guildford, United Kingdom
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9859-6/23/05.
<https://doi.org/10.1145/3558482.3590188>

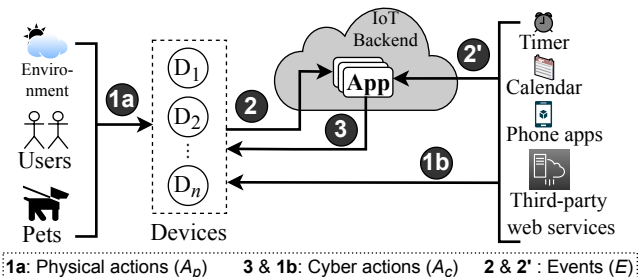


Figure 1: A modern programmable IoT system. D_1, \dots, D_n are IoT devices, **1a** are physical interactions, **1b** are cyber actions sent directly to devices, **2** and **2'** are events from IoT devices and from other services respectively, and **3** are cyber actions generated by apps on an IoT backend.

1 INTRODUCTION

Modern *programmable IoT platforms* (e.g., SmartThings, OpenHAB, IFTTT) offer a cost-effective automation platform to install customized applications (aka, *apps*), which in most cases require no explicit support from device (e.g., smart lock) manufacturers. Despite tremendous innovation, safety and security concerns in programmable IoT systems are still a pressing issue [66]. One source of this issue revolves around the automation itself. For instance, many apps can inconceivably interact with installed devices leading to a disastrous outcome, which can compound due to a chain-reaction in a smart home with multiple apps. Many solutions have been proposed to curb unexpected behavior of apps installed on IoT platforms [5, 11, 12, 15, 18, 22, 36, 38, 69, 70]. One class of prior work hinges on the enforcement policies at runtime [12, 18, 36, 38, 69, 70] – the focus of this paper.

For IoT systems, a policy expresses what aspects of the system must not be violated by the current global state of the system (aka, *system-state*, composed of the current internal state of each installed device). An *unsafe* state of the system is a system-state that violates at least one policy at runtime. During an unsafe state, the system's safety can be compromised, which can result in dire consequences. Any change in the system-state is due to a change in a device's internal state, which is in turn due to an *action/command* executed on the device (e.g., lock/unlock). Additionally, any change in a device's state results in an *event* (e.g., locked/unlocked), possibly triggering execution of one or more automation apps (**2** in Figure 1).

Some actions can be undesired in the current system-state if they can induce a system transition into an unsafe state (e.g., unlocking

the front door when nobody is at home). To ensure that the system obeys the given policies at runtime, one needs to prevent undesired actions from reaching the devices. Figure 1 shows the sources of undesired actions in a modern IoT system: *physical actions* (\mathcal{A}_p) via user interactions or environmental changes (1a) and *cyber actions* (\mathcal{A}_c) sent to devices through network messages (3 and 1b).

To the best of our knowledge, most prior work [12, 18, 36, 69, 70] assumes a weaker threat model that considers apps installed in the IoT platform's native backend (3 in Figure 1) as the only source of undesired actions. Technically, they focus on solving whether installed apps violate given policies – a *weaker version of the actual problem*. These proposals are heavily app-dependent (e.g., require source code) and thus platform-specific. Most safely assume 1a and 1b actions can never transition the system into an unsafe state, which is *unrealistic in practice*. If the system ever transitions into an unsafe state due to 1a or 1b, these solutions will either be stuck in limbo [69, 70] or inaccurately enforce policies [12].

In this paper, we tackle the *stronger version of the actual problem* where we want to ensure that the deployed IoT system satisfies the user-specified policies at runtime, independent of IoT backend apps and third party services. We expand the threat model where the sources of undesired actions include all possible cyber and physical actions (1a, 1b, and 3 in Figure 1). Using this threat model, we formally define the stronger version of the problem as the *IoT System Safety Problem* (ISSP) for a general IoT system. ISSP demands a runtime mechanism that ensures that (R1) a cyber action that transitions the IoT system into an unsafe state is prevented from the reaching the target device, and (R2) if the system ever transitions into an unsafe state due to a physical action (or a change in the environment), the mechanism can take appropriate actions to bring the system back to a safe state. We propose a runtime approach, dubbed **MAVERICK**, that ensures both R1 and R2, while being independent of installed apps and IoT platforms.

To ensure R1 under limited access to apps' source code, a solution must be able to mediate all cyber actions (1b and 3). MAVERICK tackles this issue by employing a **platform-agnostic mediator** which mediates all communication between apps and IoT devices (i.e., 2, 3, and 1b in Figure 1) and enforces user-specified policies at runtime (§ 4). Unlike prior policy enforcement work, we deliberately design this mediator to operate as an individual component (similar to PFWall[14]) independent of the platform's backend and app engine, third-party services, and installed apps. The mediator is equipped with the necessary communication interfaces to act as a man-in-the-middle, but *trusted intermediary*, that intercepts all messages between IoT devices and its backend or third-party services. For reliable runtime policy enforcement, the mediator also maintains a *shadow* system-state in-sync with IoT device state by monitoring state update messages. For R2, MAVERICK redefines the concept of **corrective actions**. A corrective action is technically a cyber action to be sent to an IoT device or a drop of a violating cyber action. Whenever MAVERICK finds that the IoT system is in an unsafe state and violates a policy (say P), MAVERICK applies a sequence of corrective actions carefully chosen for P so that the sequence can nudge the system toward a safe state that satisfies P .

Often policies for a modern IoT system cannot be enforced reliably without knowledge of system execution history capturing

temporal ordering of events and actions. Some policies are also contingent on environment/system variables. Unlike prior work, we develop a *policy language* (§ 5) that can express conditions on temporal ordering of past events/actions, rich attributes of IoT devices and their relationships, and custom predicates on system variables. MAVERICK expects the user to specify the desired policies in this language. For each policy, the user also specifies the corresponding corrective actions.

Asking users to write policy in a dedicated policy language can impede the adoption of MAVERICK in practice. To address this, we augment our policy language toolchain with two automated components for policy synthesis (§ 5.2): *invariant synthesizer* and *policy analyzer*. The invariant synthesizer aims to lift the burden of the user by generating invariants (i.e., fine-grained components of a policy) based on a user-provided set of positive and negative execution traces of the system such that the positive traces satisfy the invariants but the negative ones violate them. The policy analyzer checks the effectiveness of user-selected corrective actions.

We also prove that the decision version of ISSP is *undecidable* in general. In other words, once the system transitions to an unsafe state, whether we can nudge the system back to a safe state using corrective actions is undecidable. Despite this undecidability, it is possible to develop a *sound* solution for the ISSP decision problem under reasonable assumptions (§ 5). To evaluate if user-provided corrective actions are sufficient to bring the system back to a safe state, we utilize our policy analyzer, reducing this problem to a model checking problem (§ 5.2).

We implement a prototype of MAVERICK on top of a Mosquito MQTT server [46], which communicates with different IoT platform backends and third-party services. Our prototype leverages OpenHAB to enable support for numerous WPAN technologies (e.g., Z-wave, ZigBee) for IoT device communication. The mediator of our prototype translates messages between the MQTT network and the WPAN (e.g., ZigBee) and enforces policies or applies corrective actions as required at runtime. We have empirically evaluated our prototype in both virtual and physical testbeds with system execution examples/scenarios drawn from an existing dataset [43]. Our evaluation reports that MAVERICK is highly effective while incurring minimal overhead on latency and throughput.

Contributions. This paper has the following contributions:

- (1) We introduce the IoT system safety problem (ISSP) which captures a realistic threat model. We show that the decision problem version of ISSP is undecidable in general; however, we develop a sound algorithm under some assumptions.
- (2) We propose a runtime solution, dubbed MAVERICK, which enforces policies against undesired cyber actions from different sources by employing a platform-agnostic trusted mediator and applies corrective actions to recover the IoT system from an unsafe state.
- (3) To assist users write policies, we develop a policy language to capture temporal invariants and its necessary automated toolchain, including a policy/invariant synthesizer and a policy analyzer.
- (4) We implemented the features of MAVERICK and evaluated all its components on both virtual and physical test-beds using an existing dataset. Our evaluation shows that MAVERICK is effective and incurs a minimal overhead.

2 PRELIMINARIES

IoT Devices. *IoT devices* are the main component of an IoT system. IoT device functionality can be roughly partitioned into two categories: *sensing* and *actuation*. An IoT device uses its sensing capabilities to discern one or more environmental events (e.g., change of temperature, presence). An IoT device uses its *actuation* capabilities to carry out some tasks based on some *cyber* actions (\mathcal{A}_c) or *physical* actions (\mathcal{A}_p) (e.g., sending an on command or pressing the physical switch). An IoT device may be equipped with multiple sensing and actuation functionality, and also maintains one or more states based on such functionality. An IoT device takes one or more events (e.g., environmental change (1a), or a cyber/physical command (1a, 1b or 3)) as inputs, changes its internal state, and finally generates an output event (2) confirming the state update.

IoT Hubs/Gateways. An IoT device may communicate over the Internet directly through Wi-Fi or a wired connection, or utilize WPAN technologies such as Z-Wave [72], Zigbee [21], Bluetooth BLE [29]. WPAN connectivity is facilitated by *device hubs* which then provide IP-based connectivity with a device-specific mobile *companion app* or with *automation platforms*. These automation platforms may be hosted on *local gateways* placed within the home network (e.g. OpenHAB), or hosted on cloud backends and placed remotely (e.g. Samsung SmartThings, IFTTT). Cloud-hosted automation platforms may either communicate to the IoT device itself directly or via its device hub, or utilize a local gateway.

Automation platforms. These allow users to automate the behavior of their IoT devices through the use of *apps*. Automation services provide a centralized interface to monitor and control behavior of devices from different manufacturers, utilizing manufacturer-provided integrations. The automation services can also be categorized into *native* and *third-party* services. In the former, users have fine-grained control over how the installed apps are executed. For example, one can rewrite an installed app to include policy checking hooks, such as in OpenHAB and Samsung SmartThings. Third-party automation services provide a limited programming interface and control over app executions. In those platforms (e.g., IFTTT, Zapier), users can only choose the trigger conditions and the corresponding commands to the IoT devices. Apps on those platforms cannot be equipped with custom policy checking hooks and hence are not amenable to existing runtime approaches [12, 36, 69, 70].

Programmable IoT System. One can view an IoT system as a distributed system with IoT devices and hubs/gateways as individual computational components. Many automation platforms (e.g., IFTTT) provide programming interfaces for customizing system operation without considering low-level implementation details (e.g., event serialization). An IoT system integrated with such platforms is called a *programmable IoT system*, as illustrated in Figure 1. It primarily provides automation to orchestrate IoT devices to perform a set of high-level, complex actions when certain conditions are fulfilled, e.g. automatically opening the curtains, brewing coffee, and starting the toaster when the user wakes up in the morning (See Listing A1 in the appendix). The underlying automation unit of these automation services is an *automation application/app/rule* which can be viewed as a (stateful) event-driven program and often takes the form of a *trigger action rule*. Other platforms also support full-fledged event-driven programs, with sandboxed automation

app execution. When an event reaches the automation platform, it feeds that event to all automation apps registered for the trigger event and then executes them. Any actions generated by the automation apps (e.g., giving the toaster the `switch_on` command) are routed back to the appropriate IoT devices.

3 MOTIVATION AND PROBLEM STATEMENT

We now formally define the programmable IoT system safety problem (ISSP) along with its complexity. We start the section by discussing the threat model, a running example, and limitations of current work in solving ISSP in its entirety.

Threat Model. We consider that a user can install automation apps in various automation services, including, native (e.g., SmartThings) and third-party (e.g., IFTTT) platforms. The source code for such apps are not always accessible or modifiable. These apps can push the IoT system to transition to a unsafe state where some desired policies are violated (e.g., leaving the oven on when the user is away, opening the front door when the user is sleeping). While most prior work [11, 12, 36, 69, 70] considers the apps – installed in the native platform (3 in Figure 1) – as the source of undesired actions, we consider a stronger version where the source of undesired actions consists of native apps 3, third-party apps 1b, and physical interactions 1a. We assume that devices are tamper-proof, report state faithfully and respond faithfully to given commands. Sensor-spoofing and network-level attacks are outside the scope.

With our threat model, prior work [11, 36, 69, 70] is largely ineffective, because they only focus on 3 and rely on apps’ source code instrumentation to place hooks for collecting information and enforcing policies. While Expat [70] and Patriot [69] make the policy decision inside the instrumented apps, IoTGuard [12] and IoTSafe [18] employ a remote server for their policy decision due to the lack of access (i.e., closed-source-ness) of their chosen platform (e.g., SmartThings). Third-party apps (e.g., IFTTT) are mostly closed-source with limited user instrumentation opportunities. One can argue that IFTTT provides its users with much more control of the execution of its apps, which can be leveraged to enforce policies directly inside the apps, as in prior work for 3 through “filter code” that modifies actions or aborts executions [35]. However, it is not the same as intercepting and blocking an undesired action by an external policy enforcer at runtime. While the former needs to be in-built in the source, the latter should be performed outside the IFTTT server at runtime, unbeknownst to IFTTT.

Physical interactions 1a with devices cannot be intercepted by any IoT apps/platforms. It is possible that the system can transition to an unsafe state due a physical interaction with a device (e.g., the user turns off the water valve). Prior work considers no such physical interactions, and hence, if the system ever goes to an unsafe state, their solution does not guarantee any runtime defense. For instance, Expat [70], Patriot [69], IoTGuard [12] will inaccurately allow/block actions if the system is in an unsafe state.

All policies we consider actually capture the user’s expectation from the system. A policy violation indicates that there exists a combined execution of the installed apps (possibly, collected from unvetted sources) with or without physical interactions such that it leads the system to behave unexpectedly, which technically violates the safety of the system. In that sense, all our policies essentially evaluate the safety of the system and any violation that impacts it.

Motivating example. To demonstrate limitations in prior work, we use an example from existing work [12] with a slight deviation. Consider a programmable smart home with a water leakage sensor, a smart water system whose main valve can be switched off using a command, a smart sprinkler system [24], a temperature sensor, and a smoke sensor. Assume the user installs two apps, namely, *fire protection app* (FPA) and *water-leakage protection app* (WPA). The FPA switches on the smart sprinkler system when the smoke sensor senses smoke and the temperature is higher than some threshold suggesting potential fire. The WPA switches off the main water valve whenever the water leakage sensor senses flooding. In case of a fire, the smart sprinkler system will go off due to FPA. After sometime, however, the water leak sensor will sense flooding, and WPA will switch off the main water valve, cutting water to the sprinklers and letting the fire go on unchecked. The ideal response here is to ignore the flooding and douse the fire first, by preventing WPA from switching off the main water valve.

While some prior work [69, 70] can take the ideal measure with appropriate policy, IoTGuard [12] cannot take the ideal measure even with proper policy as its policy enforcement mechanism is based on a reachability analysis, which cannot handle app interactions through physical channels (e.g., air) [18]. We learned this counter-intuitive fact from an attempt to implement it ¹.

Now we extend our example to cover all types of cyber actions (5a and 1b in Figure 1). Suppose the user installed FPA in the smart home's native IoT backend and installed WPA in an third-party service (e.g., IFTTT) as in this app [34], where the user does not access to the WPA's actual source code.² While cyber actions sent by FPA represents 3, cyber actions by WPA represents 1b. Most prior work [12, 18, 69, 70] cannot prevent WPA from switching off the main water valve, because of their inability to mediate cyber actions from apps running on a third-party service.³

Next we incorporate some physical actions (1a in Figure 1) in our example. Assume the user has turned off the main water valve physically (perhaps for maintenance) and forgot to turn it back on; as a result, the system has moved to an unsafe state. Even if FPA turns on the sprinklers, no water will be sprayed because the main water valve is closed. Most of prior work [12, 18, 69, 70] cannot take any measure to recover the IoT system from this unsafe state.

Inadequacy of prior work. Table 1 outlines how prior work [12, 18, 36, 69, 70] closely related to MAVERICK are quite inadequate in terms of multiple perspectives, highly important to the IoT safety problem. Here we only concentrate on some key perspectives, and the rest will be covered later in the paper.

No prior work supports the stronger threat model (F1) that MAVERICK considers to address as outlined earlier. Unlike MAVERICK, the need to instrument apps' source code (F7) impedes their mediation of cyber actions 1b sent by the apps running on third-party services (F3) where the access to an app's source code is not allowed. Much

of existing work is heavily dependent and fully driven by the installed apps (F6). For instance, IoTGuard [12] and IoTSafe [18] rely on app interactions to generate dynamic interaction graphs, used in their policy enforcement step at runtime. These graphs record and manage a shadow copy of the system-state for their outsourced policy checker. A change in a device's state (aka, an event) that does not trigger any installed app, will never be recorded in the generated graph, resulting in a shadow system-state that is out-of-sync with the reality (F8), leading to inaccurate and unreliable policy enforcement even during safe states.

Automation apps dictate what actions are taken when a specific set of events occurs, while policies express invariants the system must not violate. Barring Expat [70] and Patriot [69], existing works (e.g., IoTGuard [12], IoTSafe [18], ContextIoT [36]) cannot establish a clear distinction between apps and policies (F5), and their implemented policy enforcement mechanisms become incomprehensible in many of their own test cases. However, it appears that some of IoTSafe's [18] policies can *inadvertently* recover the system from an unsafe state caused by 1a actions (F4). Further investigation reveals that those policies are much like apps initiating some cyber actions (akin to our corrective actions) rather than policies defining invariants. As IoTGuard's [12] policies can be much more general, policy enforcement can involve the user (F9) to actively resolve conflicts, while ContextIoT [36] always defers to the user when an app contemplates an action of interest.

IoT System Safety Problem (ISSP). A given programmable IoT system I can be viewed as a state machine and can be expressed as a tuple $\langle S, V, D, A, P, R \rangle$. V is the (possibly, infinite) set of variables whose values can be drawn from the (possibly, infinite) set of constants D while making sure the types of variables are respected (e.g., a variable of type real can be assigned only real values, not string constants). S is the (possibly, infinite) set of states where each of which $s \in S$ maps all variables $v \in V$ to a value from D . Suppose $V = \{\text{BedroomLightState}, \dots\}$ and $D = \{\text{lightON}, \text{lightOFF}, \dots\}$, then a possible state $s \in S$ could be $s = [\text{BedroomLightState} \mapsto \text{lightON}, \dots]$. A is a finite set of high level commands that can be issued to the devices and can be decomposed into cyber commands (i.e., A_c) to be issued by automation services and human commands to be issued by the human or environment (i.e., A_e) where $A = A_c \cup A_e$. $P : S \rightarrow \{\text{true}, \text{false}\} \times C$ is a function that takes a state and decides whether that state is safe or not. If the state is not safe, it also returns a sequence of *corrective actions* $C \in C$, each action in C is drawn from the set A_c . $R \subseteq S \times A \times S$ is the left-total transition relation, which given the current state $s \in S$ and a contemplated action $a \in A$, decides which state I moves to next. *The ISSP problem requires that I never moves to an unsafe state defined by a given policy P , and even if it moves to an unsafe state, possibly due to an action in A_e , then it eventually moves back to a safe state according to the given policy P .* As we show below, the decision version of ISSP is undecidable.

THEOREM 1 (DECISION ISSP). *Given a programmable IoT system $I = \langle S, V, D, A, P, R \rangle$ where $P \in \mathbb{P}$ and an unsafe state $s \in S$, deciding whether I can be transitioned to a safe state $s^* \in S$ according to P is an undecidable problem.*

The proof reduces the Turing machine halting problem to the decision version of ISSP (See [44]).

¹The full implementation of IoTGuard [12] is not publicly available

²Installing automation apps in a third-party service is a reasonable assumption with device-specific companion phone apps which can set up such automation. Unlike FPA, WPA will be operated, controlled and monitored through the third-party's own cloud-based backend.

³Prior work like IoTGuard [12] and others [18, 69, 70] cannot actually mediate the cyber actions sent by IFTTT apps, because they rely on the instrumentation of apps' source code to mediate actions and enforce policies, but unfortunately an IFTTT app's source code is not available and thus cannot be instrumented.

F#	Feature	IoTGuard [12]	ExPAT [70]	PatIoT [69]	ContextIoT [36]	IoTSafe [18]	MAVERICK
F1	Stronger threat model						✓
F2	Mediate cyber actions ③ in Figure 1	✓	✓	✓	✓	✓	✓
F3	Mediate cyber actions ①b in Figure 1						✓
F4	Recover from unsafe state due to ①a in Figure 1					⊖	✓
F5	Clear distinction between apps and policies		✓	✓			✓
F6	App-dependent core mechanism	✓			✓	✓	
F7	Require app instrumentation	✓	✓	✓	✓	✓	
F8	Out-of-sync state management	✓			✓	✓	
F9	Require user involvement during policy enforcement	✓			✓		

Table 1: Comparison of MAVERICK with IoTGuard [12], ExPAT [70], PatIoT [69], ContextIoT [36], and IoTSafe [18]. ✓ refers to the feature being supported/available/implemented whereas ⊖ refers to the feature being partially supported/implemented and an empty cell refers to the feature being not required/implemented/supported/applicable.

4 MAVERICK OVERVIEW

We now provide a design overview of MAVERICK, highlighting each component. As discussed in Section 2, it is possible to organize IoT devices, and different automation services and hubs into a variety of IoT system architectures. Providing policy checking in these different IoT system architectures in an app- and platform-agnostic way makes designing MAVERICK challenging. The main design decision is where to place the MAVERICK’s policy checker in an IoT system such that it has a full view of all the cyber actions and state updates without requiring explicit support from device manufacturers or automation service providers. We design MAVERICK as a




Figure 2: Overview of MAVERICK’s design

man-in-the-middle, but trusted intermediary, illustrated in Figure 2. MAVERICK can be viewed as an amalgamation of the following three components: (a) A *local device interface* to which devices connect using a supported communication protocol (e.g., ZigBee, Z-Wave, etc.) and share status updates (aka, events). MAVERICK issues commands (predominantly, relaying commands from automation apps) to connected devices using this interface. (b) An *automation service interface* to which the automation services connect for issuing commands to installed IoT devices. The same interface is used by MAVERICK to forward any events to automation services. (c) A *policy checker*, which receives events from IoT devices and commands from different services and checks for safety violation. In case of safety violation, it issues relevant corrective actions. It also has an internal timer which triggers the policy/invariant checking. The policy checker also maintains the *shadow system state* through events it receives from IoT devices.

Workflow. The user first connects their IoT devices to MAVERICK’s local device interface and registers all automation services (including, IoT platform backends and third-party services) using MAVERICK’s automation service interface. The user then writes policies by interacting with its *invariant synthesizer* and *policy analyzer* (①), which aids users in writing policies in MAVERICK’s policy language. The policy is then handed to the policy checker (②).

Status updates from devices are forwarded to MAVERICK’s device interface (③), which then forwards them to the policy checker. If all policy invariants are satisfied, then device updates are forwarded to registered automation services and the shadow state is updated. If any invariants are violated, the policy checker issues corrective actions via the device interface and updates the shadow state from resulting device state updates. When automation services issue commands for devices through MAVERICK’s automation service interface (④), the policy checker is consulted. If all invariants will be satisfied upon acting on a command, it is forwarded to the target device and the shadow state is updated upon receiving device state updates. Otherwise, commands are dropped.

Any device that can connect with a hub using Bluetooth/BLE, ZigBee, Z-wave, MQTT, and Wi-Fi is supported by MAVERICK. Connecting a device or automation service to the MAVERICK does not require any modifications to the device or automation service. However, devices that connect to automation services through SSL/TLS-protected IP connections are not supported.

5 POLICY LANGUAGE AND AUTHORING

We now describe the policy language and the associated tools we provide to help users write policies for MAVERICK. A policy in MAVERICK is not only used to define whether a given IoT system state is *safe* but also used to provide corrective actions that can nudge the system to move towards a safe state from an unsafe state.

5.1 Policy Language and Policy Checking

A *policy* $\mathcal{P} \in \mathbb{P}$ in MAVERICK is a sequence of *policy rules* φ_j where $j \in \mathbb{N}$. Each policy rule φ_j is of the form $\langle I_j : C_j \rangle$ in which I_j denotes the *invariant* that must always be maintained by the IoT system to be safe whereas C_j is the *sequence of corrective actions* that must be taken if I_j were to be violated. Any state $s \in \mathcal{S}$ of the programmable IoT system \mathcal{I} is considered *unsafe* if and only if it violates at least one of the invariants I_j . The corrective actions component of a policy rule φ_j , C_j , is a sequence and has the form $[a_0, a_1, a_2, \dots, a_{m-1}]$ in which each $a_i \in \mathcal{A}_c$ is a command to an IoT device to change its state. The available commands a_i that can be used in the C_j component of a policy rule depends on the installed IoT devices. The abstract syntax of an invariant I_j is as follows.

$$\langle \text{Invariant} \rangle \quad I ::= \text{If } \Phi \text{ Then } \Psi$$

$$\langle \text{IfCondition} \rangle \quad \Phi ::= p \mid \Phi_1 \text{ and } \Phi_2 \mid \text{not } \Phi \mid \text{true}$$

$$\langle \text{ThenCondition} \rangle \quad \Psi ::= \Phi \mid \Psi_1 \text{ since } \Psi_2 \mid \text{yesterday } \Psi$$

$$\langle \text{Predicate} \rangle \quad p ::= t_1 \text{ \# } t_2 \mid p_1 \text{ and } p_2 \mid \text{not } p \mid \text{true}$$

$$\langle \text{Term} \rangle \quad t ::= x \in \mathcal{V} \mid \langle \text{const.} \rangle \mid \text{func.}(t_1, \dots, t_n)$$

An invariant I takes the form: “If Φ is true Then Ψ must be true”. In an invariant $I \equiv \text{If } \Phi \text{ Then } \Psi$, Φ can be a predicate p , the constant true, or their logical combinations. In addition to the form of Φ , the then-condition Ψ also allows standard past temporal operators (i.e., since and yesterday) found in linear temporal logic (LTL) [55]. This allows for richer invariant expression through consideration of past executed events. One of the atomic components of the invariant language is a predicate p . A predicate can be a relational operator \boxplus (e.g., \leq , \neq , ...) applied to a pair of terms or logical combinations of multiple predicates. A term is a variable x drawn from the set of variables \mathcal{V} , a constant (e.g., “ON”, 2, 4.9), or a function applied to one or more terms (e.g., RoomTemperature + 10).

Example. A user wishes to devise a policy that ensures that their front door is locked whenever they are away from home. The invariant for this policy can be defined as “If Away Then FrontDoorLocked” with predicates defined as:

Away = (HomeMode.status == “Away”)
FrontDoorLocked = (FrontDoorLock.status == “locked”)

This invariant can be violated in two ways: (1) a command to “unlock” the door received from the automation system is allowed to reach the device or (2) the front door lock reports an “unlocked” status when the home mode is reported as away. For (1), the corrective action is to drop the “unlock” command and for (2), the front door is relocked by sending a “lock” command to it. This set of corrective actions is defined as $C = \{\text{drop}(\text{“unlock”}), \text{send}(\text{“lock”})\}$. The resultant policy is defined as $\mathcal{P} = [I, C]$.

Policy checking. MAVERICK’s policy checker gets triggered by one of three conditions: (1) a downstream command is being sent to a device (1b and 3 in Figure 1); (2) a device changed its state and a corresponding state update message is being sent upstream (2 in Figure 1); (3) the global policy checker timer expired. In all cases, the policy checker checks each invariant for violation. If a violation is found, the policy checker triggers the associated corrective actions. The corrective actions are, however, not regulated by the policy checker. The policy checking in condition (1) ensures that MAVERICK covers all possible sources of cyber actions as compared to prior work which depends upon triggering of an instrumented app or platform API call, MAVERICK can preemptively prevent the IoT system from transitioning into an unsafe state by denying the violating commands (i.e., through the special “drop” corrective action). The policy checking for condition (2) is necessary for bringing the system to safe state in the case the system transitioned to an unsafe state due to a physical action or a command issued to the device through a channel that is unobservable by MAVERICK. Policy checking in condition (3) is needed when it takes multiple iterations of corrective actions to navigate the system to a safe state.

MAVERICK’s policy checker takes as input the current state σ_i , the execution history σ , and the policy \mathcal{P} , then it generates a sequence of corrective actions (possibly, empty if all invariants hold). In addition, for condition (1), it also allows the contemplated downstream command to reach the target device. Each invariant in our policy can be represented in a restricted, first-order linear temporal logic formula (i.e., no quantifiers). One can use dynamic programming based approach from the runtime verification literature [8, 19] to check the invariant in an efficient way without having to store the whole execution history of the system (See Appendix A for details).

5.2 Policy Synthesis and Analysis

We now discuss tools we have designed to help users write policies.

Policy Synthesis Workflow. The policy synthesis workflow (1 in Figure 2) starts by the user identifying a set of scenarios that are undesirable (i.e., unsafe) for the IoT system. For each scenario, the user generates one or more diverse, example system executions/traces of changes in IoT device states in which the undesired state did not arise (i.e., *positive examples*). The users then generate a diverse set of example system execution where the undesired state occurred (i.e., *negative examples*). Such traces may vary in terms of event count or precision of the relevant state. The user feeds those example traces into the **automatic invariant synthesizer**, which outputs a set of *candidate invariants* which satisfy each of the positive examples but reject each of the negative examples. The user chooses the appropriate invariant and selects the sequence of corrective actions. Once all the invariants and their corrective actions have been generated, they are fed into the **automatic policy analyzer**, which will check whether the policy has the desired effect. If the policy achieves the desired behavior, then it is deployed.

Automatic Invariant Synthesis. The invariant synthesizer takes as input a set of positive examples, a set of negative examples, and a set of predicates to use in the invariant, then it tries to generate a set of candidate invariants that satisfy the each of the positive examples but violate each of the negative examples. The automatic invariant generation can be viewed as a restricted instance of the *language learning from the informant problem* [16]. For this, we modified an existing tool called SYSLITE [4] to only generate invariants compliant with our policy language, by removing the capability to use the past operators not utilized in our policy language. We use SYSLITE for invariant synthesis for its speed and scalability.

Policy Analyses. The policy analyzer takes as input a policy and then helps the user check whether the policy has the intended behavior by performing the following two classes of analyses: (1) checking invariants; (2) checking the corrective actions.

For class 1 policy analyses, the policy analyzer performs the following consistency checks: (a) does each invariant I_j at least designate one state as safe? (b) does each invariant I_j at least designate one state as unsafe? (c) do all invariants together at least designate one state as safe? (d) do all invariants together at least designate one state as unsafe? We solve these consistency checks by reducing it to a model checking problem in the standard way [61].

For class 2 policy analysis, the policy analyzer checks to see whether corrective actions C_j of each invariant I_j can take the system to a safe state (i.e., all invariants hold) in case I_j gets violated. This requires us to solve the undecidable decision version of the ISSP problem. We reduce the problem as a model checking problem where we use the policy checker as the model and leave the other actions (e.g., commands generated by automation services, human actions) as environment controlled. The guarantee this policy analysis provides is that there exists a path in which the system will *eventually* end up in a safe state given the policy checker takes the necessary corrective actions. In case execution of C_j leads to violation of another invariant I_k , then its set of corrective actions C_k should lead the system to a safe state, for C_j to be accepted.

Policy Analysis Assumptions. To make policy analysis tractable, we make the following assumptions: (i) the devices act faithfully; (ii)

the communication is reliable; (iii) the number of devices are finite; (iv) the number of variables are finite; (v) at most one invariant violation at each state; (vi) there is only one incoming command from automation services or one state update from the device at any point in time. Assumption (v) is only required for tractable policy analysis, especially, to simplify the underlying model checking problem to **not** maintain extra states for storing all pending corrective actions. Our runtime checking allows multiple invariants to be violated. Assumption (vi) is sound due to the fact that the underlying IoT system will *serialize* concurrent events. We leave the ordering of concurrent events as a non-deterministic choice to allow policy analysis to identify orderings leading to unsafe policy.

6 IMPLEMENTATION AND EVALUATION

6.1 MAVERICK Implementation

We develop a MAVERICK prototype by modifying an open source MQTT broker (Mosquitto [46] coded in C) to include the functionality described in Section 4. Mosquitto acts as the *automation service interface* and is responsible for consulting the policy checker whenever a command is received and forwarding commands and status updates from/to the local device interface. We utilize MQTT for its wide support by most automation systems such as OpenHAB, SmartThings and Home Assistant, with an MQTT server acting as a suitable, secure intermediary with platform-agnostic capabilities. The *policy checker* is integrated into Mosquitto, which parses the policies during configuration time, and evaluates them on-demand using the approach in Appendix A.⁴

We utilize OpenHAB v3 [51] as the *local device interface* in our prototype with its support of WPAN communications with IoT devices and its MQTT integration capabilities. We configure OpenHAB to expose all connected IoT devices via MQTT to Mosquitto, which then maintains a synchronized shadow state by monitoring all device status updates sent via the device interface.

6.2 Data

We use the dataset released with Helion [43] to evaluate MAVERICK under realistic settings. Helion utilizes *statistical language modeling* to learn patterns in home automation events to build a statistical language model which is then able to generate home automation event sequences that can be considered *natural* by human observers. Helion can generate both benign (*up*) event sequences representing typical IoT system operation and malfunctioning (*down*) event sequences where the system is behaving unexpectedly.

Helion event sequences consist of event *tokens* of the form $\langle \text{device}, \text{capability}, \text{value} \rangle$, representing an IoT *device* undertaking an action resulting in *value* based upon its specific *capability*. Depending upon the capability (e.g. sensor, actuator), an event can occur exclusively from the device side or can occur from either direction as a cyber/physical action. We ensure that Helion events are replayed from the relevant direction in our testbeds.

6.3 Experimental Setup

Deployment: Figure 3 illustrates our deployment architecture. We deploy our prototype on a Raspberry Pi 4 Model B [57] running the Linux-based Raspbian OS as a local *proxy* in the IoT environment,

with all devices communicating to the programmable IoT platform through it. We use 2 automation platforms – OpenHAB and SmartThings – for our deployment. OpenHAB is running on a separate Raspberry Pi with a stable version (3.0.7), while we use a Samsung SmartThings Hub (STH-ETH-250) for the SmartThings platform running SmartThings V2. While SmartThings does not have built-in MQTT support, we use a community-developed SmartApp [37] for MQTT capabilities. For third-party services, we select IFTTT and integrate it with the SmartThings backend using the web interface. The proxy and the automation platform’s backend (for OpenHAB) or hub (for SmartThings) are connected to the same local network.




Figure 3: MAVERICK deployment over our testbed with automation platform, event generator and virtual IoT devices.

Testbeds: We prepare a virtual testbed using OpenHAB as the automation platform. We replicate Manandhar *et al.*’s [43] virtual testbed which defined 70 home IoT devices and their capabilities, including devices such as smart locks, switches, security systems. We replicate device behavior virtually through the device interface on our proxy, such as change in status through commands sent by the automation platform or through physical interactions using the OpenHAB REST API. We do not use automation apps directly in our automation platforms for most of our evaluations as MAVERICK’s design does not require it to have prior knowledge of installed apps. To simulate app behavior, we generate automation platform commands using their REST APIs. This testbed allows us to virtually run event sequences generated by Helion.

We also prepare a smaller proof-of-concept physical testbed for SmartThings to demonstrate the platform-agnostic capability of MAVERICK. This testbed includes only a smart lock, a smart outlet and a smart bulb. These devices were easily connected to MAVERICK, due to support provided by our OpenHAB-based device interface. We ensure that devices in our physical testbed do not connect to the SmartThings hub directly so that all device commands and status updates between SmartThings and devices are routed through MAVERICK. IFTTT is integrated with this testbed to evaluate MAVERICK’s ability to intercept commands sent by third-party platforms.

6.4 Research Questions

We evaluate the following research questions in our evaluation.

RQ1: *Can MAVERICK synthesize invariants that match user expectations precisely?* Recall that MAVERICK synthesizes policy *invariants* using positive and negative examples of event sequences. We evaluate this question by providing MAVERICK with two types of event

⁴This modified version is available at <https://github.com/hammadmazhar1/MAVERICK>

ID	Invariant	Corrective Action(s)	Invariant Objective
*11	If User Away, on Vacation or Sleeping Then Door is Locked	(Drop Unlock Command), (Door Lock, unlock)	Physical Security
*12	If User Away, on Vacation or Sleeping Then Gas Stove is off	(Drop Turn on Command), (Gas Stove, off)	Physical Safety
13	If Fire Sprinkler is on Then Water Valve is on	(Drop Shut off Command), (Water Valve, open)	Physical Safety
14	If User Away or on Vacation Then Induction Cooktop is off	(Drop Turn on Command), (Induction Cooktop, off)	Physical Safety, Efficiency
*15	If User Away or on Vacation Then Coffee Maker is off	(Drop Turn on Command), (Coffee Maker, off)	Efficiency
16	If User Away or on Vacation Air Conditioner is off while Away or on Vacation	(Drop Turn on command), (Air Conditioner, off)	Efficiency
17	If Heater is on Then Air Conditioner is off	(Drop Turn on command)	Efficiency
18	If Air Conditioner is on Then Heater is off	(Drop Turn on command)	Efficiency

Table 2: Policy used in evaluation. Each invariant is derived from an issue identified by the researcher in Manadhar *et al.* [43]. Invariants marked with * are used as ground truth to measure effectiveness of policy synthesis. Corrective actions are defined by authors when configuring MAVERICK for policy enforcement.

sequences; (1) *natural* event sequences generated by Helion, representing a user providing past examples of positive and negative event sequences and (2) hand-crafted *synthetic* event sequences that only provide targeted positive and negative examples.

RQ2: *How effective is MAVERICK in enforcing user policies in IoT systems?* We replay event sequences generated by Helion on our virtual testbed after synthesizing policies and configuring MAVERICK with them. As part of their evaluation, Manadhar *et al.* [43] had a security expert analyze event sequences to identify safety issues and develop policies to mitigate them. We use the same event sequences to evaluate MAVERICK.

RQ3: *What are the performance characteristics of MAVERICK?* We measure MAVERICK’s performance by measuring its throughput rate, and also the latency it imposes as an intermediary.

6.5 Evaluation Results

RQ1: We select a subset of traces from researcher-vetted *down* traces from the Helion dataset that were identified with the same problem to represent natural negative examples. For instance, in some traces the smart door lock was left unlocked while the user was away or had left on vacation, leading to a physical security issue and recommending invariant I1 in Table 2 to safeguard against it. Natural positive examples were generated by providing Helion with the same histories used for the negative examples, but with Helion set to *up* mode. In contrast, synthetic positive and negative examples were hand-crafted with only the relevant events in sequence.

We identify three sets of sequences of positive/negative examples for the invariants marked in Table 2, and construct the same number of synthetic event sequences. We then provide these sets to MAVERICK for invariant synthesis, asking it to generate 6 candidate invariants, noting differences between the generated and intended invariants. We observe that *event sequence structure affects invariant structure*. Invariants generated via natural event sequences were more specific to the sequences compared to those generated via synthetic sequences. This is likely because natural event sequences are non-uniform and therefore provide a larger sample space to the invariant synthesis algorithm to search through. For natural negative sequences involving invariant I1, multiple events occurred between the event of the door lock unlocking and the user leaving the home, MAVERICK synthesized invariants that required the door lock to be locked n events before the user left, with n dependent on the length of the provided sequences. In contrast, with synthetic negative sequences where the door being unlocked and the user leaving had very little variance in distance, MAVERICK was able to derive a more generalized invariant which only required the door

to be locked *when* the user left. In subsequent experiments, we use 1 invariant from the 6 candidate invariants generated for each set of synthetic event sequences, totaling 3 invariants marked in Table 2. All other invariants were specified by the authors themselves.

We observe that MAVERICK is able to synthesize better *candidate invariants* when the user provides event sequences that precisely describe their positive and negative experiences while excluding other irrelevant events. We also observe that more diverse set of negative examples result in more generalized invariants. While more natural event sequences may be easier to provide through identifying periods in the system’s execution history when the experiences occur, synthesizing policies using more variable-length event sequences is more difficult and can even lead to user confusion with the suggested policies.

Scenario	Event Sequence	Targeted Invariant
*Physical	Door Lock is left unlocked when user leaves.	I1
Security	Unlock Door Lock while user is away	
Physical	Turn on Gas Stove while user is away	I2
Safety	Fire Sprinkler turns on while Water valve is closed	I3
Efficiency	Coffee Maker is turned on while user is on vacation	I5
	Turn on Heater while Air Conditioner is on	I8

Table 3: Scenario-specific event sequences selected for evaluation, with the invariant they target. Scenarios marked with * were also evaluated on the physical testbed

RQ2: We select two types of sequences from the Helion dataset; (1) *scenario-specific* sequences that are marked with an unsafe sequence of events by the dataset, and (2) *longitudinal* sequences with longer event sequences generated using Helion. We initialize the virtual testbed before replaying any event sequence to ensure consistency.

6.5.1 Scenario-specific sequences. Scenario-specific sequences target the set of invariants synthesized in Table 2, with relevant events highlighted in Table 3. These invariants and their associated corrective actions are fed to our policy analyzer before configuration to ensure that they will lead the system to a safe state. We use these sequences to evaluate the effectiveness of MAVERICK in enforcing safety policies when the IoT system attempts to transition into an unsafe state. We note MAVERICK’s behavior during each sequence to observe how it evaluates the IoT system state and reacts to unsafe states. We observe that MAVERICK *is effective at maintaining the IoT system in a safe state when commands from the automation system could lead the system to an unsafe state*. For instance, in the *physical safety* scenario the automation system sends a command to turn on the gas stove while the user is away. MAVERICK evaluates the

effect of the command on the current state to derive the future state which is then evaluated with the configured policy for satisfaction. Since the resulting state does not satisfy invariant I2, MAVERICK undertakes the associated corrective action *i.e.* block the command to turn the gas stove, ensuring a safe state for the IoT system.

We observe that MAVERICK *can rectify the IoT system after it transitions to an unsafe state*. For instance, a physical safety scenario event sequence shuts off a water valve after which the fire sprinkler turns on upon detecting smoke. Upon receiving the state update from the sprinkler device and updating its internal representation of system state, MAVERICK's policy evaluation finds invariant I5 to be violated by the new state. MAVERICK generates a corrective action to open the water valve, allowing water to flow to the sprinkler and preventing an unsafe situation. We observe the same in our physical testbed with the physical security scenario where the smart lock is left unlocked when the user leaves. MAVERICK detects that invariant I1 is violated, and generates an corrective action to lock the door.

6.5.2 Longitudinal sequences. Longitudinal sequences are generated by providing short histories (3 events) to Helion and then having it predict subsequent events (100 predictions per sequence). We generate event sequences in both *up* and *down* flavors and replay them on our virtual testbed to evaluate MAVERICK in an in-the-wild setting with the invariants in Table 2 configured. As these sequences are run, we note the number of times MAVERICK undertakes corrective actions, and inspect the recorded state of the testbed before and after the corrective action for correctness. Table 4 presents the results. We note that across both sequence flavors, MAVERICK initiated corrective actions for 10-14% of events, depending upon the nature of the policy MAVERICK is configured with during the system's operation as well as the events themselves. Inspection of these corrective actions revealed that most of them were generated due to violation of invariant I1, attempting to unlock the door after the user left the home. In general, we note that across the longitudinal event sequences, MAVERICK maintains the IoT system in a safe state as defined by the user's safety policy, allowing only commands that trigger safe state transitions.

Event Flavor	Events	Actions triggered	Correct Action Rate
<i>up (benign)</i>	500	10	100%
<i>down (malfunctioning)</i>	500	70	100%

Table 4: Results of running longitudinal sequences on the virtual testbed with MAVERICK configured.

RQ3: We measure performance overhead for MAVERICK in terms of *added one-way latency* and *message throughput*. We isolate MAVERICK and send MQTT messages through it under different sets of invariants installed. Invariants for this evaluation are written as the simplest form; checking if a particular value (as indicated by a MQTT topic) has been set. Therefore each invariant has an associated MQTT topic on which it applies. We use such invariants to limit the compounding factors introduced by more complex invariants in our measurements. We use the same topics to send messages over an unmodified version of the broker as baseline.

Figure 4a shows the average one way latency incurred by MAVERICK. We note that MAVERICK's latency increases with the number

of invariants. At 1000 invariants, MAVERICK incurs 5 milliseconds of extra latency, and the overall trend indicates that latency will increase linearly with invariant count. Figure 4b shows the average throughput of both MAVERICK and an unmodified MQTT broker, measured as messages per second processed by each. In general, we observe a decrease in MAVERICK's throughput as the number of invariants are increased. MAVERICK also achieves much lower throughput as compared to the unmodified broker. Figure 4c illustrates this by deriving the *throughput reduction rate* caused by MAVERICK as invariants are increased, calculated as the ratio of the unmodified MQTT broker's bandwidth over MAVERICK. With 100 invariants, MAVERICK incurs 5x throughput reduction, which increases to 42x at 1000 invariants.

Across our measurements, we see that MAVERICK incurs minimal *latency* overhead even at relatively high invariant counts. However, it does incur significant reduction in the message bandwidth, which may be of concern in deployments involving devices at an order of magnitude. In such situations, it may be prudent to deploy MAVERICK in a more localized fashion with subsets of local invariants divided across multiple MAVERICK instances instead of a globalized deployment of a central MAVERICK instance. This would reduce the number of invariants a given MAVERICK instance has to monitor at any given time. We also note that this is just a proof-of-concept implementation of MAVERICK and as such would likely benefit from more performance based optimizations. One such optimization would be to pinpoint the subset of configured invariants affected by each device state variable. Upon an update on said variable, only the identified invariants are evaluated for satisfaction, thereby further reducing the number of invariants evaluated at any time.

7 DISCUSSION

Invariant Expressiveness: MAVERICK invariants are expressed in a fragment of first-order temporal logic (FOTL). While it can express *temporal invariants* that refer to past states, it cannot express temporal invariants with explicit time. As an example, MAVERICK cannot currently capture invariants of the form: *if the user has been away for at least an hour, all security cameras must be turned on*. MAVERICK can be easily augmented to express and evaluate invariants of this form by using Metric-time Temporal Logic [6, 7, 52, 69] instead of FOTL. This increase in expressive power, however, comes with a price as it makes both invariant synthesis and policy analysis intractable. Thus, we opt for a policy language that is amenable to automated invariant synthesis and policy analysis.

Usability: In order for MAVERICK to be effective, it has to be usable by a typical user in IoT environments. Deploying MAVERICK involves the following two steps: (1) installing the MAVERICK trusted intermediary in the IoT environment, and (2) configuring invariants and corrective actions. Step (1) further involves setting up the proxy, and configuring devices and the automation system to connect to it. This is more involved than plug-and-play. Future MAVERICK versions will focus on making installation easier.

Configuring MAVERICK also requires writing invariants and corrective actions in a given policy language (Step 2). However, an expressive policy language may not necessarily be usable for users unfamiliar with its semantics, leading to ineffective policies. We address this problem by utilizing automated invariant synthesis and policy analysis techniques to assist users in developing policies.




Figure 4: Throughput and latency characteristics of MAVERICK.

We note that prior work does not provide such approaches, and our work is a first step in attempting to make a usable policy enforcement approach. We also note that improving usability for such systems itself is a big challenge due to the fact that no prior knowledge of user requirements can be assumed at the time of policy specification, as well as the myriad possible system configurations possible. Improved means of policy generation from automatically generated event sequences and development of User Interfaces which are able to guide users through the policy synthesis process would be required to improve MAVERICK usability. We primarily focus on developing MAVERICK's core technology and leave the evaluation of its usability as future work to limit scope, as developing a usable policy specification interface requires a full process of user interface development with user studies and feedback.

8 RELATED WORK

There is a substantial amount of prior work in IoT security that focuses on device security [20, 32, 33, 39, 49, 59, 65], communication and authentication protocols [5, 25, 27, 42, 60, 71, 73], and IoT platforms as well as automation apps [3, 10–13, 15, 17, 22, 31, 36, 40, 47, 56, 58, 64, 67, 69, 70, 74]. Some of them aim to ensure that only policy compliant actions are allowed during a programmable IoT system execution either through static analysis [11, 15, 22, 47, 64, 66, 74] or through runtime monitoring [12, 18, 36, 38, 69, 70].

Static analysis based approaches that try to identify policy violations before system deployment suffer from the following limitations: (a) they suffer from imprecision due to approximate analysis, (b) they do not prescribe any solution to the inexperienced user on how to mitigate any identified issue, (c) they suffer from false positives due to lack of contextual information during analysis, and (d) modifications to the IoT environment (e.g., adding an app) warrant a new analysis. On the contrary, MAVERICK being a runtime approach can avoid these limitations.

Prior runtime approaches aim to dynamically enforce policies, but they have multiple shortcomings: (a) They are specialized for a particular IoT platform (i.e., SmartThings or OpenHAB) and cannot readily be ported to other heterogeneous platforms. (b) They are deployed in the platform's backend (to be precise, as either *instrumented apps* or *instrumented APIs*). Thus, third-party cloud services (like IFTTT, Flow) that use the external web API to directly communicate with IoT devices can potentially bypass these runtime approaches when these instrumented automation apps are never executed; leaving IoT devices at risk. (c) They all expect the IoT system to never end up in an unsafe state by any environmental interference. If the system transitions to an unsafe state ever, these solutions can no longer operate. (d) Some of the proposed policy languages are mostly rudimentary, often platform-specific, and barely

useful to express rich policies, such as temporal invariants with custom predicates. MAVERICK addresses all these shortcomings.

There has been work [41] to provide remedial actions when user expectations are violated. Unlike MAVERICK, they involve analysis of applications to determine when conflicts may occur, and then suggest alternative actions/rules to the user to remove the conflict.

Prior work has developed a run-time mediation of messages sent by IoT devices and automation systems to limit the data sent by IoT devices and thus prevent private information leakage [14]. This work has employed similar methods as ours to mediate platform-device communication, by inserting a trusted intermediary that decides what is allowed to pass between the IoT device and platform. However, they seek to address the *privacy* problem in the IoT ecosystem which is orthogonal to our efforts. Their approach also requires analysis of apps/rules on the automation platform to ensure that restrictions on data release do not impact app operations. MAVERICK, in contrast, is completely app-agnostic.

There have been also efforts on forensics analysis [67, 68] and risk-based analysis [2, 17, 56], verification of IoT device events [53], designing access control [28, 31, 40], repairing or synthesizing automation apps [74], addressing information leakage [10, 23, 62], machine learning-based anomaly detection [1, 9, 26, 30, 45, 48, 50, 54, 63, 75] and fuzzing smartphone companion apps for IoT [13, 58]. These efforts are orthogonal to MAVERICK.

9 CONCLUSION

We presented MAVERICK, which ensures that deployed IoT devices maintain some user-defined state invariants in an app- and platform-agnostic way while taking physical interactions into consideration. We implemented MAVERICK using a modified MQTT broker and OpenHAB and evaluated it on testbeds which showed that it can be easily deployed in IoT environments to maintain the desired state invariants in real-time. Our performance evaluation also showed that while MAVERICK may induce a lower message throughput than standard MQTT, it adds minimal latency to individual messages even when configured with a large number of invariants.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Fareed Arif for his assistance in modification of SYSLITE[4]. The authors would also like to thank the anonymous reviewers for their valuable comments and suggestions and the shepherd for their assistance. This material is based upon work supported by the National Science Foundation under Grant Numbers: CNS 2006556 and CNS 2007512

REFERENCES

- [1] Abbas Acar, Hossein Fereidooni, Tigist Abera, Amit Kumar Sikder, Markus Miettinen, Hidayet Aksu, Mauro Conti, Ahmad-Reza Sadeghi, and Selcuk Uluagac. 2020. Peek-a-boo: I see your smart home activities, even encrypted!. In *ACM WiSec '20*.
- [2] Noga Agmon, Asaf Shabtai, and Rami Puzis. 2019. Deployment optimization of IoT devices through attack graph analysis. In *ACM WiSec '19*.
- [3] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. 2019. SoK: Security evaluation of home-based IoT deployments. In *2019 IEEE Symposium on Security and Privacy (S&P)*. IEEE.
- [4] M. Fareed Arif, Daniel Larraz, Mitziu Echeverria, Andrew Reynolds, Omar Chowdhury, and Cesare Tinelli. 2020. SYSLITE: Syntax-Guided Synthesis of PTL Formulas from Finite Traces. In *2020 Formal Methods in Computer Aided Design (FMCAD)*.
- [5] David Barrera, Ian Molloy, and Heqing Huang. 2018. Standardizing IoT network security policy enforcement. In *Workshop on decentralized IoT security and standards (DISS)*.
- [6] David Basin, Felix Klaedtke, Samuel Müller, and Birgit Pfizmann. 2008. Runtime monitoring of metric first-order temporal properties. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*.
- [7] David Basin, Felix Klaedtke, Samuel Müller, and Eugen Zălinescu. 2015. Monitoring metric first-order temporal properties. *Journal of the ACM (JACM)* (2015).
- [8] David A Basin, Felix Klaedtke, and Eugen Zălinescu. 2017. The MonPoly Monitoring Tool. *RV-CuBES* (2017).
- [9] Bruhadeshwar Bezawada, Maalvika Bachani, Jordan Peterson, Hossein Shirazi, Indrakshi Ray, and Indrajit Ray. 2018. Behavioral fingerprinting of IoT devices. In *Proceedings of the 2018 workshop on attacks and solutions in hardware security*.
- [10] Z. Berkay Celik, Leonardo Babun, Amit Kumar Sikder, Hidayet Aksu, Gang Tan, Patrick McDaniel, and A. Selcuk Uluagac. 2018. Sensitive Information Tracking in Commodity IoT. In *27th USENIX Security Symposium (USENIX Security 18)*.
- [11] Z Berkay Celik, Patrick McDaniel, and Gang Tan. 2018. Soteria: Automated IoT Safety and Security Analysis. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*.
- [12] Z Berkay Celik, Gang Tan, and Patrick D McDaniel. 2019. IoTGuard: Dynamic Enforcement of Security and Safety Policy in Commodity IoT. In *NDSS*.
- [13] Jiongyi Chen, Wenrui Diao, Qingchuan Zhao, Chaoshun Zuo, Zhiqiang Lin, Xiaofeng Wang, Wing Cheong Lau, Menghan Sun, Ronghai Yang, and Kehuan Zhang. 2018. IoTfuzzer: Discovering Memory Corruptions in IoT Through App-based Fuzzing. In *NDSS*.
- [14] Haotian Chi, Qiang Zeng, Xiaojiang Du, and Lannan Luo. 2021. PFirewall: Semantics-Aware Customizable Data Flow Control for Smart Home Privacy Protection. *Network and Distributed System Security Symposium*.
- [15] Haotian Chi, Qiang Zeng, Xiaojiang Du, and Jiaping Yu. 2020. Cross-app interference threats in smart homes: Categorization, detection and handling. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE.
- [16] Colin De la Higuera. 2010. *Grammatical inference: learning automata and grammars*. Cambridge University Press.
- [17] Wenbo Ding and Hongxin Hu. 2018. On the safety of IoT device physical interaction control. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*.
- [18] Wenbo Ding, Hongxin Hu, and Long Cheng. 2021. IoTSafe: Enforcing Safety and Security Policy with Real IoT Physical Interaction Discovery. In *NDSS*.
- [19] Mitziu Echeverria, Zeeshan Ahmed, Bincheng Wang, M Fareed Arif, Syed Rafiul Hussain, and Omar Chowdhury. 2021. PHOENIX: Device-Centric Cellular Network Protocol Monitoring using Runtime Verification. In *Proceedings of the 2021 Network and Distributed Systems Symposium*.
- [20] Sam Edwards and Ioannis Profetis. 2016. Hajime: Analysis of a decentralized internet worm for IoT devices. *Rapidity Networks* (2016).
- [21] Sinem Coleri Ergen. 2004. ZigBee/IEEE 802.15. 4 Summary. *UC Berkeley, September 10, 17 (2004)*, 11.
- [22] E. Fernandes, J. Jung, and A. Prakash. 2016. Security Analysis of Emerging Smart Home Applications. In *2016 IEEE Symposium on Security and Privacy (S&P)*. IEEE.
- [23] Earlene Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato, Mauro Conti, and Atul Prakash. 2016. FlowFence: Practical Data Protection for Emerging IoT Application Frameworks. In *USENIX Security '16*.
- [24] FM Global. 2016. *UNCOMMON SENSORS : Researchers pioneer smart sprinklers, with greater detection accuracy and speed of response*.
- [25] Behrang Fouladi and Sahand Ghanoun. 2013. Honey, I'm Home!, Hacking ZWave Home Automation Systems. *Black Hat USA* (2013).
- [26] Chenglong Fu, Qiang Zeng, and Xiaojiang Du. 2021. HAWatcher: Semantics-aware anomaly detection for appified smart homes. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*.
- [27] Neil Zhenqiang Gong, Altay Ozen, Yu Wu, Xiaoyu Cao, Richard Shin, Dawn Song, Hongxia Jin, and Xuan Bao. 2017. PIANO: Proximity-based user authentication on voice-powered internet-of-things devices. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE.
- [28] Sanket Goutam, William Enck, and Bradley Reaves. 2019. Hestia: simple least privilege network policies for smart homes. In *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*. 215–220.
- [29] Naresh Kumar Gupta. 2016. *Inside Bluetooth low energy*. Artech House.
- [30] Ayyoob Hamza, Hassan Habibi Gharakheili, Theophilus A Benson, and Vijay Sivaraman. 2019. Detecting volumetric attacks on IoT devices via sdn-based monitoring of mud activity. In *Proceedings of the 2019 ACM Symposium on SDN Research*.
- [31] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlene Fernandes, and Blase Ur. 2018. Rethinking Access Control and Authentication for the Home Internet of Things (IoT). In *27th USENIX Security Symposium (USENIX Security 18)*.
- [32] Grant Ho, Derek Leung, Pratyush Mishra, Ashkan Hosseini, Dawn Song, and David Wagner. 2016. Smart locks: Lessons for securing commodity internet of things devices. In *Proceedings of the 11th ACM on Asia conference on computer and communications security*.
- [33] Manos howpublishedkakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztin, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. 2017. Understanding the Mirai botnet. In *26th USENIX security symposium (USENIX Security 17)*.
- [34] IFTTT. 2022. Turn water off if D-Link water sensor detects water.
- [35] IFTTT. 2023. Filter Code - IFTTT Help Center.
- [36] Yunhan Jack Jia, Qi Alfred Chen, Shiqi Wang, Amir Rahmati, Earlene Fernandes, Z. Morley Mao, and Atul Prakash. 2017. ContextIoT: Towards Providing Contextual Integrity to Appified IoT Platforms. In *21st Network and Distributed Security Symposium (NDSS)*.
- [37] St. John Johnson. 2021. *SmartThings MQTT Bridge*.
- [38] Kaushal Kafle, Kirti Jagtap, Mansoor Ahmed-Rengers, Trent Jaeger, and Adwait Nadkarni. 2021. Towards Practical Integrity in the Smart Home with HomeEndorser. <https://doi.org/10.48550/ARXIV.2109.05139>
- [39] Deepak Kumar, Kelly Shen, Benton Case, Deepali Garg, Galina Alperovich, Dmitry Kuznetsov, Rajarshi Gupta, and Zakir Durumeric. 2019. All Things Considered: An Analysis of IoT Devices on Home Networks. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association.
- [40] Sanghak Lee, Jiwon Choi, Jihun Kim, Beumjin Cho, Sangho Lee, Hanjun Kim, and Jong Kim. 2017. FACT: Functionality-centric access control system for IoT programming frameworks. In *Proceedings of the 22nd ACM on Symposium on Access Control Models and Technologies*. ACM.
- [41] Renju Liu, Ziqi Wang, Luis Garcia, and Mani Srivastava. 2019. RemedIoT: Remedial actions for Internet-of-Things conflicts. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*.
- [42] Natasha Lomas. 2015. Critical Flaw IDed In ZigBee Smart Home Devices. <https://techcrunch.com/2015/08/07/critical-flaw-ided-in-zigbee-smart-home-devices/>.
- [43] Sunil Manandhar, Kevin Moran, Kaushal Kafle, Ruhao Tang, Denys Poshyvanyk, and Adwait Nadkarni. 2020. Towards a natural perspective of smart homes for practical security and safety analyses. In *2020 IEEE Symposium on Security and Privacy (S&P)*. IEEE.
- [44] M. Hammad Mazhar, Li Li, Endadul Hoque, and Omar Chowdhury. 2023. MAVERICK: An App-independent and Platform-agnostic Approach to Enforce Policies in IoT Systems at Runtime. [arXiv:2302.01452 \[cs.CR\]](https://arxiv.org/abs/2302.01452)
- [45] Markus Miettinen, Samuel Marchal, Ibbad Hafeez, N Asokan, Ahmad-Reza Sadeghi, and Sasu Tarkoma. 2017. IoT Sentinel: Automated device-type identification for security enforcement in IoT. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE.
- [46] Mosquito. 2022. *Eclipse Mosquitto*.
- [47] Dang Tu Nguyen, Chengyu Song, Zhiyun Qian, Srikanth V Krishnamurthy, Edward JM Colbert, and Patrick McDaniel. 2018. IoTSan: fortifying the safety of IoT systems. In *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*. ACM.
- [48] Thien Duc Nguyen, Samuel Marchal, Markus Miettinen, Hossein Fereidooni, N Asokan, and Ahmad-Reza Sadeghi. 2019. DiOT: A federated self-learning anomaly detection system for IoT. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE.
- [49] Sukhvir Notra, Muhammad Siddiqi, Hassan Habibi Gharakheili, Vijay Sivaraman, and Roksana Boreli. 2014. An experimental study of security and privacy risks with emerging household appliances. In *2014 IEEE Conference on Communications and Network Security (CNS)*. IEEE.
- [50] TJ O'Connor, Reham Mohamed, Markus Miettinen, William Enck, Bradley Reaves, and Ahmad-Reza Sadeghi. 2019. HomeSnitch: behavior transparency and control for smart home IoT devices. In *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*. ACM.
- [51] openHAB. 2019. .
- [52] Muslim Ozgur Ozmen, Xuansong Li, Andrew Chu, Z Berkay Celik, Bardh Hoxha, and Xiangyu Zhang. 2022. Discovering IoT Physical Channel Vulnerabilities. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*.

- [53] Muslum Ozgur Ozmen, Ruoyu Song, Habiba Farrukh, and Z Berkay Celik. 2023. Evasion attacks and defenses on smart home physical event verification. In *Network and Distributed System Security (NDSS)*'23.
- [54] Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. 2015. IoTPOT: analysing the rise of IoT compromises. In *9th USENIX Workshop on Offensive Technologies (WOOT 15)*.
- [55] Amir Pnueli. 1977. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*. IEEE.
- [56] Amir Rahmati, Earleene Fernandes, Kevin Eykholt, and Atul Prakash. 2018. Tyche: A Risk-Based Permission Model for Smart Homes. In *2018 IEEE Cybersecurity Development (SecDev)*. IEEE.
- [57] Raspberry Pi Foundation. 2022. *Teach, Learn and Make with Raspberry Pi*.
- [58] Nilo Redini, Andrea Continella, Dipanjan Das, Giulio De Pasquale, Noah Spahn, Aravind Machiry, Antonio Bianchi, Christopher Kruegel, and Giovanni Vigna. 2021. DIANE: Identifying Fuzzing Triggers in Apps to Generate Under-constrained Inputs for IoT Devices. In *42nd IEEE Symposium on Security and Privacy 2021*.
- [59] Eyal Ronen and Adi Shamir. 2016. Extended functionality attacks on IoT devices: The case of smart lights. In *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*. IEEE.
- [60] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin O'Flynn. 2017. IoT goes nuclear: Creating a ZigBee chain reaction. In *2017 IEEE Symposium on Security and Privacy (S&P)*.
- [61] Kristin Y. Rozier and Moshe Y. Vardi. 2007. LTL Satisfiability Checking. In *Model Checking Software*, Dragan Bošnački and Stefan Edelkamp (Eds.).
- [62] Letian Sha, Fu Xiao, Wei Chen, and Jing Sun. 2018. IIoT-SIDefender: Detecting and defense against the sensitive information leakage in industry IoT. *World Wide Web* (2018).
- [63] Amit Kumar Sikder, Leonardo Babun, Hidayet Aksu, and A Selcuk Uluagac. 2019. Aegis: A Context-aware Security Framework for Smart Home Systems. (2019).
- [64] Yuan Tian, Nan Zhang, Yueh-Hsun Lin, XiaoFeng Wang, Blase Ur, Xianzheng Guo, and Patrick Tague. 2017. SmartAuth: User-Centered Authorization for the Internet of Things. In *26th USENIX Security Symposium (USENIX Security 17)*.
- [65] Blase Ur, Jaeyeon Jung, and Stuart Schechter. 2013. The current state of access control for smart devices in homes. In *Workshop on Home Usable Privacy and Security (HUPS)*. HUPS 2014.
- [66] Qi Wang, Pubali Datta, Wei Yang, Si Liu, Adam Bates, and Carl A. Gunter. 2019. Charting the Attack Surface of Trigger-Action IoT Platforms. In *The 26th ACM Conference on Computer and Communications (CCS 19)* (2019-11).
- [67] Qi Wang, Wajih Ul Hassan, Adam Bates, and Carl Gunter. 2018. Fear and Logging in the Internet of Things. In *ISOC NDSS*.
- [68] Judson Wilson, Riad S Wahby, Henry Corrigan-Gibbs, Dan Boneh, Philip Levis, and Keith Winstein. 2017. Trust but verify: Auditing the secure internet of things. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. ACM.
- [69] Moosa Yahyazadeh, Syed Rafiul Hussain, Endadul Hoque, and Omar Chowdhury. 2020. PatIoT: Policy Assisted Resilient Programmable IoT System. In *International Conference on Runtime Verification*. Springer.
- [70] Moosa Yahyazadeh, Proyash Podder, Endadul Hoque, and Omar Chowdhury. 2019. EXPAT: Expectation-based policy analysis and enforcement for appified smart-home platforms. In *Proceedings of the 24th ACM Symposium on Access Control Models and Technologies*.
- [71] Tianlong Yu, Vyas Sekar, Srinivasan Seshan, Yuvraj Agarwal, and Chenren Xu. 2015. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the internet-of-things. In *ACM Workshop on Hot Topics in Networks*.
- [72] Z-Wave. 2022. *Better and Safer Homes are built on Z-Wave*. Z-Wave Alliance.
- [73] Jiansong Zhang, Zeyu Wang, Zhice Yang, and Qian Zhang. 2017. Proximity based IoT device authentication. In *IEEE INFOCOM 2017*. IEEE.
- [74] Lefan Zhang, Weijia He, Jesse Martinez, Noah Brackenburg, Shan Lu, and Blase Ur. 2019. AutoTap: Synthesizing and Repairing Trigger-Action Programs Using LTL Properties. In *ICSE*.
- [75] Wei Zhang, Yan Meng, Yugeng Liu, Xiaokuan Zhang, Yinqian Zhang, and Haojin Zhu. 2018. Homonit: Monitoring smart home apps from encrypted traffic. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM.

A POLICY CHECKING

MAVERICK's policy checking module conceptually takes as input a policy \mathcal{P} , the system's execution history σ , and the current system state σ_i . It then checks to see whether the current state σ_i satisfies all the invariants of the policy \mathcal{P} .

The main insight that MAVERICK uses for policy checking is hinged upon the observation that any policy written in MAVERICK's

policy language can be represented as a restricted-fragment of first-order logic formula. This allows us to use techniques for runtime verification for efficient policy checking in MAVERICK.

Note that, invariants in \mathcal{P} can have standard past temporal operators such as *yesterday* and *since*. *yesterday* is a unary temporal operator that can be applied to a linear temporal logic formula φ (i.e., written *yesterday* φ), which evaluates to true if and only if the formula φ held in the immediately previous state. Similarly, *since* is a binary temporal operator that can be applied to two linear temporal logic formulas φ_1 and φ_2 (i.e., written φ_1 *yesterday* φ_2), which evaluates to true if and only if the formula φ_2 held in *some* previous state, then φ_1 must hold in every state following the state in which φ_2 held until the current state.

Since invariants in MAVERICK can refer to arbitrary past system states through the use of temporal operators *yesterday* and *since*, it seems natural that one has to keep track of all the previous states. However, as we will show, using insights from runtime verification, one does not need to keep track of the whole execution history σ and can rather cache all previous state concisely into a bit-vector of size σ (i.e., the size of a formula is the number of sub-formulas in it). This is particularly crucial because σ keeps growing as the system executes and would have made policy checking prohibitively expensive. The policy checking complexity of MAVERICK, fortunately, is linear to the size of the invariant, thanks to its dynamic programming algorithm that is presented below.

Suppose the temporal logic formula corresponding to a policy \mathcal{P} is φ . One can use dynamic programming to check whether φ holds in the current position i , denoted by $\llbracket \varphi \rrbracket^i = \text{True}$. Instead of storing the full history of system execution σ , this approach only stores the values of $\llbracket \pi \rrbracket^{i-1}$ and $\llbracket \pi \rrbracket^i$, for all sub-formulas π of φ . Thus, we need two bits of information for each sub-formula of φ to check whether φ holds in the current position i , denoted by $\llbracket \varphi \rrbracket^i = \text{True}$, in the following way.

$$\begin{aligned}
 \llbracket p \rrbracket^i &= p \in \sigma_i \text{ (} p \text{ holds in } \sigma_i \text{)} \\
 \llbracket \varphi_1 \wedge \varphi_2 \rrbracket^i &= \llbracket \varphi_1 \rrbracket^i \wedge \llbracket \varphi_2 \rrbracket^i \\
 \llbracket \neg \varphi \rrbracket^i &= \neg \llbracket \varphi \rrbracket^i \\
 \llbracket \text{yesterday } \varphi \rrbracket^i &= i > 0 \wedge \llbracket \varphi \rrbracket^{i-1} \\
 \llbracket \varphi_1 \text{ since } \varphi_2 \rrbracket^i &= \llbracket \varphi_2 \rrbracket^i \vee (\llbracket \varphi_1 \rrbracket^i \wedge \llbracket \varphi_1 \text{ since } \varphi_2 \rrbracket^{i-1})
 \end{aligned}$$

As an example, *yesterday* φ holds in the current position i of the system execution history, denoted by $\llbracket \text{yesterday } \varphi \rrbracket^i$ if and only if $i > 0$ and φ held in state $i - 1$, denoted by $\llbracket \varphi \rrbracket^{i-1}$. This is what is captured above in MAVERICK's policy checking algorithm presented just above.

```

1 rule "Morning Routine"
2   when
3     Item BedSensor received update OutOfBed
4   then
5     WindowBlinds.sendCommand(UP)
6     CoffeeMachine.sendCommand(Brew)
7     Toaster.sendCommand(Start)
8   end
9

```

Listing A1: An example trigger-action rule in OpenHAB